



US00RE48443E

(19) **United States**
(12) **Reissued Patent**
Drury

(10) **Patent Number: US RE48,443 E**
(45) **Date of Reissued Patent: Feb. 16, 2021**

(54) **SYSTEMS AND METHODS FOR DETERMINING A TOPOLOGY OF AN ETHERNET RING IN A BUILDING MANAGEMENT SYSTEM**

(71) Applicant: **Johnson Controls Technology Company, Auburn Hills, MI (US)**

(72) Inventor: **Eric J. Drury, Lake Zurich, IL (US)**

(73) Assignee: **Johnson Controls Technology Company, Auburn Hills, MI (US)**

(21) Appl. No.: **16/206,614**

(22) Filed: **Nov. 30, 2018**

Related U.S. Patent Documents

Reissue of:

(64) Patent No.: **10,116,517**
Issued: **Oct. 30, 2018**
Appl. No.: **15/583,901**
Filed: **May 1, 2017**

U.S. Applications:

(66) Substitute for application No. 15/583,901, filed on May 1, 2017, now Pat. No. 10,116,517.

(51) **Int. Cl.**
H04L 12/28 (2006.01)
H04L 12/24 (2006.01)
H04L 29/12 (2006.01)
H04L 12/42 (2006.01)

(52) **U.S. Cl.**
CPC **H04L 41/12** (2013.01); **H04L 12/42** (2013.01)

(58) **Field of Classification Search**
CPC **H04L 41/12**; **H04L 12/42**; **H04L 61/6022**
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

8,244,838 B2 * 8/2012 Balasubramanian
G05B 19/0423
370/216
8,520,508 B2 * 8/2013 Zinjuvadia H04L 45/00
370/222
9,037,748 B2 * 5/2015 Nguyen H04L 43/0811
709/242
9,413,552 B2 * 8/2016 Forguites H04L 61/2015
9,450,808 B2 * 9/2016 Lee H04L 12/437
9,537,776 B2 * 1/2017 Alsup H04L 47/13

(Continued)

OTHER PUBLICATIONS

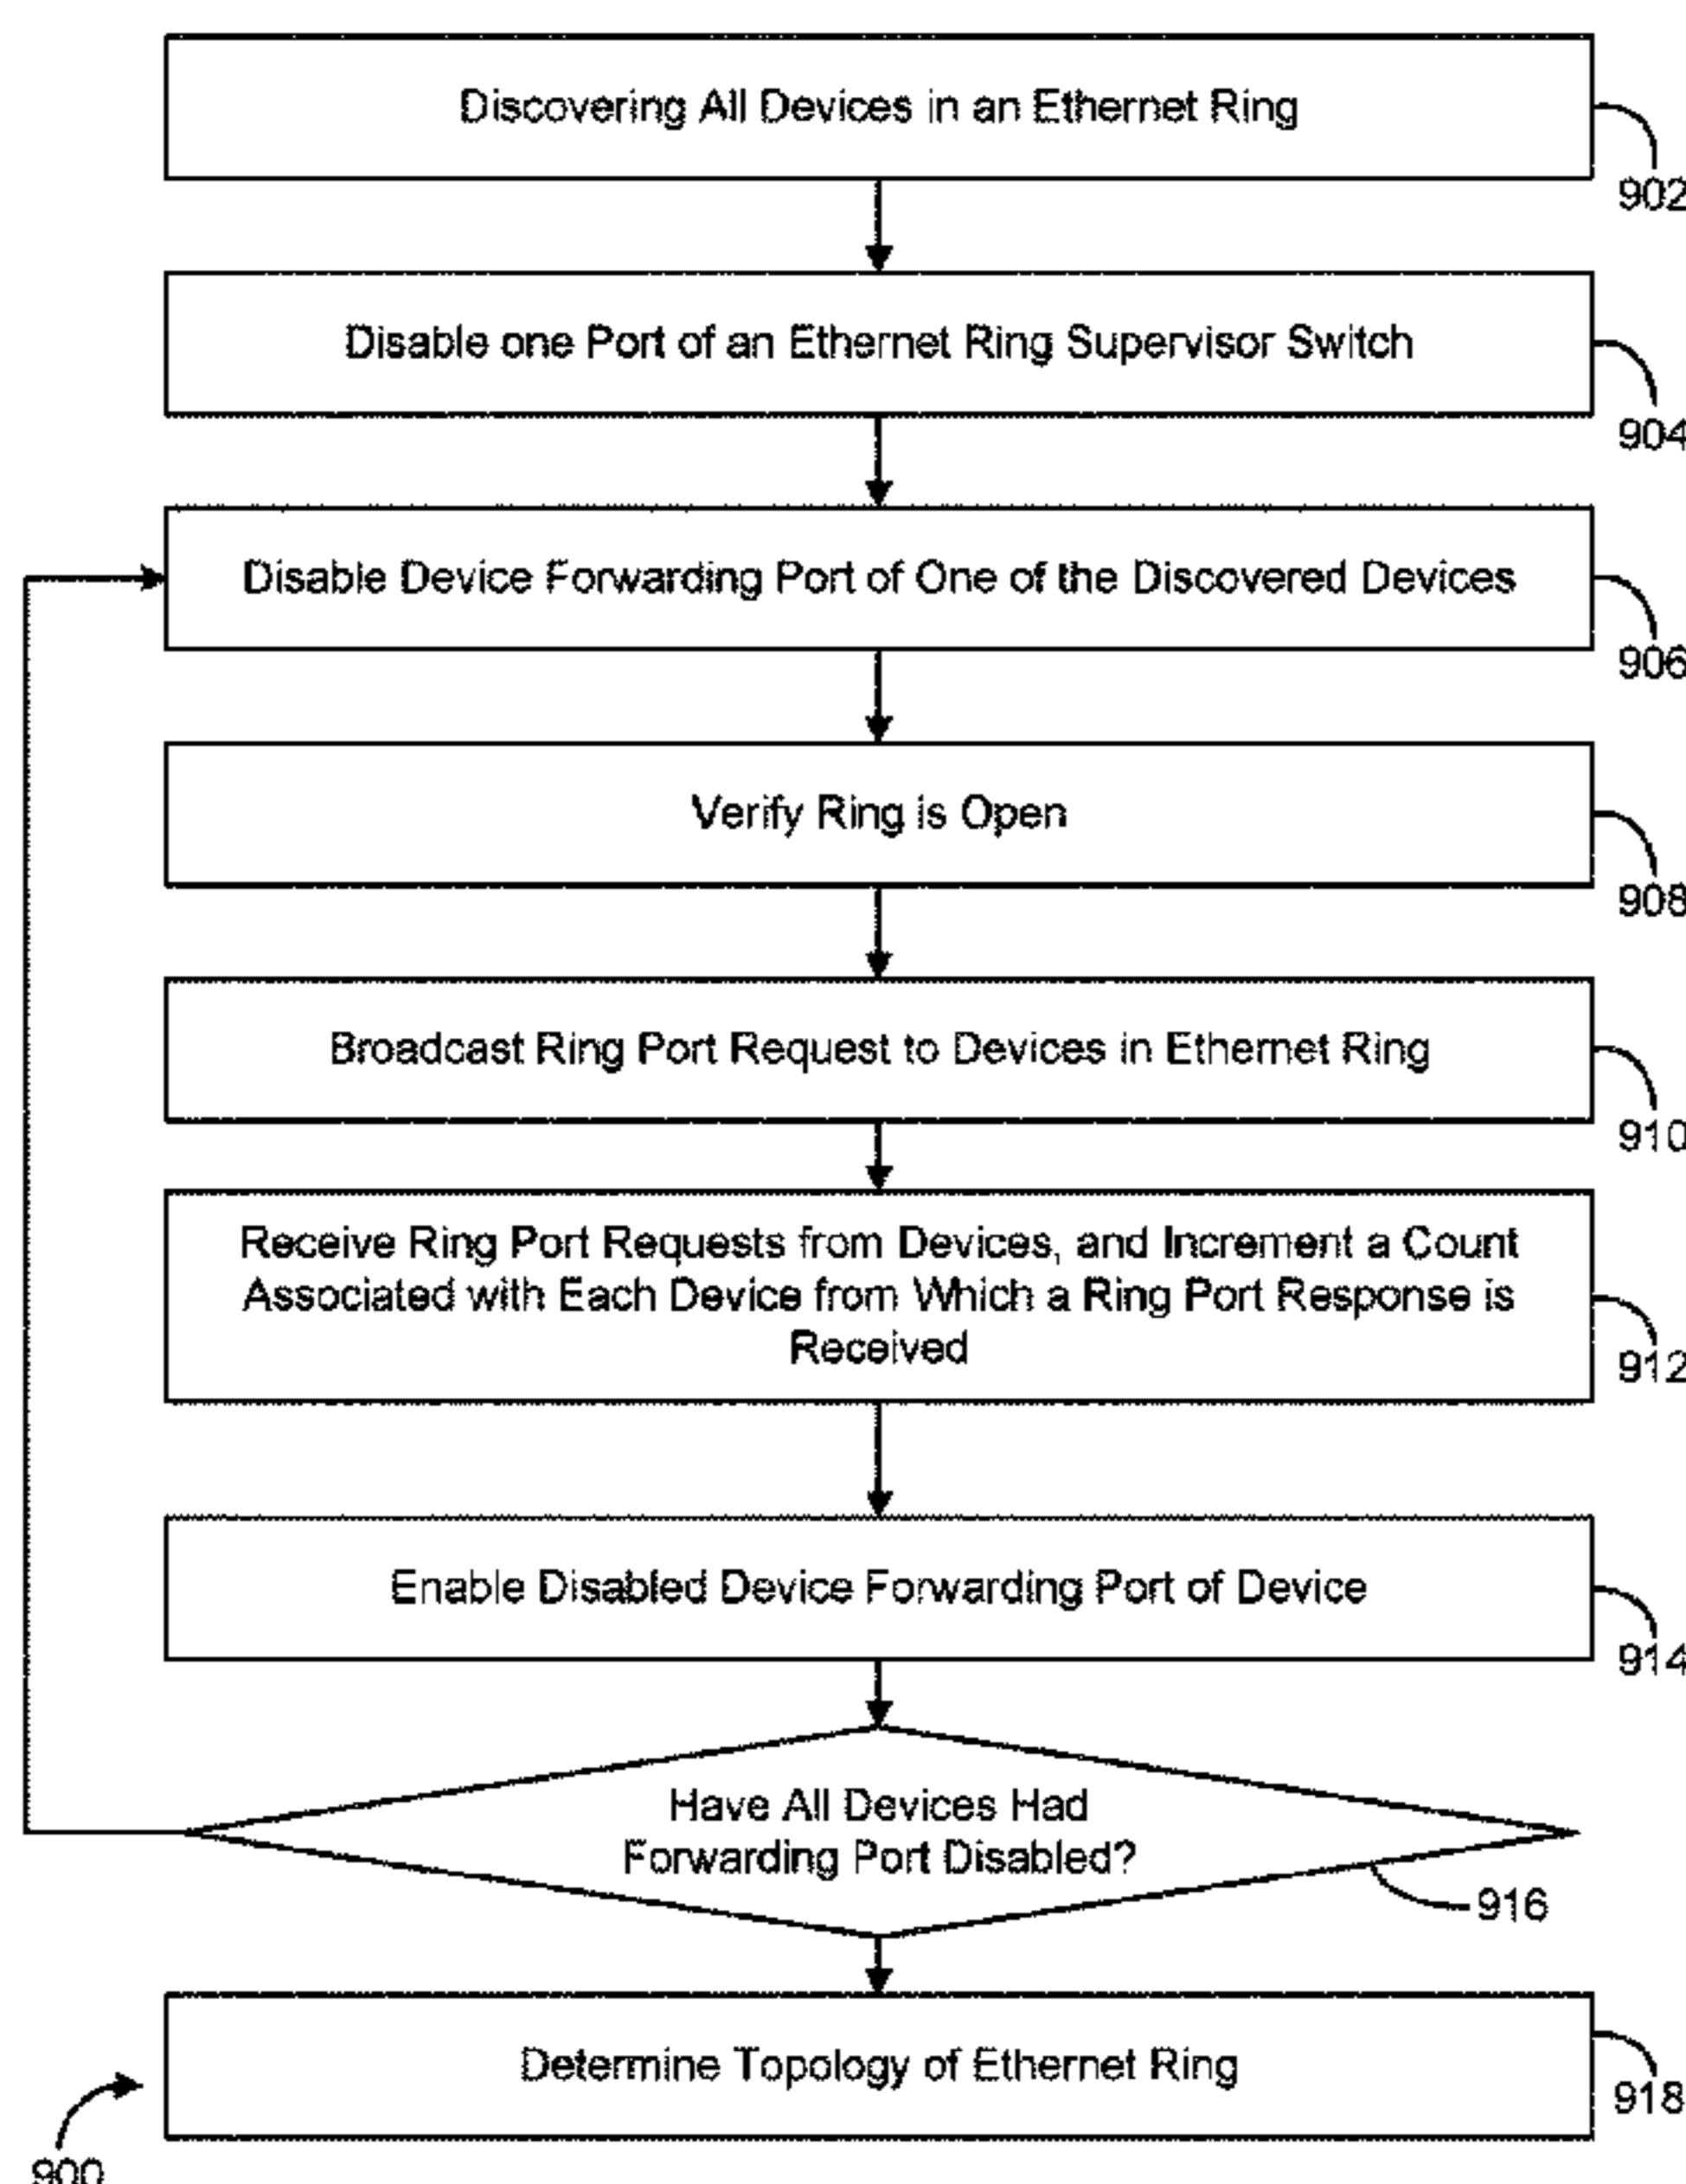
Notice of Allowance on U.S. Appl. No. 15/583,901 dated Aug. 24, 2018. 9 pages.

Primary Examiner — Colin M LaRose
(74) *Attorney, Agent, or Firm* — Foley & Lardner LLP

(57) **ABSTRACT**

A system and method of determining a topology of devices in an Ethernet network of a building management system (BMS). The method includes discovering a number and identification of devices within the Ethernet ring. The method further includes, for each device discovered in the Ethernet ring, disabling an Ethernet port of one of the device in the Ethernet ring; broadcasting a ring port request onto the Ethernet ring; receiving a ring port response from each of the devices having received the ring port request; and incrementing a count associated with each device based on receiving a ring port response from the device. The method then repeats the above steps until each device in the Ethernet ring has had an Ethernet port disabled. The method further includes determining the topology of the devices in the Ethernet ring by ordering the devices in the Ethernet ring based on the value of the count associated with each device.

40 Claims, 9 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

9,655,232 B2 * 5/2017 Saxena H04L 41/0803
 10,116,517 B1 10/2018 Drury
 10,411,984 B2 9/2019 Hoglund
 2005/0076140 A1 * 4/2005 Fung H04L 12/4625
 709/245
 2005/0243739 A1 * 11/2005 Anderson H04L 29/12783
 370/254
 2008/0056281 A1 * 3/2008 Denecheau H04L 12/42
 370/400
 2008/0239970 A1 * 10/2008 Lu H04L 12/40182
 370/245
 2010/0165883 A1 * 7/2010 Holness H04L 45/28
 370/255
 2011/0019536 A1 * 1/2011 Kim H04L 12/437
 370/223
 2014/0321261 A1 * 10/2014 Lee H04L 12/437
 370/216
 2016/0036602 A1 * 2/2016 Forguites H04L 12/4637
 370/404
 2017/0012694 A1 * 1/2017 Kaku H04B 7/14
 2019/0349218 A1 11/2019 Hoglund et al.
 2020/0076694 A1 3/2020 Drury et al.
 2020/0076735 A1 3/2020 Hoglund et al.

* cited by examiner




FIG. 1
10

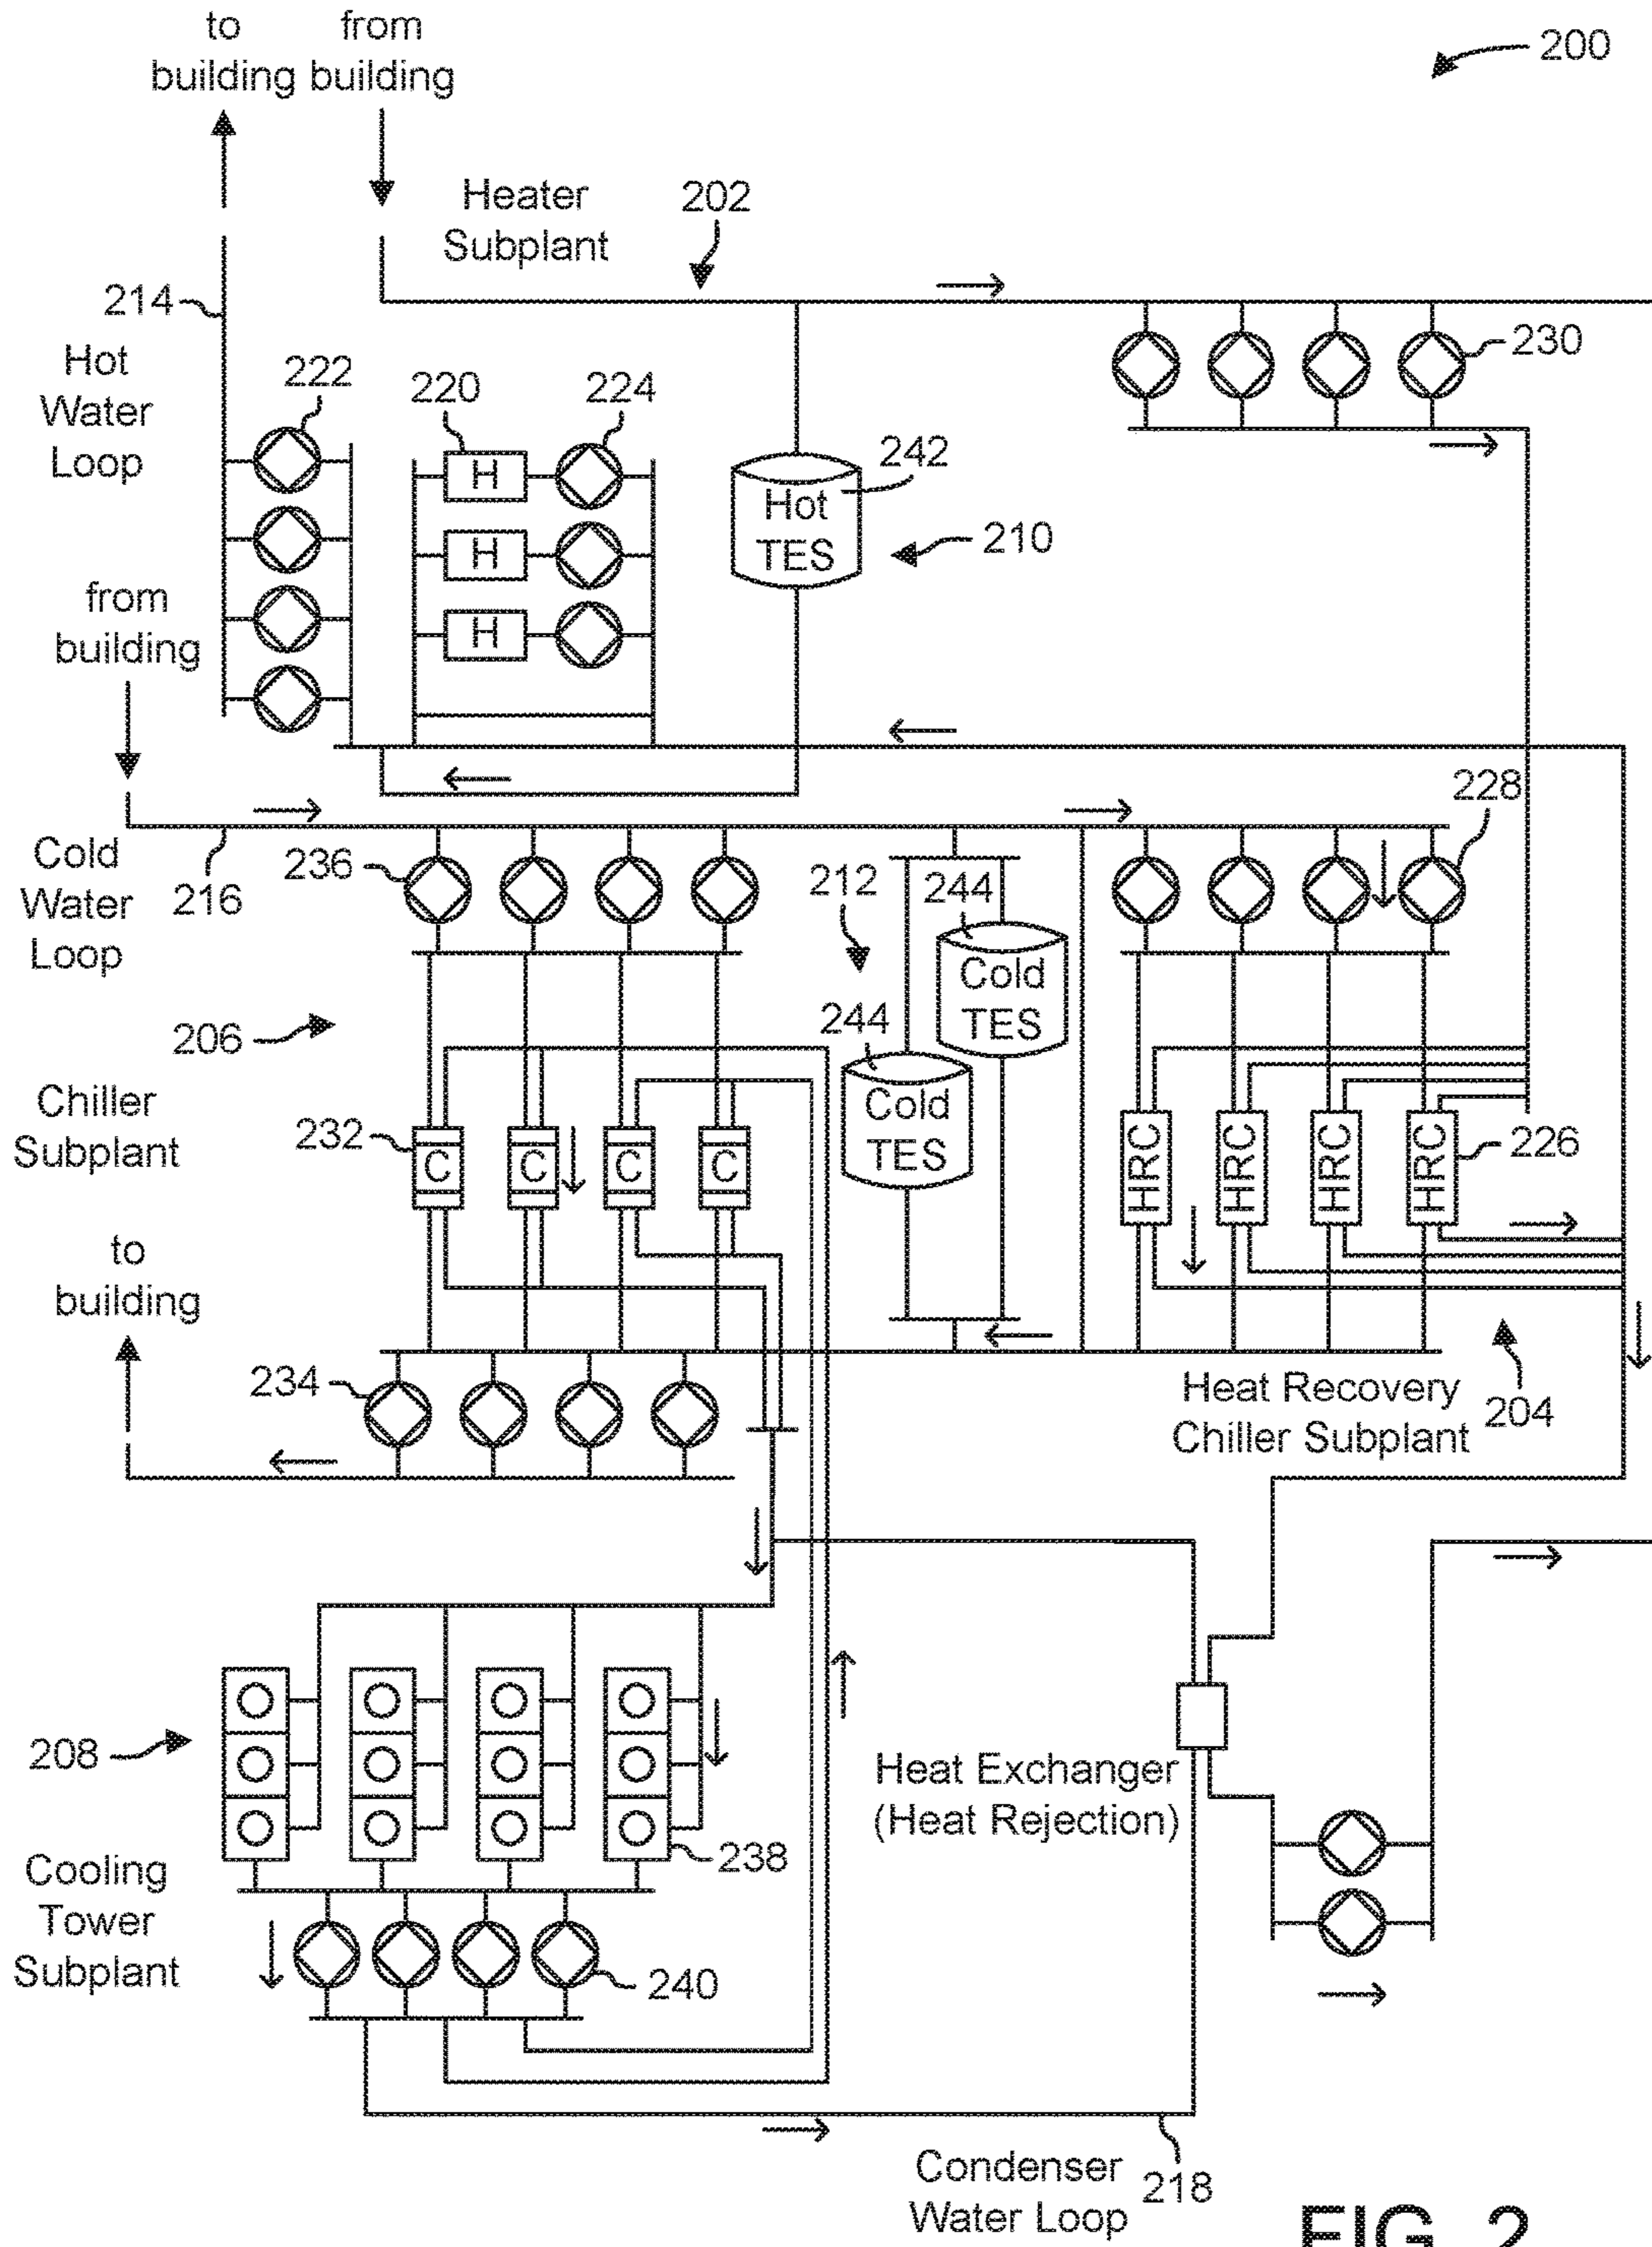


FIG. 2

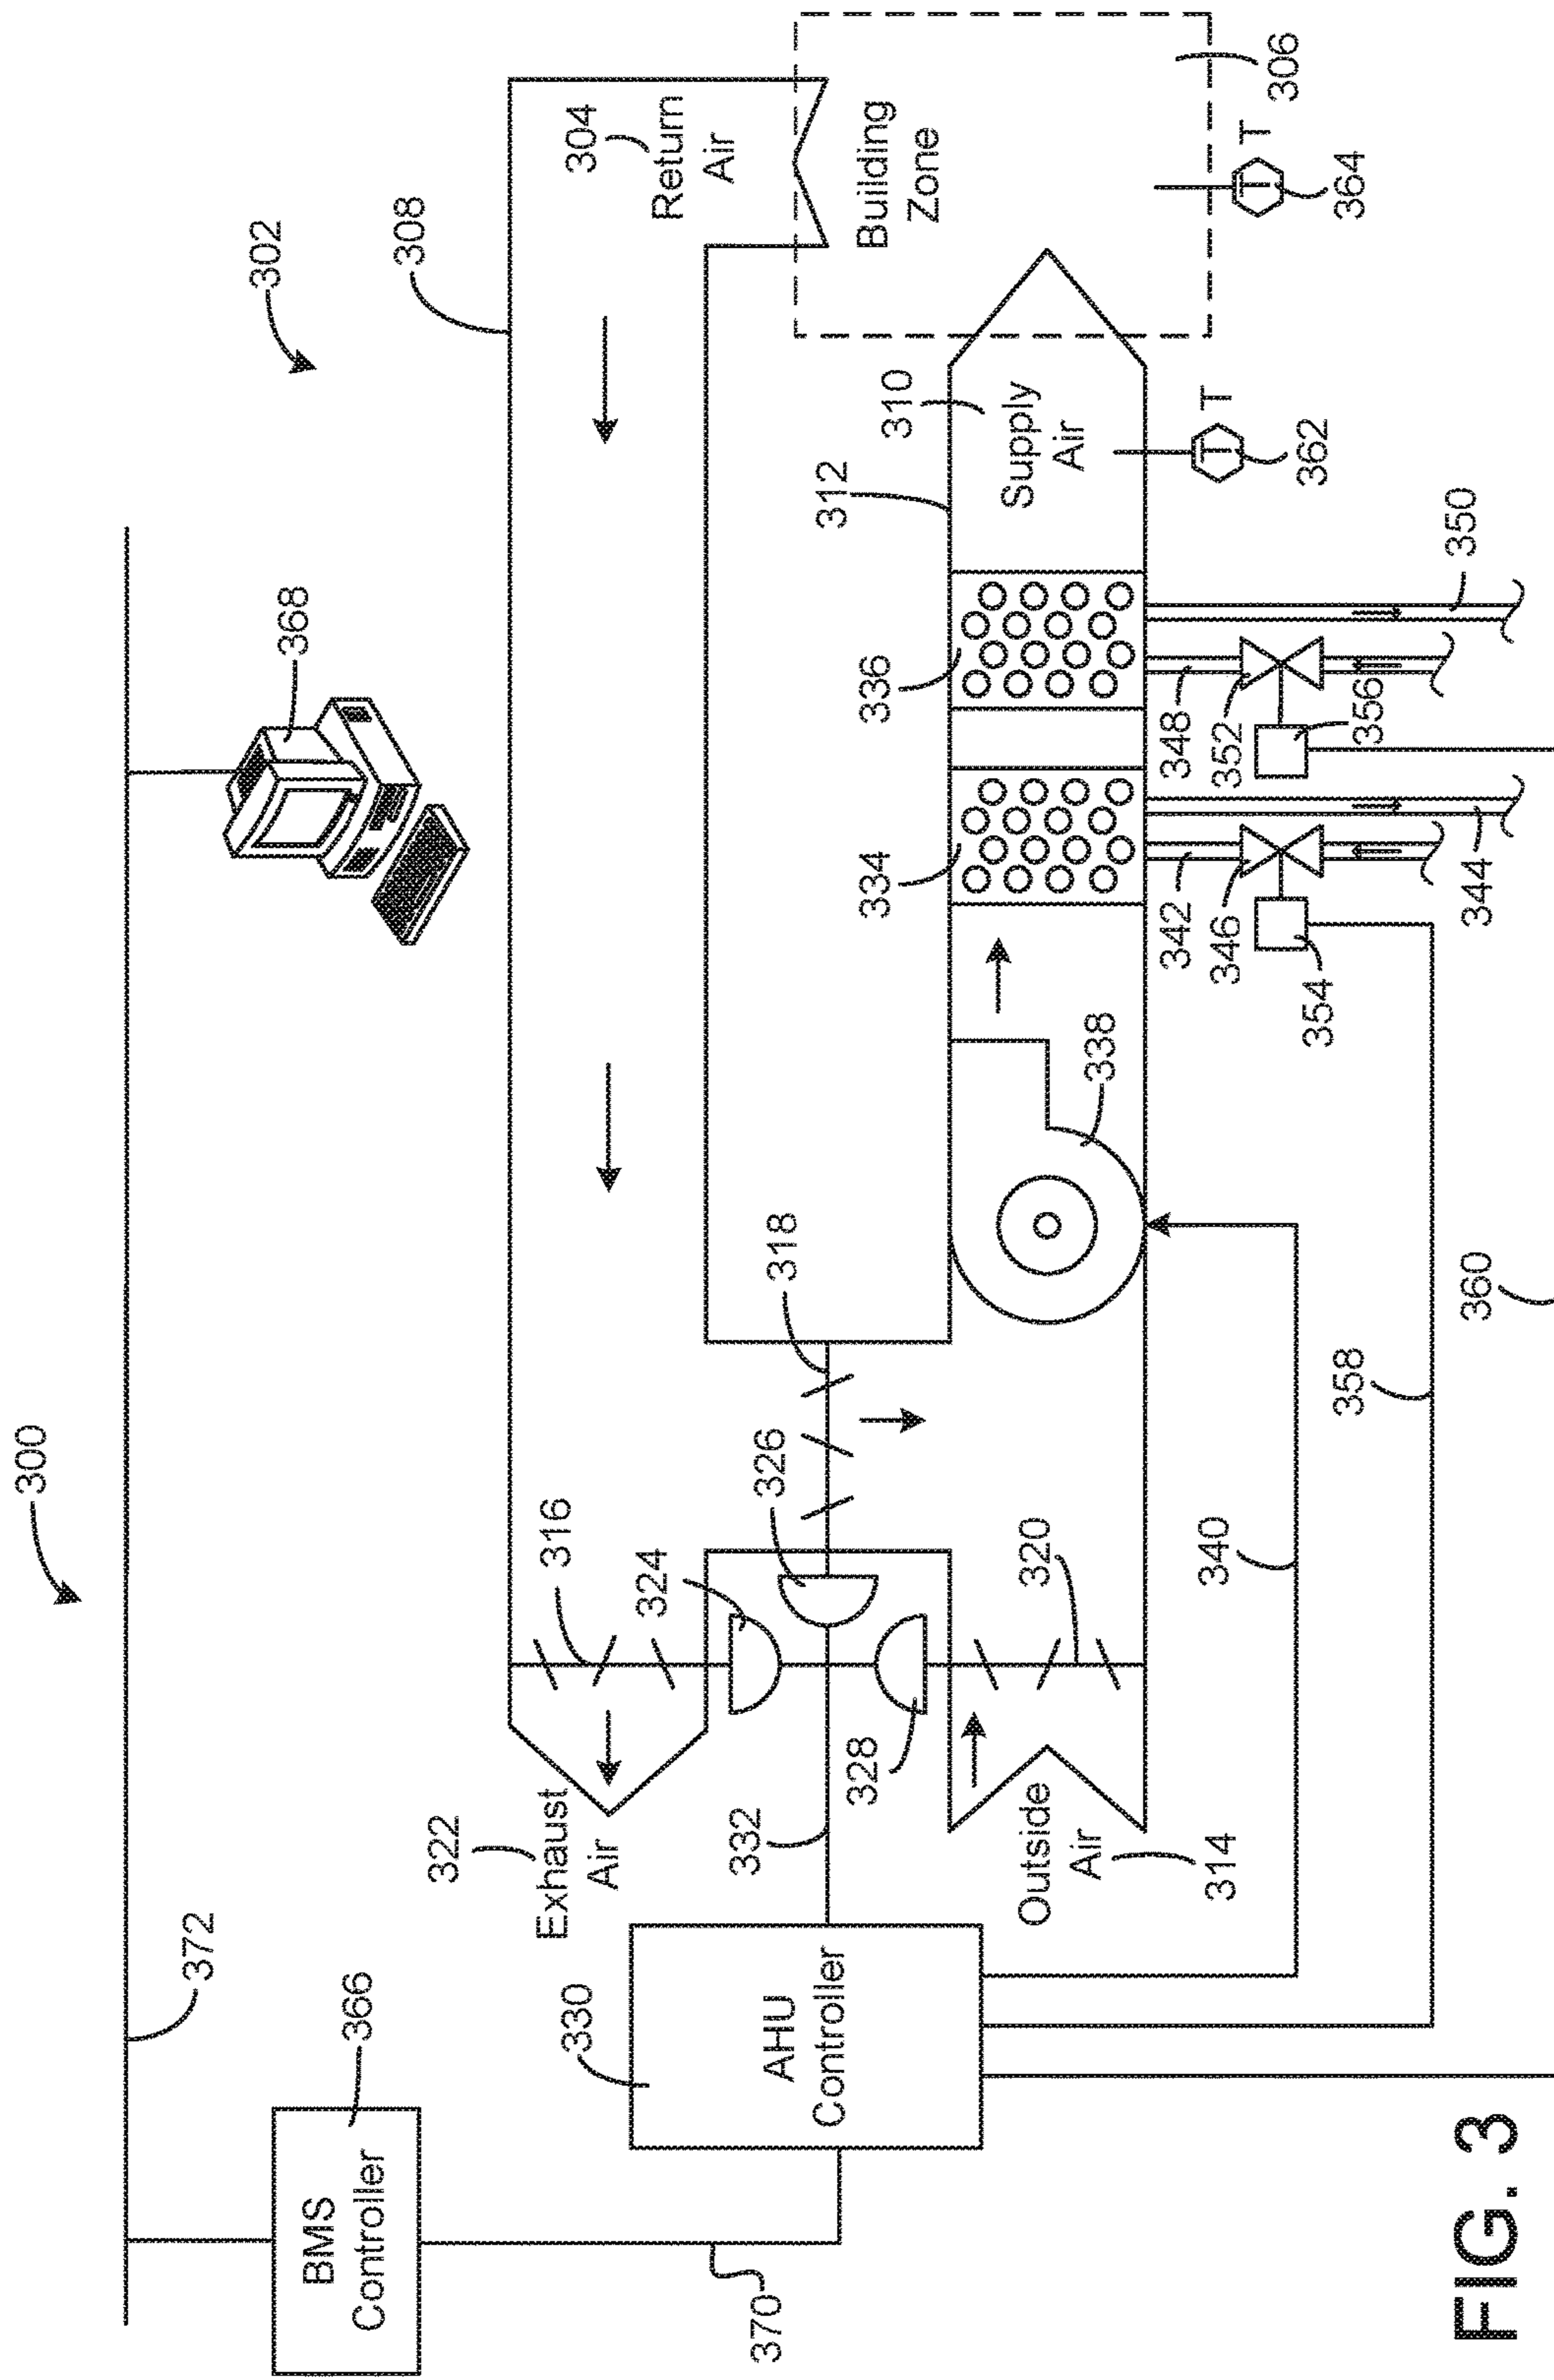


FIG. 3




FIG. 4

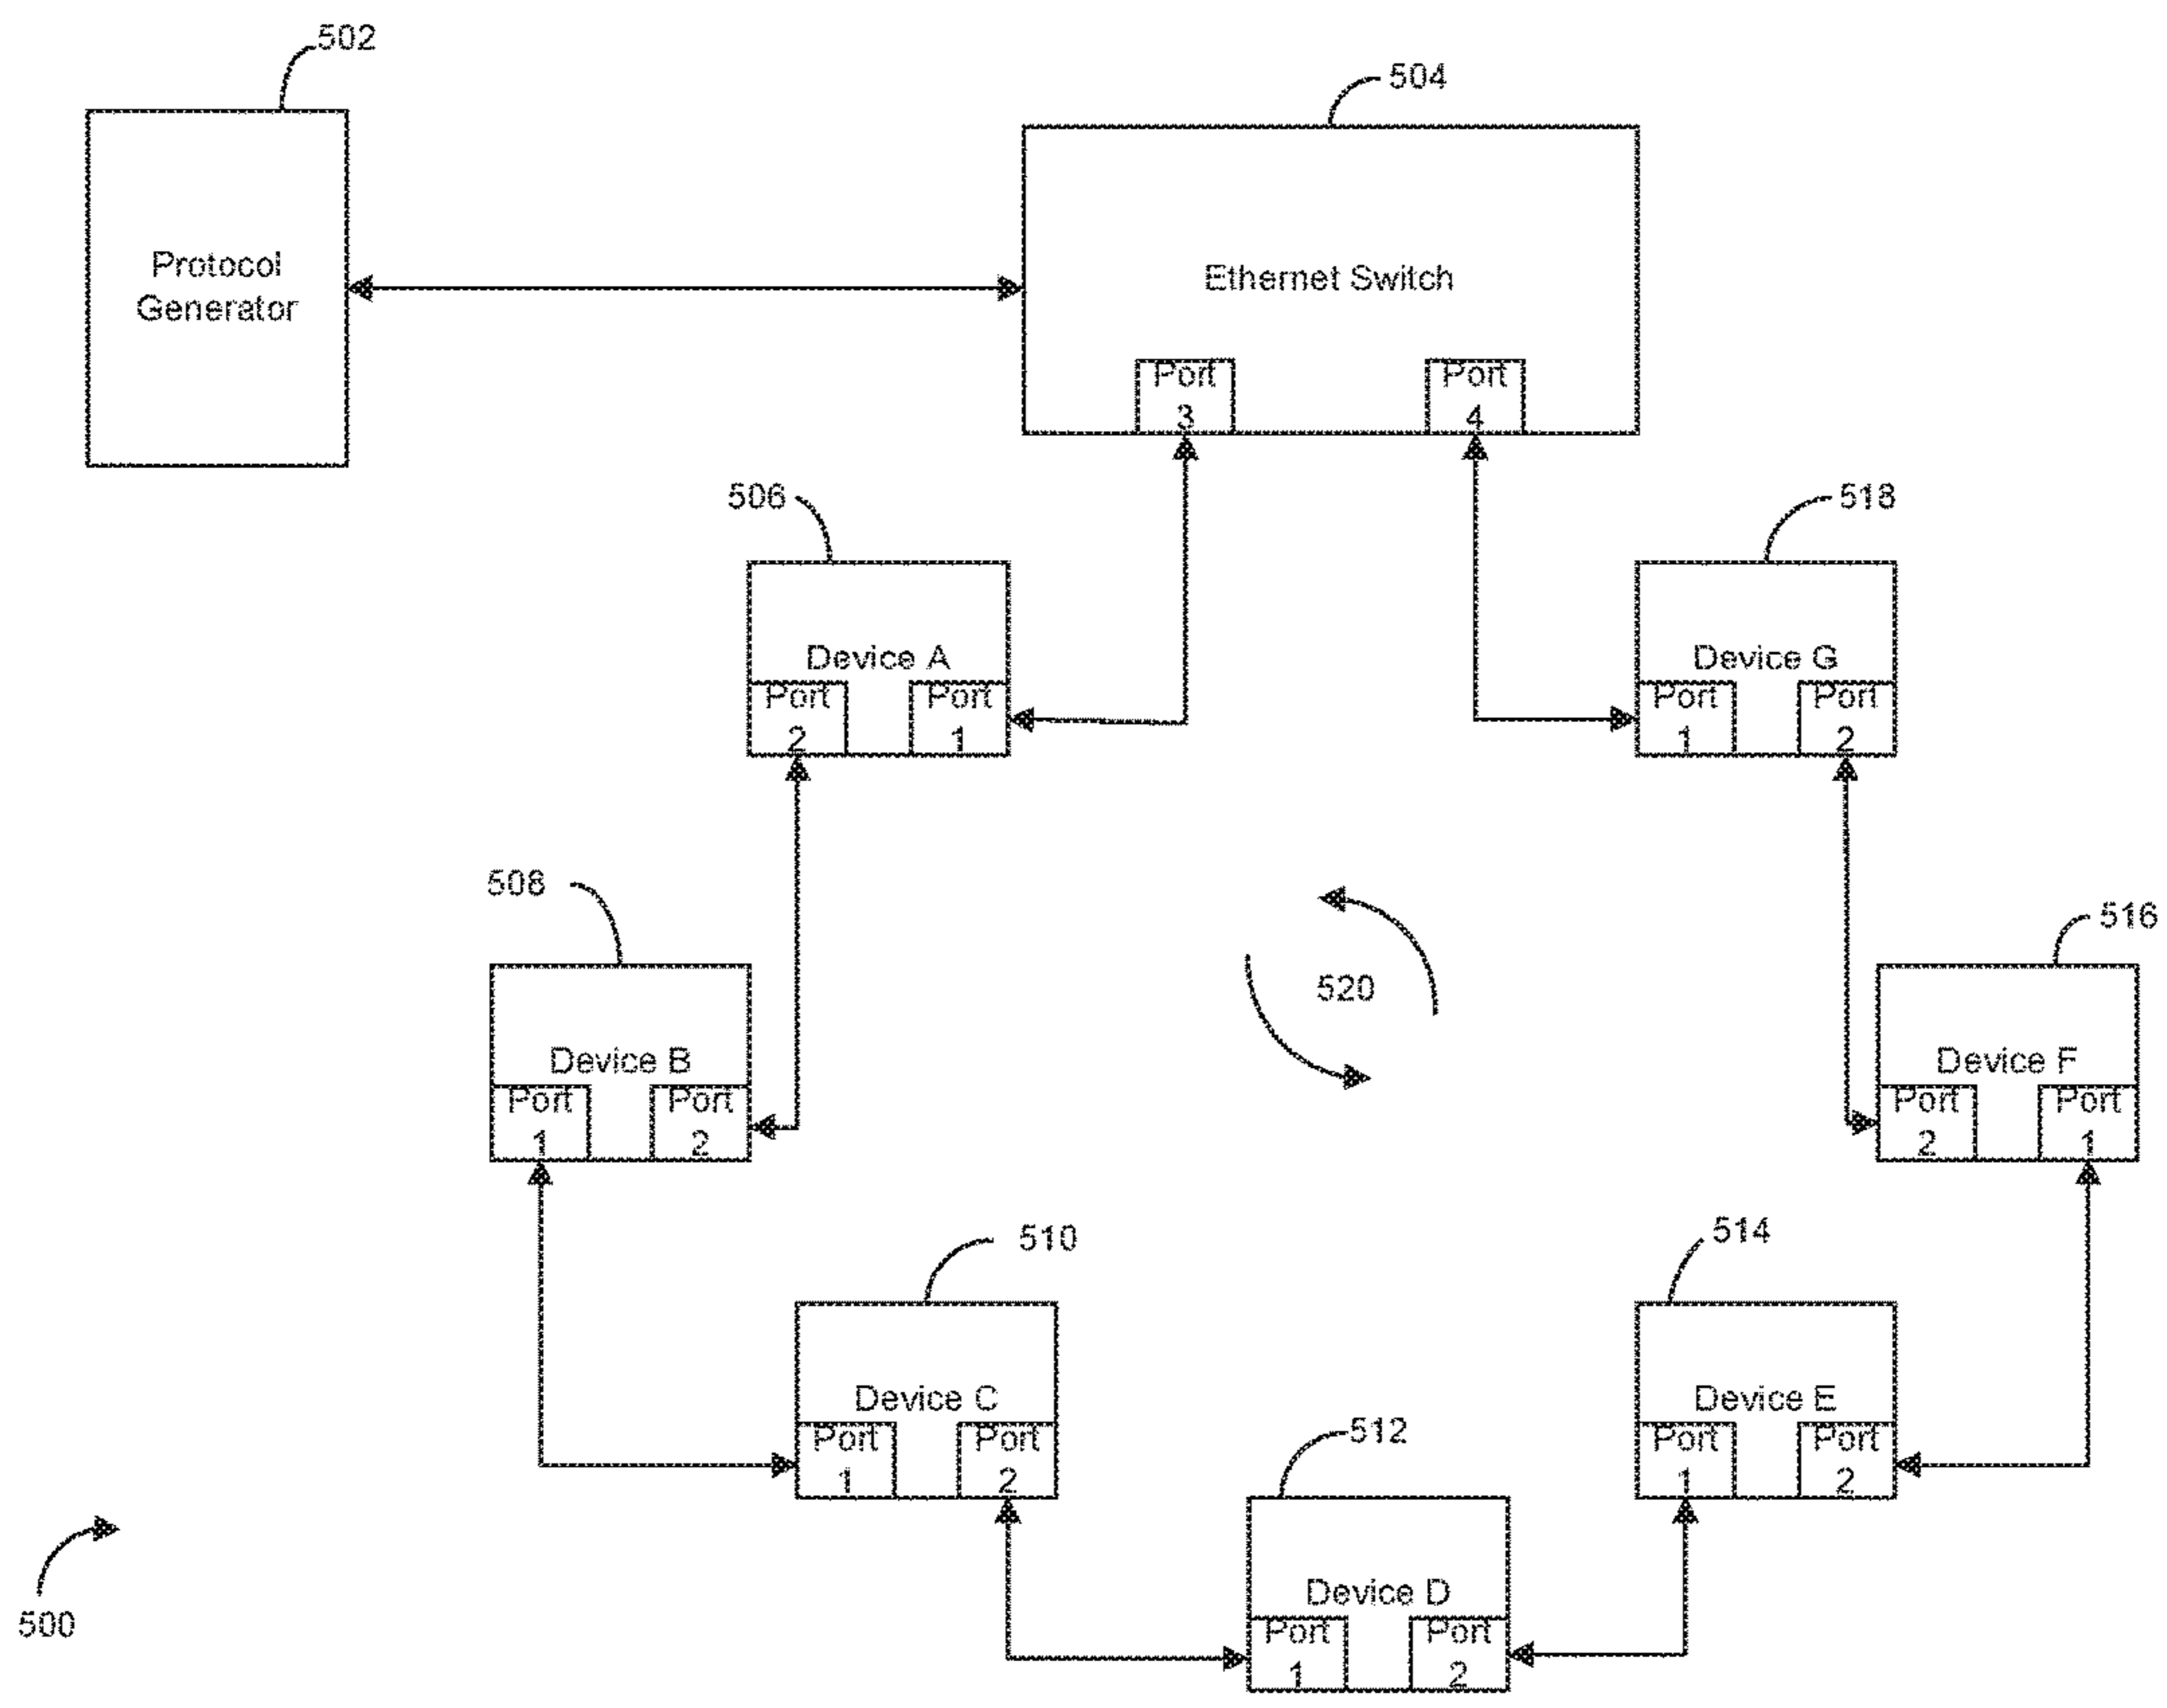


FIG. 5

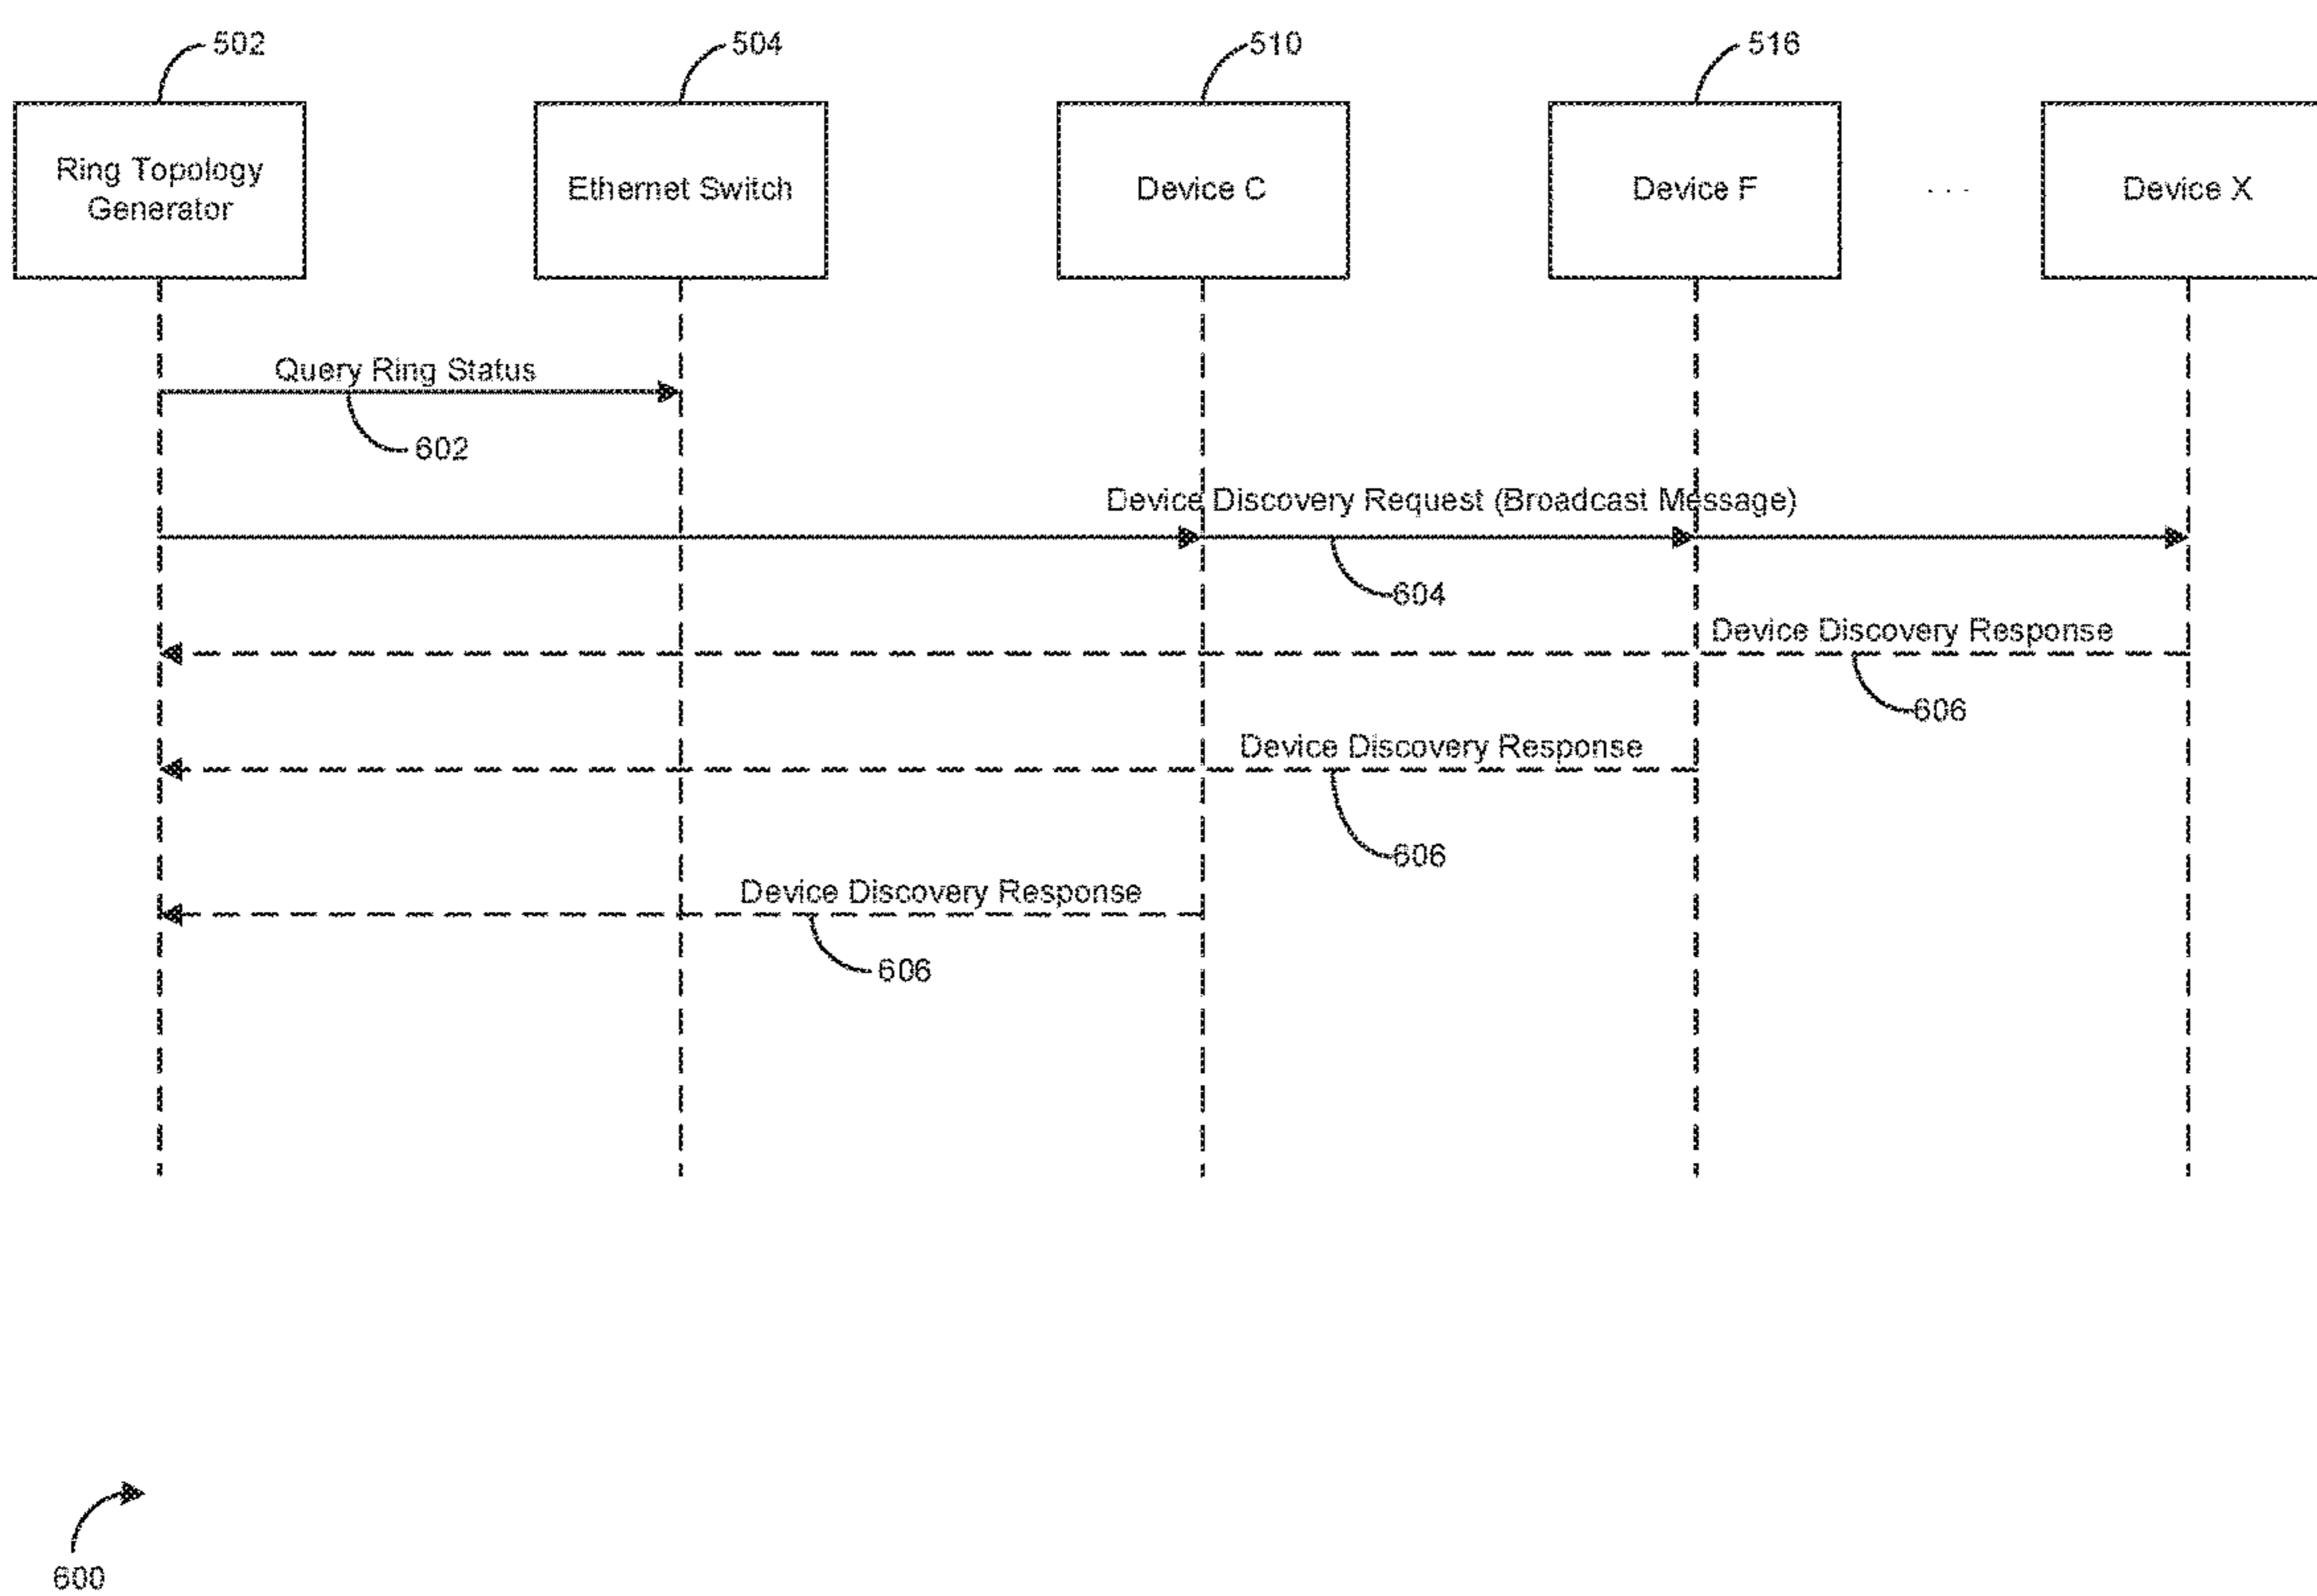


FIG. 6

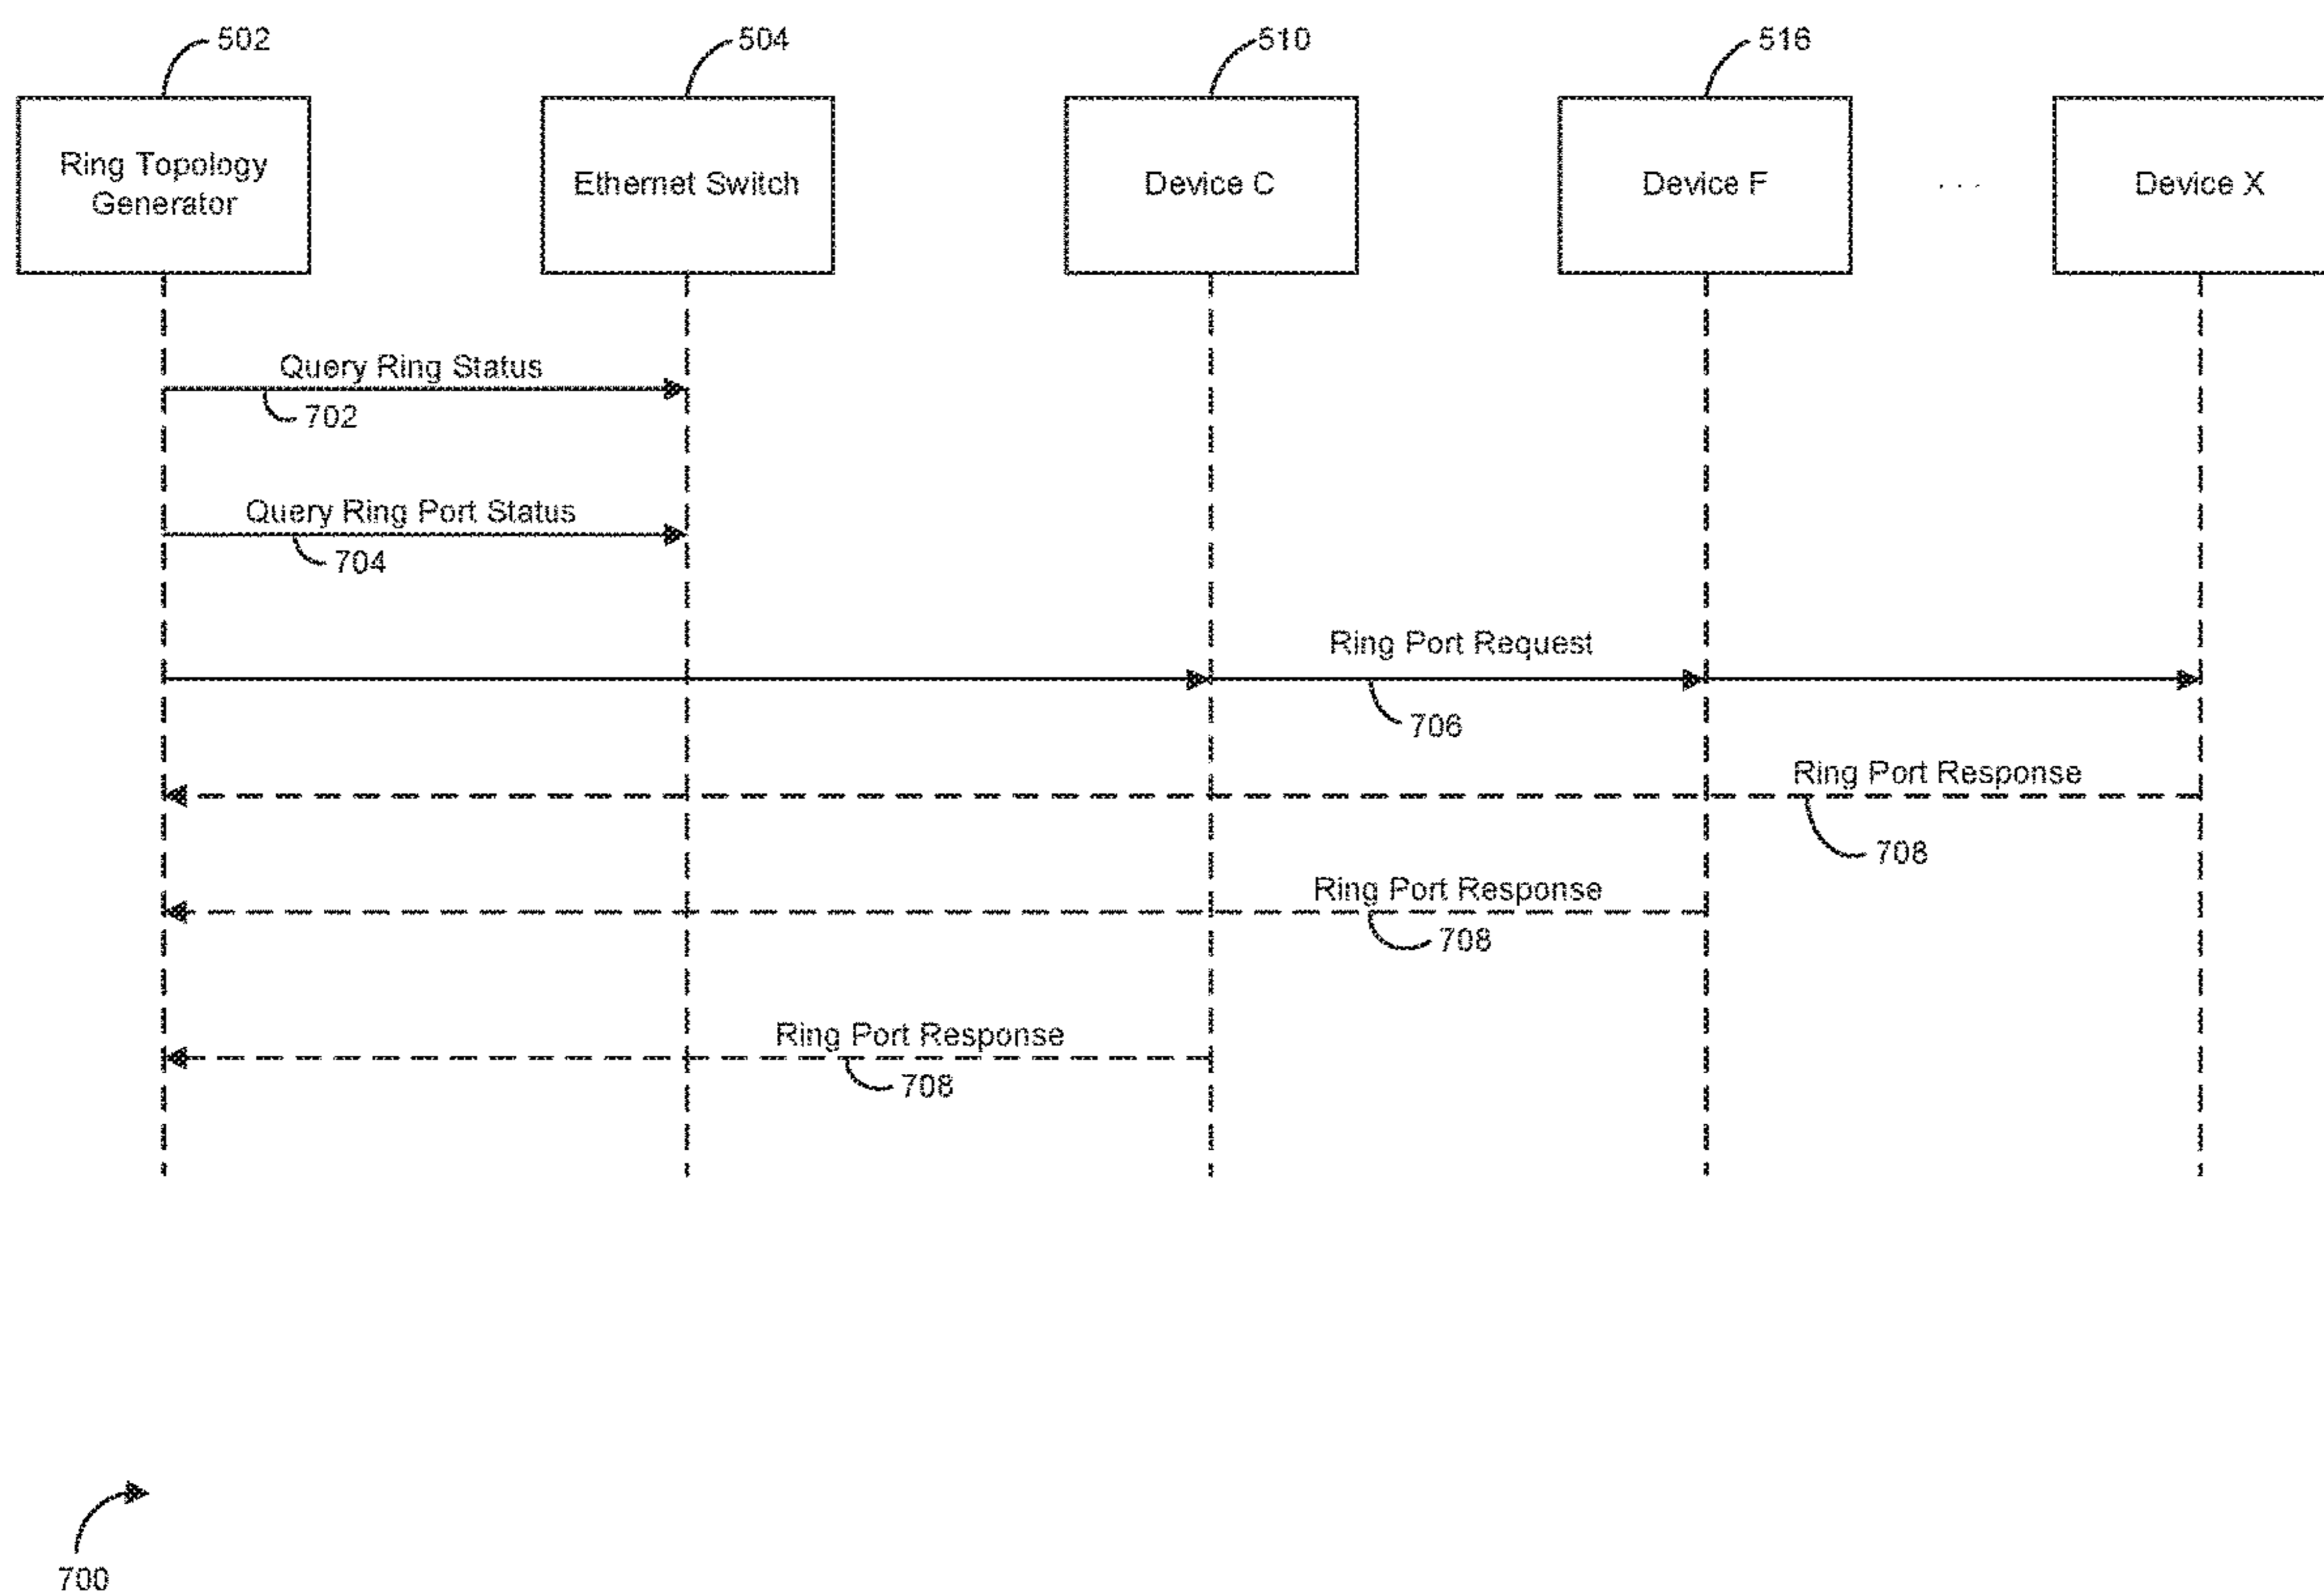


FIG. 7

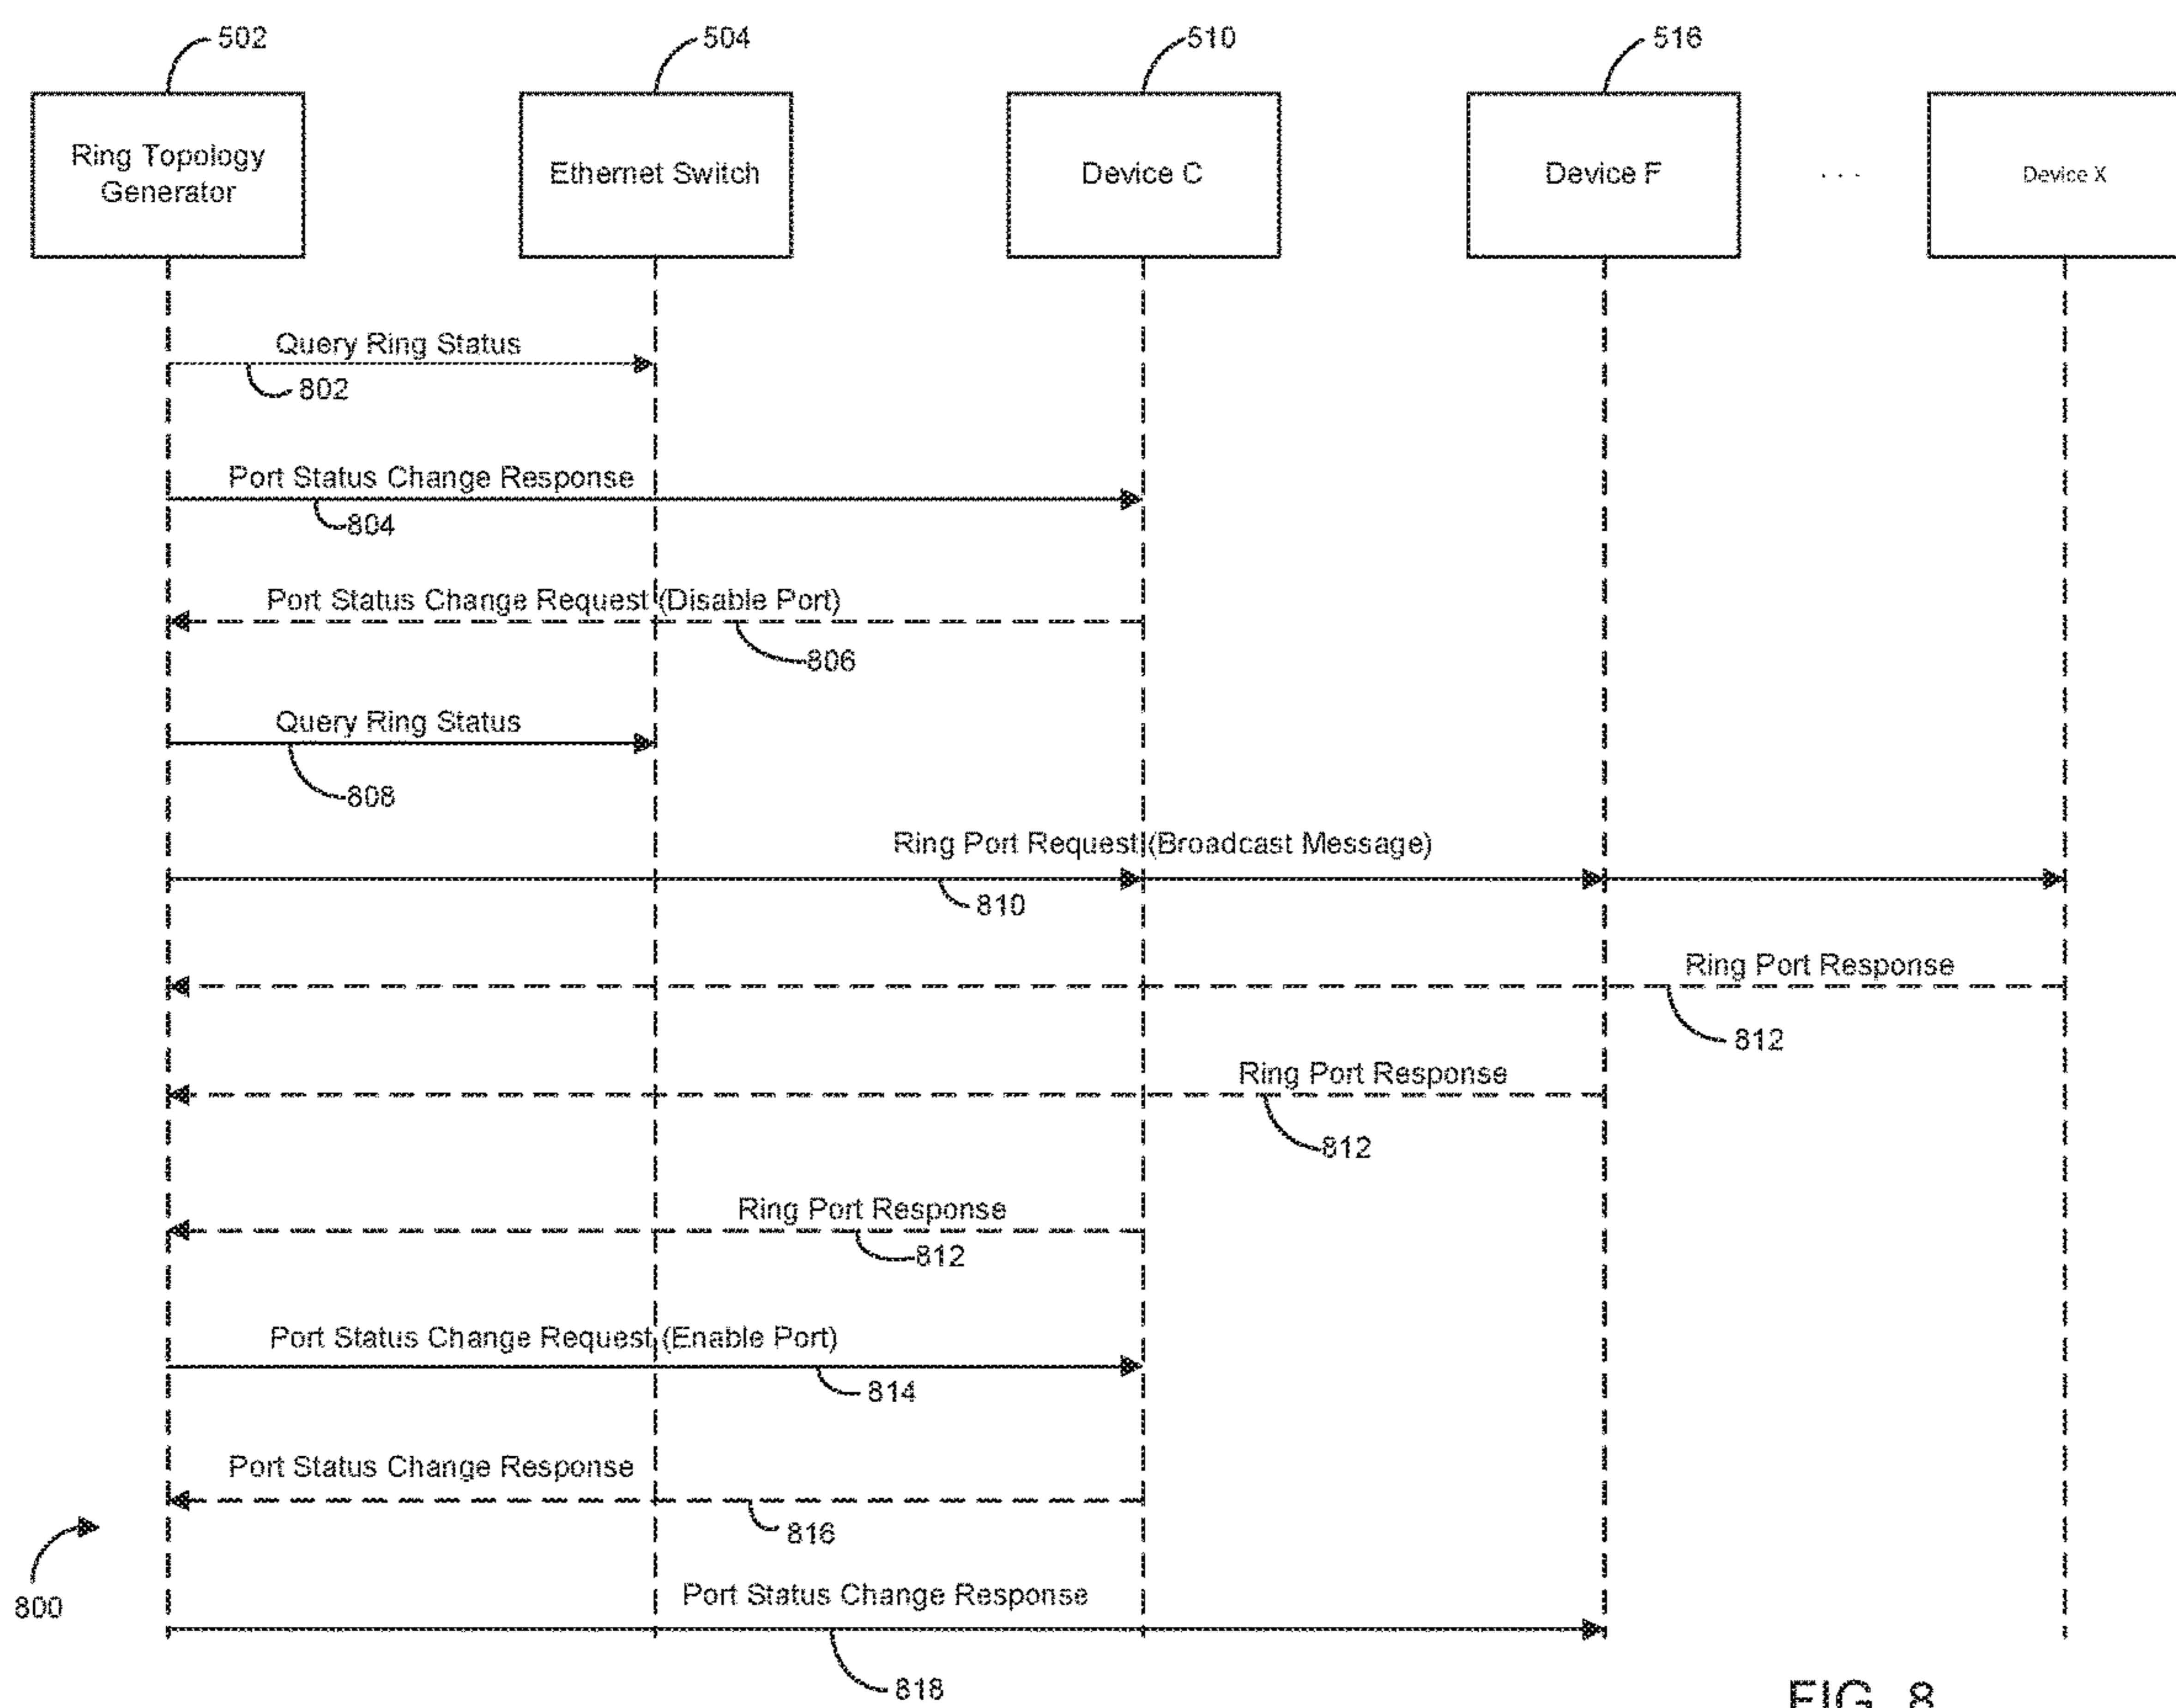


FIG. 8

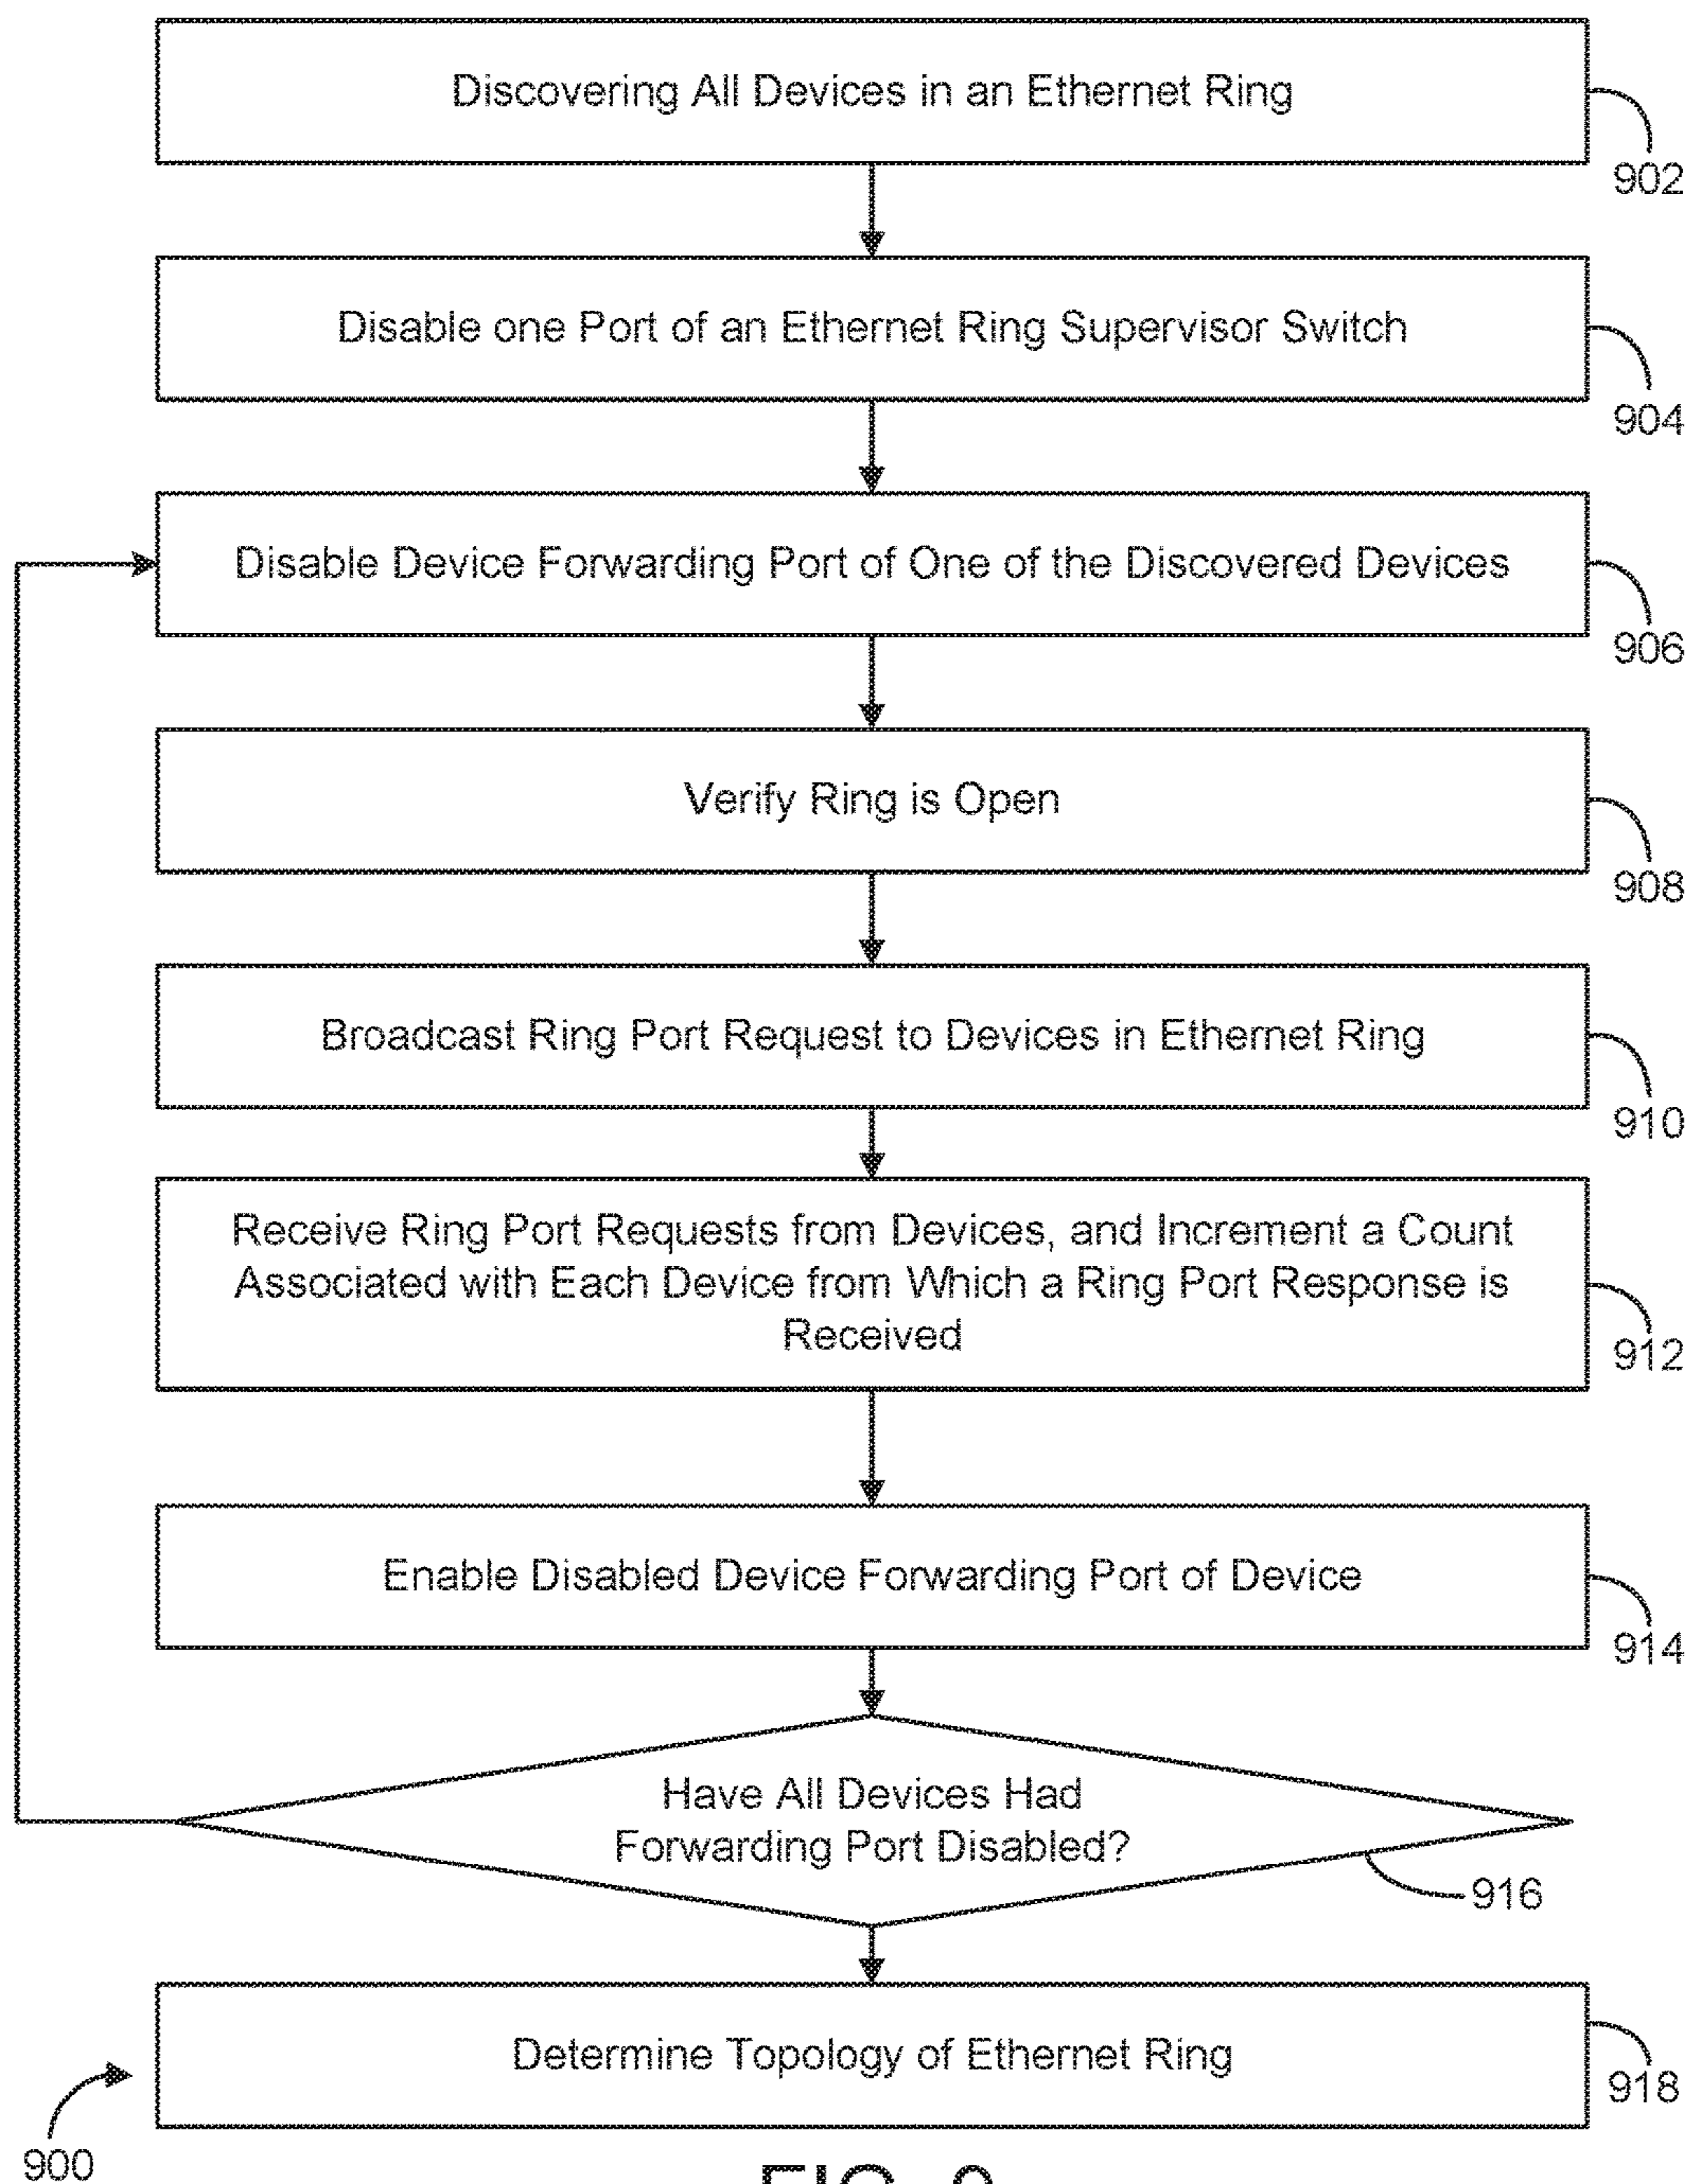


FIG. 9

**SYSTEMS AND METHODS FOR
DETERMINING A TOPOLOGY OF AN
ETHERNET RING IN A BUILDING
MANAGEMENT SYSTEM**

Matter enclosed in heavy brackets [] appears in the original patent but forms no part of this reissue specification; matter printed in italics indicates the additions made by reissue; a claim printed with strikethrough indicates that the claim was canceled, disclaimed, or held invalid by a prior post-patent action or proceeding.

*CROSS-REFERENCE TO RELATED
APPLICATION*

This application is a reissue application of U.S. Pat. No. 10,116,517, filed May 1, 2018, the entire disclosure of which is incorporated by reference herein.

BACKGROUND

The present disclosure relates generally to building management systems. The present disclosure relates more particularly to systems and methods for identifying faults in a building management system Ethernet ring network, such as device faults, cabling faults, and/or cabling errors.

A building management system (BMS) is, in general, a system of devices configured to control, monitor, and manage equipment in or around a building or building area. A BMS can include a heating, ventilation, and air conditioning (HVAC) system, a security system, a lighting system, a fire alerting system, another system that is capable of managing building functions or devices, or any combination thereof. BMS devices may be installed in any environment (e.g., an indoor area or an outdoor area) and the environment may include any number of buildings, spaces, zones, rooms, or areas. A BMS may include a variety of devices (e.g., HVAC devices, controllers, chillers, fans, sensors, etc.) configured to facilitate monitoring and controlling the building space. Throughout this disclosure, such devices are referred to as BMS devices or building equipment.

In some BMS systems, at least some of the devices are connected in a network and particularly, in an Ethernet network. In many system, the Ethernet network may be configured in a ring topology such as a ring using a Media Redundancy Protocol (MRP) ring protocol. MRP rings can provide for single fault tolerances such that connectivity within the ring can be broken (e.g. a device fails or goes offline, or an Ethernet cable between devices is cut or disconnected) without impacting the network connectivity of the (other) devices in the ring.

While MRP rings provide a robust ring topology for a network there is currently no mechanism to accurately determine the order of the devices in the ring or how the devices are connected to each other (e.g. which switch port on a device is connected to a specific neighboring device). This data could be useful to verify that the devices in the MRP ring are connected correctly as per the system design. Further, device order and connection data can be useful in troubleshooting an MRP ring when errors or faults occur. Thus, it would be advantageous to have systems and methods that could automatically determine an order and connection scheme of devices connected in an MRP or other Ethernet ring topology.

SUMMARY

One implementation of the present disclosure is a method for determining a topology of devices in an Ethernet ring.

The method includes broadcasting a device discovery command from a ring topology generator to the devices in the Ethernet ring; and receiving device discovery information from one or more of the devices in the Ethernet ring at the ring topology generator. The method also includes querying a ring supervisor to confirm that the Ethernet ring is closed and to determine which Ethernet port of the ring supervisor is connected to Ethernet ring as the forwarding port. Additionally, the ring topology generator configured to, for each discovered device in the Ethernet ring: (i) disabling an Ethernet port of one of the devices in the Ethernet ring; (ii) verifying the Ethernet ring is open; (iii) broadcasting a ring port request onto the Ethernet ring from the ring supervisor; (iv) receiving a ring port response from each of the devices having received the ring port request; (v) incrementing a count associated with each device based on receiving a ring port response from the device; (vi) re-enabling the Ethernet port of the one of the devices in the Ethernet ring; and (vii) repeating steps (i) through (vi) until each device in the Ethernet ring has had an Ethernet port disabled. The method further includes determining the topology of the devices of the Ethernet ring by ordering the devices in the Ethernet ring based on the value of the count associated with each device. The device having the highest count value being closest to the forwarding port of the ring supervisor and the device having the lowest count value being furthest from the forwarding port of the ring supervisor.

A further implementation of the present disclosure is a system for determining a topology of an Ethernet ring. The system includes a number of Ethernet devices configured in an Ethernet ring. The system further includes a ring supervisor. The ring supervisor is integral to the Ethernet ring and configured to manage data flow within the Ethernet ring. The system further includes a ring topology generator configured to broadcast a device discovery command from a ring topology generator to the devices in the Ethernet ring. The ring topology generator is further configured to receive device discovery information from one or more of the plurality of devices in the Ethernet ring, the received device discovery information indicating a discovered device. The ring topology generator further configured to instruct the ring supervisory to initialize one Ethernet port connected to the ring as the forwarding port, and disable the other port connected to the Ethernet ring. The ring topology generator further configured to execute the following instructions for each of the discovered devices: (i) disable an Ethernet port of one of the devices in the Ethernet ring; (ii) verify the Ethernet ring is open; (iii) broadcast a ring port request onto the Ethernet ring from the forwarding port of the ring supervisor; (iv) receive a ring port response from each of the devices having received the ring port request; (v) increment a count associated with each device based on receiving a ring port response from the device; (vi) re-enable the Ethernet port of the one of the devices in the Ethernet ring; and (vii) repeat steps (i) through (vi) until each device in the Ethernet ring has had an Ethernet port disabled. The ring topology generator further configured to determine the topology of the devices of the Ethernet ring by ordering the devices in the Ethernet ring based on the value of the count associated with each devices. The device having the highest count value being closest to the forwarding port of the ring supervisor and the device having the lowest count value being furthest from the forwarding port of the ring supervisor.

A further implementation of the present disclosure is a method of determining a topology of devices in an Ethernet ring. The method includes discovering a number and iden-

tification of devices within an Ethernet ring. The method further includes, for each devices discovered in the Ethernet ring: (i) disabling an Ethernet port of one of the devices in the Ethernet ring; (ii) verifying the Ethernet ring is open; (iii) broadcasting a ring port request onto the Ethernet ring from the forwarding port of the ring supervisor; (iv) receiving a ring port response from each of the devices having received the ring port request; (v) incrementing a count associated with each device based on receiving a ring port response from the device; (vi) re-enabling the Ethernet port of the one of the devices in the Ethernet ring; and (vii) repeating steps (i) through (vi) until each device in the Ethernet ring has had an Ethernet port disabled. The method further includes determining the topology of the devices of the Ethernet ring by ordering the devices in the Ethernet ring based on the value of the count associated with each devices. The device having the highest count value being closest to the forwarding port of the ring supervisor and the device having the lowest count value being furthest from the forwarding port of the ring supervisor.

Those skilled in the art will appreciate that the summary is illustrative only and is not intended to be in any way limiting. Other aspects, inventive features, and advantages of the devices and/or processes described herein, as defined solely by the claims, will become apparent in the detailed description set forth herein and taken in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a drawing of a building equipped with a building management system (BMS) and a HVAC system, according to some embodiments.

FIG. 2 is a schematic of a waterside system which can be used as part of the HVAC system of FIG. 1, according to some embodiments.

FIG. 3 is a block diagram illustrating an airside system which can be used as part of the HVAC system of FIG. 1, according to some embodiments.

FIG. 4 is a block diagram illustrating a BMS which can be used in the building of FIG. 1, according to some embodiments.

FIG. 5 is a block diagram illustrating a network connected to a plurality of devices, according to some embodiments.

FIG. 6 is a data flow chart illustrating a process for discovering devices in an Ethernet ring, according to some embodiments.

FIG. 7 is a data flow chart illustrating an Ethernet port orientation process, according to some embodiments.

FIG. 8 is a data flow chart illustrating a diagnostic flow to establish that the Ethernet network of FIG. 5 is in a closed status, according to some embodiments.

FIG. 9 is a flow diagram illustrating a process for determining a topology of an Ethernet ring, according to some embodiments.

DETAILED DESCRIPTION

Building Management System and HVAC System

Referring now to FIGS. 1-4, an exemplary building management system (BMS) and a heating, ventilation, and air conditioning (HVAC) system in which the systems and methods of the present disclosure can be implemented are shown, according to an exemplary embodiment. Referring particularly to FIG. 1, a perspective view of a building 10 is shown. Building 10 is served by a BMS. A BMS is, in general, a system of devices configured to control, monitor,

and manage equipment in or around a building or building area. A BMS can include, for example, a HVAC system, a security system, a lighting system, a fire alerting system, any other system that is capable of managing building functions or devices, or any combination thereof.

The BMS that serves building 10 includes an HVAC system 100. HVAC system 100 can include a plurality of HVAC devices (e.g., heaters, chillers, air handling units, pumps, fans, thermal energy storage, etc.) configured to provide heating, cooling, ventilation, or other services for building 10. For example, HVAC system 100 is shown to include a waterside system 120 and an airside system 130. Waterside system 120 can provide a heated or chilled fluid to an air handling unit of airside system 130. Airside system 130 can use the heated or chilled fluid to heat or cool an airflow provided to building 10. An exemplary waterside system and airside system which can be used in HVAC system 100 are described in greater detail with reference to FIGS. 2-3.

HVAC system 100 is shown to include a chiller 102, a boiler 104, and a rooftop air handling unit (AHU) 106. Waterside system 120 can use boiler 104 and chiller 102 to heat or cool a working fluid (e.g., water, glycol, etc.) and can circulate the working fluid to AHU 106. In various embodiments, the HVAC devices of waterside system 120 can be located in or around building 10 (as shown in FIG. 1) or at an offsite location such as a central plant (e.g., a chiller plant, a steam plant, a heat plant, etc.). The working fluid can be heated in boiler 104 or cooled in chiller 102, depending on whether heating or cooling is required in building 10. Boiler 104 can add heat to the circulated fluid, for example, by burning a combustible material (e.g., natural gas) or using an electric heating element. Chiller 102 can place the circulated fluid in a heat exchange relationship with another fluid (e.g., a refrigerant) in a heat exchanger (e.g., an evaporator) to absorb heat from the circulated fluid. The working fluid from chiller 102 and/or boiler 104 can be transported to AHU 106 via piping 108.

AHU 106 can place the working fluid in a heat exchange relationship with an airflow passing through AHU 106 (e.g., via one or more stages of cooling coils and/or heating coils). The airflow can be, for example, outside air, return air from within building 10, or a combination of both. AHU 106 can transfer heat between the airflow and the working fluid to provide heating or cooling for the airflow. For example, AHU 106 can include one or more fans or blowers configured to pass the airflow over or through a heat exchanger containing the working fluid. The working fluid can then return to chiller 102 or boiler 104 via piping 110.

Airside system 130 can deliver the airflow supplied by AHU 106 (i.e., the supply airflow) to building 10 via air supply ducts 112 and can provide return air from building 10 to AHU 106 via air return ducts 114. In some embodiments, airside system 130 includes multiple variable air volume (VAV) units 116. For example, airside system 130 is shown to include a separate VAV unit 116 on each floor or zone of building 10. VAV units 116 can include dampers or other flow control elements that can be operated to control an amount of the supply airflow provided to individual zones of building 10. In other embodiments, airside system 130 delivers the supply airflow into one or more zones of building 10 (e.g., via supply ducts 112) without using intermediate VAV units 116 or other flow control elements. AHU 106 can include various sensors (e.g., temperature sensors, pressure sensors, etc.) configured to measure attributes of the supply airflow. AHU 106 can receive input from sensors located within AHU 106 and/or within the building

zone and can adjust the flow rate, temperature, or other attributes of the supply airflow through AHU 106 to achieve set-point conditions for the building zone.

Referring now to FIG. 2, a block diagram of a waterside system 200 is shown, according to an exemplary embodiment. In various embodiments, waterside system 200 can supplement or replace waterside system 120 in HVAC system 100 or can be implemented separate from HVAC system 100. When implemented in HVAC system 100, waterside system 200 can include a subset of the HVAC devices in HVAC system 100 (e.g., boiler 104, chiller 102, pumps, valves, etc.) and can operate to supply a heated or chilled fluid to AHU 106. The HVAC devices of waterside system 200 can be located within building 10 (e.g., as components of waterside system 120) or at an offsite location such as a central plant.

In FIG. 2, waterside system 200 is shown as a central plant having a plurality of subplants 202-212. Subplants 202-212 are shown to include a heater subplant 202, a heat recovery chiller subplant 204, a chiller subplant 206, a cooling tower subplant 208, a hot thermal energy storage (TES) subplant 210, and a cold thermal energy storage (TES) subplant 212. Subplants 202-212 consume resources (e.g., water, natural gas, electricity, etc.) from utilities to serve the thermal energy loads (e.g., hot water, cold water, heating, cooling, etc.) of a building or campus. For example, heater subplant 202 can be configured to heat water in a hot water loop 214 that circulates the hot water between heater subplant 202 and building 10. Chiller subplant 206 can be configured to chill water in a cold water loop 216 that circulates the cold water between chiller subplant 206 and the building 10. Heat recovery chiller subplant 204 can be configured to transfer heat from cold water loop 216 to hot water loop 214 to provide additional heating for the hot water and additional cooling for the cold water. Condenser water loop 218 can absorb heat from the cold water in chiller subplant 206 and reject the absorbed heat in cooling tower subplant 208 or transfer the absorbed heat to hot water loop 214. Hot TES subplant 210 and cold TES subplant 212 can store hot and cold thermal energy, respectively, for subsequent use.

Hot water loop 214 and cold water loop 216 can deliver the heated and/or chilled water to air handlers located on the rooftop of building 10 (e.g., AHU 106) or to individual floors or zones of building 10 (e.g., VAV units 116). The air handlers push air past heat exchangers (e.g., heating coils or cooling coils) through which the water flows to provide heating or cooling for the air. The heated or cooled air can be delivered to individual zones of building 10 to serve the thermal energy loads of building 10. The water then returns to subplants 202-212 to receive further heating or cooling.

Although subplants 202-212 are shown and described as heating and cooling water for circulation to a building, it is understood that any other type of working fluid (e.g., glycol, CO₂, etc.) can be used in place of or in addition to water to serve the thermal energy loads. In other embodiments, subplants 202-212 can provide heating and/or cooling directly to the building or campus without requiring an intermediate heat transfer fluid. These and other variations to waterside system 200 are within the teachings of the present invention.

Each of subplants 202-212 can include a variety of equipment configured to facilitate the functions of the subplant. For example, heater subplant 202 is shown to include a plurality of heating elements 220 (e.g., boilers, electric heaters, etc.) configured to add heat to the hot water in hot water loop 214. Heater subplant 202 is also shown to include

several pumps 222 and 224 configured to circulate the hot water in hot water loop 214 and to control the flow rate of the hot water through individual heating elements 220. Chiller subplant 206 is shown to include a plurality of chillers 232 configured to remove heat from the cold water in cold water loop 216. Chiller subplant 206 is also shown to include several pumps 234 and 236 configured to circulate the cold water in cold water loop 216 and to control the flow rate of the cold water through individual chillers 232.

Heat recovery chiller subplant 204 is shown to include a plurality of heat recovery heat exchangers 226 (e.g., refrigeration circuits) configured to transfer heat from cold water loop 216 to hot water loop 214. Heat recovery chiller subplant 204 is also shown to include several pumps 228 and 230 configured to circulate the hot water and/or cold water through heat recovery heat exchangers 226 and to control the flow rate of the water through individual heat recovery heat exchangers 226. Cooling tower subplant 208 is shown to include a plurality of cooling towers 238 configured to remove heat from the condenser water in condenser water loop 218. Cooling tower subplant 208 is also shown to include several pumps 240 configured to circulate the condenser water in condenser water loop 218 and to control the flow rate of the condenser water through individual cooling towers 238.

Hot TES subplant 210 is shown to include a hot TES tank 242 configured to store the hot water for later use. Hot TES subplant 210 can also include one or more pumps or valves configured to control the flow rate of the hot water into or out of hot TES tank 242. Cold TES subplant 212 is shown to include cold TES tanks 244 configured to store the cold water for later use. Cold TES subplant 212 can also include one or more pumps or valves configured to control the flow rate of the cold water into or out of cold TES tanks 244.

In some embodiments, one or more of the pumps in waterside system 200 (e.g., pumps 222, 224, 228, 230, 234, 236, and/or 240) or pipelines in waterside system 200 include an isolation valve associated therewith. Isolation valves can be integrated with the pumps or positioned upstream or downstream of the pumps to control the fluid flows in waterside system 200. In various embodiments, waterside system 200 can include more, fewer, or different types of devices and/or subplants based on the particular configuration of waterside system 200 and the types of loads served by waterside system 200.

Referring now to FIG. 3, a block diagram of an airside system 300 is shown, according to an exemplary embodiment. In various embodiments, airside system 300 can supplement or replace airside system 130 in HVAC system 100 or can be implemented separate from HVAC system 100. When implemented in HVAC system 100, airside system 300 can include a subset of the HVAC devices in HVAC system 100 (e.g., AHU 106, VAV units 116, ducts 112-114, fans, dampers, etc.) and can be located in or around building 10. Airside system 300 can operate to heat or cool an airflow provided to building 10 using a heated or chilled fluid provided by waterside system 200.

In FIG. 3, airside system 300 is shown to include an economizer-type air handling unit (AHU) 302. Economizer-type AHUs vary the amount of outside air and return air used by the air handling unit for heating or cooling. For example, AHU 302 can receive return air 304 from building zone 306 via return air duct 308 and can deliver supply air 310 to building zone 306 via supply air duct 312. In some embodiments, AHU 302 is a rooftop unit located on the roof of building 10 (e.g., AHU 106 as shown in FIG. 1) or otherwise positioned to receive return air 304 and outside air 314.

AHU 302 can be configured to operate an exhaust air damper 316, mixing damper 318, and outside air damper 320 to control an amount of outside air 314 and return air 304 that combine to form supply air 310. Any return air 304 that does not pass through mixing damper 318 can be exhausted from AHU 302 through exhaust damper 316 as exhaust air 322.

Each of dampers 316-320 can be operated by an actuator. For example, exhaust air damper 316 can be operated by actuator 324, mixing damper 318 can be operated by actuator 326, and outside air damper 320 can be operated by actuator 328. Actuators 324-328 can communicate with an AHU controller 330 via a communications link 332. Actuators 324-328 can receive control signals from AHU controller 330 and can provide feedback signals to AHU controller 330. Feedback signals can include, for example, an indication of a current actuator or damper position, an amount of torque or force exerted by the actuator, diagnostic information (e.g., results of diagnostic tests performed by actuators 324-328), status information, commissioning information, configuration settings, calibration data, and/or other types of information or data that can be collected, stored, or used by actuators 324-328. AHU controller 330 can be an economizer controller configured to use one or more control algorithms (e.g., state-based algorithms, extremum seeking control (ESC) algorithms, proportional-integral (PI) control algorithms, proportional-integral-derivative (PID) control algorithms, model predictive control (MPC) algorithms, feedback control algorithms, etc.) to control actuators 324-328.

Still referring to FIG. 3, AHU 302 is shown to include a cooling coil 334, a heating coil 336, and a fan 338 positioned within supply air duct 312. Fan 338 can be configured to force supply air 310 through cooling coil 334 and/or heating coil 336 and provide supply air 310 to building zone 306. AHU controller 330 can communicate with fan 338 via communications link 340 to control a flow rate of supply air 310. In some embodiments, AHU controller 330 controls an amount of heating or cooling applied to supply air 310 by modulating a speed of fan 338.

Cooling coil 334 can receive a chilled fluid from waterside system 200 (e.g., from cold water loop 216) via piping 342 and can return the chilled fluid to waterside system 200 via piping 344. Valve 346 can be positioned along piping 342 or piping 344 to control a flow rate of the chilled fluid through cooling coil 334. In some embodiments, cooling coil 334 includes multiple stages of cooling coils that can be independently activated and deactivated (e.g., by AHU controller 330, by BMS controller 366, etc.) to modulate an amount of cooling applied to supply air 310.

Heating coil 336 can receive a heated fluid from waterside system 200 (e.g., from hot water loop 214) via piping 348 and can return the heated fluid to waterside system 200 via piping 350. Valve 352 can be positioned along piping 348 or piping 350 to control a flow rate of the heated fluid through heating coil 336. In some embodiments, heating coil 336 includes multiple stages of heating coils that can be independently activated and deactivated (e.g., by AHU controller 330, by BMS controller 366, etc.) to modulate an amount of heating applied to supply air 310.

Each of valves 346 and 352 can be controlled by an actuator. For example, valve 346 can be controlled by actuator 354 and valve 352 can be controlled by actuator 356. Actuators 354-356 can communicate with AHU controller 330 via communications links 358-360. Actuators 354-356 can receive control signals from AHU controller 330 and can provide feedback signals to controller 330. In

some embodiments, AHU controller 330 receives a measurement of the supply air temperature from a temperature sensor 362 positioned in supply air duct 312 (e.g., downstream of cooling coil 334 and/or heating coil 336). AHU controller 330 can also receive a measurement of the temperature of building zone 306 from a temperature sensor 364 located in building zone 306.

In some embodiments, AHU controller 330 operates valves 346 and 352 via actuators 354-356 to modulate an amount of heating or cooling provided to supply air 310 (e.g., to achieve a set-point temperature for supply air 310 or to maintain the temperature of supply air 310 within a set-point temperature range). The positions of valves 346 and 352 affect the amount of heating or cooling provided to supply air 310 by cooling coil 334 or heating coil 336 and may correlate with the amount of energy consumed to achieve a desired supply air temperature. AHU controller 330 can control the temperature of supply air 310 and/or building zone 306 by activating or deactivating coils 334-336, adjusting a speed of fan 338, or a combination thereof.

Still referring to FIG. 3, airside system 300 is shown to include a building management system (BMS) controller 366 and a client device 368. BMS controller 366 can include one or more computer systems (e.g., servers, supervisory controllers, subsystem controllers, etc.) that serve as system level controllers, application or data servers, head nodes, or master controllers for airside system 300, waterside system 200, HVAC system 100, and/or other controllable systems that serve building 10. BMS controller 366 can communicate with multiple downstream building systems or subsystems (e.g., HVAC system 100, a security system, a lighting system, waterside system 200, etc.) via a communications link 370 according to like or disparate protocols (e.g., LON, BACnet, etc.). In various embodiments, AHU controller 330 and BMS controller 366 can be separate (as shown in FIG. 3) or integrated. In an integrated implementation, AHU controller 330 can be a software module configured for execution by a processor of BMS controller 366.

In some embodiments, AHU controller 330 receives information from BMS controller 366 (e.g., commands, setpoints, operating boundaries, etc.) and provides information to BMS controller 366 (e.g., temperature measurements, valve or actuator positions, operating statuses, diagnostics, etc.). For example, AHU controller 330 can provide BMS controller 366 with temperature measurements from temperature sensors 362-364, equipment on/off states, equipment operating capacities, and/or any other information that can be used by BMS controller 366 to monitor or control a variable state or condition within building zone 306.

Client device 368 can include one or more human-machine interfaces or client interfaces (e.g., graphical user interfaces, reporting interfaces, text-based computer interfaces, client-facing web services, web servers that provide pages to web clients, etc.) for controlling, viewing, or otherwise interacting with HVAC system 100, its subsystems, and/or devices. Client device 368 can be a computer workstation, a client terminal, a remote or local interface, or any other type of user interface device. Client device 368 can be a stationary terminal or a mobile device. For example, client device 368 can be a desktop computer, a computer server with a user interface, a laptop computer, a tablet, a smartphone, a PDA, or any other type of mobile or non-mobile device. Client device 368 can communicate with BMS controller 366 and/or AHU controller 330 via communications link 372.

Referring now to FIG. 4, a block diagram of a building management system (BMS) 400 is shown, according to an

exemplary embodiment. BMS 400 can be implemented in building 10 to automatically monitor and control various building functions. BMS 400 is shown to include BMS controller 366 and a plurality of building subsystems 428. Building subsystems 428 are shown to include a building electrical subsystem 434, an information communication technology (ICT) subsystem 436, a security subsystem 438, a HVAC subsystem 440, a lighting subsystem 442, a lift/escalators subsystem 432, and a fire safety subsystem 430. In various embodiments, building subsystems 428 can include fewer, additional, or alternative subsystems. For example, building subsystems 428 can also or alternatively include a refrigeration subsystem, an advertising or signage subsystem, a cooking subsystem, a vending subsystem, a printer or copy service subsystem, or any other type of building subsystem that uses controllable equipment and/or sensors to monitor or control building 10. In some embodiments, building subsystems 428 include waterside system 200 and/or airside system 300, as described with reference to FIGS. 2-3.

Each of building subsystems 428 can include any number of devices, controllers, and connections for completing its individual functions and control activities. HVAC subsystem 440 can include many of the same components as HVAC system 100, as described with reference to FIGS. 1-3. For example, HVAC subsystem 440 can include a chiller, a boiler, any number of air handling units, economizers, field controllers, supervisory controllers, actuators, temperature sensors, and other devices for controlling the temperature, humidity, airflow, or other variable conditions within building 10. Lighting subsystem 442 can include any number of light fixtures, ballasts, lighting sensors, dimmers, or other devices configured to controllably adjust the amount of light provided to a building space. Security subsystem 438 can include occupancy sensors, video surveillance cameras, digital video recorders, video processing servers, intrusion detection devices, access control devices (e.g., card access, etc.) and servers, or other security-related devices.

Still referring to FIG. 4, BMS controller 366 is shown to include a communications interface 407 and a BMS interface 409. Interface 407 can facilitate communications between BMS controller 366 and external applications (e.g., monitoring and reporting applications 422, enterprise control applications 426, remote systems and applications 444, applications residing on client devices 448, etc.) for allowing user control, monitoring, and adjustment to BMS controller 366 and/or subsystems 428. Interface 407 can also facilitate communications between BMS controller 366 and client devices 448. BMS interface 409 can facilitate communications between BMS controller 366 and building subsystems 428 (e.g., HVAC, lighting security, lifts, power distribution, business, etc.).

Interfaces 407, 409 can be or include wired or wireless communications interfaces (e.g., jacks, antennas, transmitters, receivers, transceivers, wire terminals, etc.) for conducting data communications with building subsystems 428 or other external systems or devices. In various embodiments, communications via interfaces 407, 409 can be direct (e.g., locally wired or wireless communications) or via a communications network 446 (e.g., a WAN, the Internet, a cellular network, etc.). For example, interfaces 407, 409 can include an Ethernet card and port for sending and receiving data via an Ethernet-based communications link or network. In another example, the interfaces 407, 409 can include a Wi-Fi transceiver for communicating via a wireless communications network. In another example, one or more of interfaces 407, 409 can include cellular or mobile phone

communications transceivers. In one embodiment, communications interface 407 is a power line communications interface and BMS interface 409 is an Ethernet interface. In other embodiments, communications interface 407 and BMS interface 409 are Ethernet interfaces or are the same Ethernet interface.

Still referring to FIG. 4, BMS controller 366 is shown to include a processing circuit 404 including a processor 406 and memory 408. Processing circuit 404 can be communicably connected to BMS interface 409 and/or communications interface 407 such that processing circuit 404 and the various components thereof can send and receive data via interfaces 407, 409. Processor 406 can be implemented as a general purpose processor, an application specific integrated circuit (ASIC), one or more field programmable gate arrays (FPGAs), a group of processing components, or other suitable electronic processing components.

Memory 408 (e.g., memory, memory unit, storage device, etc.) can include one or more devices (e.g., RAM, ROM, Flash memory, hard disk storage, etc.) for storing data and/or computer code for completing or facilitating the various processes, layers and modules described in the present application. Memory 408 can be or include volatile memory or non-volatile memory. Memory 408 can include database components, object code components, script components, or any other type of information structure for supporting the various activities and information structures described in the present application. According to an exemplary embodiment, memory 408 is communicably connected to processor 406 via processing circuit 404 and includes computer code for executing (e.g., by processing circuit 404 and/or processor 406) one or more processes described herein.

In some embodiments, BMS controller 366 is implemented within a single computer (e.g., one server, one housing, etc.). In various other embodiments BMS controller 366 can be distributed across multiple servers or computers (e.g., that can exist in distributed locations). Further, while FIG. 4 shows applications 422 and 426 as existing outside of BMS controller 366, in some embodiments, applications 422 and 426 can be hosted within BMS controller 366 (e.g., within memory 408).

Still referring to FIG. 4, memory 408 is shown to include an enterprise integration layer 410, an automated measurement and validation (AM&V) layer 412, a demand response (DR) layer 414, a fault detection and diagnostics (FDD) layer 416, an integrated control layer 418, and a building subsystem integration later 420. Layers 410-420 can be configured to receive inputs from building subsystems 428 and other data sources, determine optimal control actions for building subsystems 428 based on the inputs, generate control signals based on the optimal control actions, and provide the generated control signals to building subsystems 428. The following paragraphs describe some of the general functions performed by each of layers 410-420 in BMS 400.

Enterprise integration layer 410 can be configured to serve clients or local applications with information and services to support a variety of enterprise-level applications. For example, enterprise control applications 426 can be configured to provide subsystem-spanning control to a graphical user interface (GUI) or to any number of enterprise-level business applications (e.g., accounting systems, user identification systems, etc.). Enterprise control applications 426 can also or alternatively be configured to provide configuration GUIs for configuring BMS controller 366. In yet other embodiments, enterprise control applications 426 can work with layers 410-420 to optimize building

performance (e.g., efficiency, energy use, comfort, or safety) based on inputs received at interface 407 and/or BMS interface 409.

Building subsystem integration layer 420 can be configured to manage communications between BMS controller 366 and building subsystems 428. For example, building subsystem integration layer 420 can receive sensor data and input signals from building subsystems 428 and provide output data and control signals to building subsystems 428. Building subsystem integration layer 420 can also be configured to manage communications between building subsystems 428. Building subsystem integration layer 420 translate communications (e.g., sensor data, input signals, output signals, etc.) across a plurality of multi-vendor/multi-protocol systems.

Demand response layer 414 can be configured to optimize resource usage (e.g., electricity use, natural gas use, water use, etc.) and/or the monetary cost of such resource usage in response to satisfy the demand of building 10. The optimization can be based on time-of-use prices, curtailment signals, energy availability, or other data received from utility providers, distributed energy generation systems 424, from energy storage 427 (e.g., hot TES 242, cold TES 244, etc.), or from other sources. Demand response layer 414 can receive inputs from other layers of BMS controller 366 (e.g., building subsystem integration layer 420, integrated control layer 418, etc.). The inputs received from other layers can include environmental or sensor inputs such as temperature, carbon dioxide levels, relative humidity levels, air quality sensor outputs, occupancy sensor outputs, room schedules, and the like. The inputs can also include inputs such as electrical use (e.g., expressed in kWh), thermal load measurements, pricing information, projected pricing, smoothed pricing, curtailment signals from utilities, and the like.

According to an exemplary embodiment, demand response layer 414 includes control logic for responding to the data and signals it receives. These responses can include communicating with the control algorithms in integrated control layer 418, changing control strategies, changing setpoints, or activating/deactivating building equipment or subsystems in a controlled manner. Demand response layer 414 can also include control logic configured to determine when to utilize stored energy. For example, demand response layer 414 can determine to begin using energy from energy storage 427 just prior to the beginning of a peak use hour.

In some embodiments, demand response layer 414 includes a control module configured to actively initiate control actions (e.g., automatically changing setpoints) which minimize energy costs based on one or more inputs representative of or based on demand (e.g., price, a curtailment signal, a demand level, etc.). In some embodiments, demand response layer 414 uses equipment models to determine an optimal set of control actions. The equipment models can include, for example, thermodynamic models describing the inputs, outputs, and/or functions performed by various sets of building equipment. Equipment models can represent collections of building equipment (e.g., sub-plants, chiller arrays, etc.) or individual devices (e.g., individual chillers, heaters, pumps, etc.).

Demand response layer 414 can further include or draw upon one or more demand response policy definitions (e.g., databases, XML files, etc.). The policy definitions can be edited or adjusted by a user (e.g., via a graphical user interface) so that the control actions initiated in response to demand inputs can be tailored for the user's application, desired comfort level, particular building equipment, or

based on other concerns. For example, the demand response policy definitions can specify which equipment can be turned on or off in response to particular demand inputs, how long a system or piece of equipment should be turned off, what setpoints can be changed, what the allowable set point adjustment range is, how long to hold a high demand set-point before returning to a normally scheduled set-point, how close to approach capacity limits, which equipment modes to utilize, the energy transfer rates (e.g., the maximum rate, an alarm rate, other rate boundary information, etc.) into and out of energy storage devices (e.g., thermal storage tanks, battery banks, etc.), and when to dispatch on-site generation of energy (e.g., via fuel cells, a motor generator set, etc.).

Integrated control layer 418 can be configured to use the data input or output of building subsystem integration layer 420 and/or demand response layer 414 to make control decisions. Due to the subsystem integration provided by building subsystem integration layer 420, integrated control layer 418 can integrate control activities of the subsystems 428 such that the subsystems 428 behave as a single integrated supersystem. In an exemplary embodiment, integrated control layer 418 includes control logic that uses inputs and outputs from a plurality of building subsystems to provide greater comfort and energy savings relative to the comfort and energy savings that separate subsystems could provide alone. For example, integrated control layer 418 can be configured to use an input from a first subsystem to make an energy-saving control decision for a second subsystem. Results of these decisions can be communicated back to building subsystem integration layer 420.

Integrated control layer 418 is shown to be logically below demand response layer 414. Integrated control layer 418 can be configured to enhance the effectiveness of demand response layer 414 by enabling building subsystems 428 and their respective control loops to be controlled in coordination with demand response layer 414. This configuration may advantageously reduce disruptive demand response behavior relative to conventional systems. For example, integrated control layer 418 can be configured to assure that a demand response-driven upward adjustment to the set-point for chilled water temperature (or another component that directly or indirectly affects temperature) does not result in an increase in fan energy (or other energy used to cool a space) that would result in greater total building energy use than was saved at the chiller.

Integrated control layer 418 can be configured to provide feedback to demand response layer 414 so that demand response layer 414 checks that constraints (e.g., temperature, lighting levels, etc.) are properly maintained even while demanded load shedding is in progress. The constraints can also include set-point or sensed boundaries relating to safety, equipment operating limits and performance, comfort, fire codes, electrical codes, energy codes, and the like. Integrated control layer 418 is also logically below fault detection and diagnostics layer 416 and automated measurement and validation layer 412. Integrated control layer 418 can be configured to provide calculated inputs (e.g., aggregations) to these higher levels based on outputs from more than one building subsystem.

Automated measurement and validation (AM&V) layer 412 can be configured to verify that control strategies commanded by integrated control layer 418 or demand response layer 414 are working properly (e.g., using data aggregated by AM&V layer 412, integrated control layer 418, building subsystem integration layer 420, FDD layer 416, or otherwise). The calculations made by AM&V layer

412 can be based on building system energy models and/or equipment models for individual BMS devices or subsystems. For example, AM&V layer 412 can compare a model-predicted output with an actual output from building subsystems 428 to determine an accuracy of the model.

Fault detection and diagnostics (FDD) layer 416 can be configured to provide on-going fault detection for building subsystems 428, building subsystem devices (i.e., building equipment), and control algorithms used by demand response layer 414 and integrated control layer 418. FDD layer 416 can receive data inputs from integrated control layer 418, directly from one or more building subsystems or devices, or from another data source. FDD layer 416 can automatically diagnose and respond to detected faults. The responses to detected or diagnosed faults can include providing an alert message to a user, a maintenance scheduling system, or a control algorithm configured to attempt to repair the fault or to work-around the fault.

FDD layer 416 can be configured to output a specific identification of the faulty component or cause of the fault (e.g., loose damper linkage) using detailed subsystem inputs available at building subsystem integration layer 420. In other exemplary embodiments, FDD layer 416 is configured to provide "fault" events to integrated control layer 418 which executes control strategies and policies in response to the received fault events. According to an exemplary embodiment, FDD layer 416 (or a policy executed by an integrated control engine or business rules engine) can shut-down systems or direct control activities around faulty devices or systems to reduce energy waste, extend equipment life, or ensure proper control response.

FDD layer 416 can be configured to store or access a variety of different system data stores (or data points for live data). FDD layer 416 can use some content of the data stores to identify faults at the equipment level (e.g., specific chiller, specific AHU, specific terminal unit, etc.) and other content to identify faults at component or subsystem levels. For example, building subsystems 428 can generate temporal (i.e., time-series) data indicating the performance of BMS 400 and the various components thereof. The data generated by building subsystems 428 can include measured or calculated values that exhibit statistical characteristics and provide information about how the corresponding system or process (e.g., a temperature control process, a flow control process, etc.) is performing in terms of error from its set-point. These processes can be examined by FDD layer 416 to expose when the system begins to degrade in performance and alert a user to repair the fault before it becomes more severe.

Ethernet Ring Topology Generation

The BMS, as described above, can have multiple individual components within the BMS. Example components may include control devices, such as field equipment controllers (FECs), advanced application field equipment controllers (FACs), network control engines (NCEs), input/output modules (IOMs), and variable air volume (VAV) modular assemblies. However, other control device types are contemplated. For example, the BMS may include multiple devices such as sensors, actuators, valves, beacons, switches, thermostats, etc. In some embodiments, some or many of these devices may be configured as or connected to a network and controlled through that network. For example, the network may be an Ethernet network arranged in a ring topology. The present disclosure describes a mechanism to dynamically and efficiently determine the topology of the devices within the ring, thereby allowing for the system configuration to be easily verified automatically. Further, the

mechanism for determining the topology of the devices within the ring can allow for more efficient identification of a malfunctioning device within the Ethernet ring when a fault is detected.

Referring now to FIG. 5, a block diagram illustrating a networked system 500 is shown. The system includes an Ethernet ring topology generator (RTG) 502, a ring supervisor 504, and a number of devices 506-518. As described above, the devices 506-518 may be any type of devices within a BMS that are connected via an Ethernet connection. The RTG 502 may be a remote server, a personal computer, a laptop computer, or other device configured to operate as the topology generator. In one embodiment, the RTG 502 is a Metasys server from Johnson Controls, Inc. In some embodiments, the RTG 502 must be a device outside of an Ethernet ring in order to accurately generate the topology of the devices within the Ethernet ring. The RTG 502 is configured to transmit one or more commands to the devices 506-518 in the Ethernet ring via the ring supervisor 504. The RTG 502 may further be configured to generate and manage one or more topology tables associated with an Ethernet ring, and which may be used to determine an order and configuration of the devices 506-518 arranged in an Ethernet ring 520. In some examples, the RTG 502 may also store the topology information in lists, arrays or other data structures, as applicable. In one embodiment, to determine the topology of the devices 506-518 of the devices within the Ethernet ring 520, the RTG 502 must determine the following information: the devices in the ring, the order of the devices in the ring relative to each other; and how the devices are connected to each other (i.e., which Ethernet ports connect two devices in the ring, or which Ethernet ports of a device connects the device to the ring supervisor 504. The processes for determining the above information will be described in more detail below.

The ring supervisor 504 may be standard Ethernet switch configured to manage a number of Ethernet-based devices in ring topology. In one embodiment, the Ring supervisor 504 is configured to host a Dynamic Host Control Protocol (DHCP) server and a MRP ring manager. The ring supervisor 504 may have a number of Ethernet ports for communicating with one or more Ethernet devices or networks. In one embodiment, the ring supervisor 504 uses two ports, such as port A and port B to communicate with the devices 506-518 and generate an Ethernet ring. As stated above, the Ethernet ring 520 of system 500 may be configured as an MRP ring. However, the herein disclosed topology determination methods and systems are intended to be protocol agnostic, and could be used in various other Ethernet ring protocols, such as a spanning tree protocol (STP) rings, daisy chain configurations, and/or other Ethernet ring protocols. In an MRP ring, an MRP beacon frame may be transmitted to each of the devices 506-518 in the Ethernet ring 520 to inform the devices 506-518 that they are a member of an MRP ring. The devices 506-518, upon receiving the MRP beacon flag, flag themselves as being a part of an MRP ring. In an MRP ring, one device in the ring (e.g. the ring supervisor 504) manages the Ethernet ring 520 to prevent continuous/infinite looping of messages within the Ethernet ring 520.

Each of the devices 506-518 may include a first port and a second port. For example, device A 506 is shown to include Port 1 and Port 2. One port may be designated as a receiving port and the other port may be a forwarding port. In one example, when the ring is closed (e.g. there is no break in the connectivity in the ring), the ring supervisor 504 may block one of the ports to prevent infinite looping within

the ring and forwards all packets received outside the ring to the device on the ring. However, both Port 1 and Port 2 may be capable of sending and receiving data packets, as needed. For example, for device A 506, Port 1 is shown as a receiving port, and port 2 is shown as a forwarding port. However, the ports (e.g. port 1 and port 2) may be interchanged between being a forwarding port and a receiving port based on which port on the ring supervisor 504 is being used as a forwarding port. For example, in the system 500, if Port 4 of the Ring supervisor is used as the forwarding port of the ring supervisor 504, then Port 2 would be the receiving port and Port 1 would be the forwarding port of device A 506. In a typical ring topology, one port of the ring supervisor 504 is set as the forwarding port to prevent redundant messages from being provided to the Ethernet ring 520.

Turning now to FIG. 6, a data flow chart illustrating a process 600 for discovering the devices in an Ethernet ring is shown, according to some embodiments. The process 600 is described in reference to the system 500 described above. However, it is contemplated that the process 600 can be used with various other Ethernet networks, as described above. As stated above, before the topology of the devices 506-518 can be determined, the RTG 502 needs to discover the devices 506-518 that are participating in the Ethernet ring 520. In some embodiments, the discovery process 600 includes discovering the devices 506-518 as well as the ring supervisor 504. More specifically, the discovery process 600 determines which ports on the ring supervisor 504 are hosting the Ethernet ring 520. In some embodiments, the RTG 502 can be pre-configured with basic information about the Ethernet ring 520. For example, the RTG 502 may be configured with the IP address of the ring supervisor 504 hosting the Ethernet ring 520, and the type of ring (MRP, STP, etc.) being used. This information can be used to determine which ring supervisor 504 commands need to be issued to obtain a status of the Ethernet ring 520. Additionally, the RTG 502 may know the switch ports of the ring supervisor 504 used to operate the ring Ethernet 520, and a subnetwork of the devices in the Ethernet ring 520. In some embodiments, a user may provide this information to the RTG 502. In other embodiments, the RTG may interrogate the Ring supervisor 504 to obtain the necessary information.

The RTG 502 first transmits a ring status query 602 to determine if the Ethernet ring 520 is open or closed. If the ring is closed, the RTG 502 determines that all devices 506-518 are reachable and broadcasts a device discovery broadcast message 604. Upon receiving the device discovery broadcast message 604, the devices 506-518 respond by transmitting a device discovery response 606. The device discovery response 606 may include a device ID, and a device IP address.

Upon receiving a device discovery response 606 from each of the devices 506-518, the RTG 502 may create a ring topology data structure and populate the data structure with the information within the received device discovery responses 606. In one embodiment, the data structure is a Ring Topology Table, such as Table 1, shown below.

TABLE 1

Ring Topology Table with Device Discovery Response Data			
Ring Topology Table			
Ethernet Supervisor Port/Ring Device ID	Port to Previous Ring Device	Port to Next Ring Device	Ring Order Count
Supervisor Ring Port	N/A		N/A
Device C			
Device F			

TABLE 1-continued

Ring Topology Table with Device Discovery Response Data			
Ring Topology Table			
Ethernet Supervisor Port/Ring Device ID	Port to Previous Ring Device	Port to Next Ring Device	Ring Order Count
Device B			
Device G			
Device D			
Device A			
Device E			
Supervisor Ring Port		N/A	N/A

After the devices 506-518 of the Ethernet ring are discovered, the RTG 502 initiates an Ethernet port orientation process 700, as shown in FIG. 7 to orient the ports of the devices 506-518 relative to each other. In one embodiment, the RTG 502 orients each device's 506-518 ports relative to the ring supervisor's 504 forwarding ring port. In order to prevent infinite looping when the ring is closed, the ring supervisor 504 may block one of its ring ports (Port 3 or Port 4) such that data from outside the Ethernet ring 520 is always routed into the Ethernet ring 520 through the other ring port (the forwarding port). Therefore, all of the device 506-518 will receive messages on their data port which is closest to the ring supervisor's 504 forwarding port.

In one embodiment, the RTG 502 first transmits a ring status query 702 to the ring supervisor 504 to determine if the Ethernet ring 520 is open or closed. If the ring is open, the RTG 502 may abort the Ethernet port orientation process 700. If the ring is closed, the RTG 502 may transmit a ring port status query 704 to the ring supervisor 504 to determine which port is currently configured as the forwarding port of the ring supervisor 504, and therefore which port of the ring supervisor 504 is blocked. In one embodiment, the ring status query 702 and the ring port status query 704 may be combined into a single query. For purposes of this example, the forwarding port is determined to be Port 3, and the blocked port is determined to be Port 4.

The RTG 502 may then broadcast a ring port request message 706 to all of the devices 506-518. The devices 506-518 may then respond back to the RTG 502 with a ring port response message 708 containing their Device ID, and the port on which they received the ring port request message (e.g. port 1 or port 2). In one example, messages from the RTG 502 will be received on the Ethernet port closest to the forwarding port of the ring supervisor 504 for each of the devices 506-518 where the Ethernet ring 520 is closed. The port receiving the ring port request message 706 therefore is either directly connected to the forwarding port of the ring supervisor 504 (port 3), or is connected to the previous device 506-518 in the ring relative to the forwarding port of the ring supervisor 504 (port 3). Therefore, the device's 506-518 other port is therefore connected to either the blocked port of the ring supervisor 504 (port 4) or to the next device 506-518 in the Ethernet ring 520, according to some embodiments.

The RTG 502 may then update the Ring Topology table with the data received in the ring port response messages 708 received from the devices 506-518. In one embodiment, the RTG 502 may populate the Port to Previous Ring Device with the Ethernet port (1 or 2) of each device 504-518 based on the data received in the ring port response message 708. The Port to Previous Ring Device data points indicate which

port of a given device **506-518** is closest to the ring supervisor's **504** forwarding port. The RTG **502** may also populate the Port to Next Ring Device with the Ethernet port (1 or 2) of each device **506-518** based on the data received in the ring port response message **708**. The Port to Next Ring Device data points indicate which port of a device **506-518** is closest to the Ethernet ring's **520** blocked port (port 4). An example updated Ring Topology Table is shown in Table 2, below.

TABLE 2

Ring Topology Table with Device Port Data Ring Topology Table			
Ethernet Supervisor Port/Ring Device ID	Port to Previous Ring Device	Port to Next Ring Device	Ring Order Count
Supervisor Ring Port	N/A	3	N/A
Device C	2	1	0
Device F	1	2	0
Device B	1	2	0
Device G	2	1	0
Device D	1	2	0
Device A	1	2	0
Device E	1	2	0
Supervisor Ring Port	4	N/A	N/A

Once the RTG **502** has determined the orientation of each of the devices' **506-518** ports relative to the ring supervisor **504**, the RTG **502** can determine an order of the devices in the ring relative to the ring supervisor's **504** forwarding port. Turning now to FIG. **8**, a data flow chart illustrating a process **800** for determining a position of devices in an Ethernet ring is shown, according to some embodiments. The RTG **502** first transmits a ring status query **802** to the ring supervisor **504** to determine if the Ethernet ring **520** is open or closed. If the Ethernet ring **520** is determined to be open, the RTG **502** aborts the process **800**. If the Ethernet ring **520** is determined to be closed, the RTG **502** then sends a port status change request message **804** to a first device **506-518** of the Ethernet ring **520**. In some embodiments, the RTG **502** may transmit the port status change request message **804** to the device which is listed first in the generated ring topology table. However, in other embodiments, the RTG **502** may transmit the port status change request message to a device based on other criteria.

In one embodiment, the port status change request message **804** includes a command to disable one ports of the device **506-518** receiving the message **804**. In one embodiment, the port status change request message **804** instructs the device **506-518** receiving the message **804** to disable the port associated with the port to next ring device parameter in the ring topology table for the associated device **506-518**. As shown in FIG. **8**, device C **510** may be the first device to receive the port status change request message **804**. In some embodiments, the device **506-518** currently receiving the port status change request message **804** is referred to as the reference device. After the device C **510** disables a port, the device C **510** may transmit a port status change request message **806** to the RTG **502** indicating that the port has been disabled. Upon receiving the port status change request message **806**, the RTG **502** again transmits a ring status query **808** to the ring supervisor **504** to verify that the ring is now open. Upon verifying that the ring is open, the RTG **502** broadcasts a ring port request message **810** to devices **506-518** via the forwarding port of the ring supervisor **504**.

Each device **506-518** receiving the ring port request message **810** may transmit a ring port response message **812** to the RTG **502**.

Upon receiving the ring port response messages **812**, the RTG **502** first determines if the ring port response message **812** is from the reference device (e.g. the device with the currently disabled port). If the ring port response message **812** is from the reference device, the RTG **502** ignores the ring port response message **812**. For the ring port response messages **812** not from the reference device, the RTG **502** then determines if the ring port response message for each of the other non-reference devices **506-518** is received from the Ethernet port which matches the Port to Previous Ring Device parameter for each associated device in the ring topology table. If the ring port response message **812** is not from the reference device, and the Ethernet port from which the ring port response message **812** was received matches the port to previous ring device parameter within the ring topology table, the RTG **502** increments a ring order count parameter in the ring topology table for the associated device. As only the forwarding port of the ring supervisor **504** is active, the RTG **502** will only receive the ring port response messages from device **506-518** between the forwarding port of the ring supervisor **504** and the reference device. At this point the ring is open such that the ring supervisor **504** can send the ring port request message **810** over both its ring ports (e.g. port 1 and port 2). However, only the devices which receive the ring port request message **810** via the previously defined forwarding port (ie. The Port to Previous Ring Device port) will have their Ring Order Count incremented. Thus, as in the above example where the device C is the initial reference device, after all of the ring port response messages **812** have been received, the ring order count parameters values within the ring topology chart will be as shown in Table 3, below.

TABLE 3

Ring Topology Table with Initial Ring Order Count Ring Topology Table			
Ethernet Supervisor Port/Ring Device ID	Port to Previous Ring Device	Port to Next Ring Device	Ring Order Count
Supervisor Ring Port	N/A	3	N/A
Device C	2	1	0
Device F	1	2	0
Device B	1	2	1
Device G	2	1	0
Device D	1	2	0
Device A	1	2	1
Device E	1	2	0
Supervisor Ring Port	4	N/A	N/A

Once the ring port response messages **812** have been received, the RTG **502** may transmit a port status change request **814** to the reference device, device C **510**. The port status change request **814** can instruct the reference device, device C **510** to re-enable its Ethernet port. The reference device, device C **510**, can then transmit a port status change response **816** to the RTG **502** indicating that the port has been re-enabled. The port status change response **816** may include a message ID indicating that the message is a port status change response, an Ethernet port that was updated, and a result bit associated with the updated Ethernet port (e.g. 1—port enable, 2—port disabled.). Upon receiving the port status change response **816**, the RTG **502** may transmit a port status change request **818** to a subsequent device **506-518** in the ring that is different from the previous

reference device. After receiving confirmation that the device's Ethernet port has transitioned back to enabled, the RTG **502** may perform another Query Ring Status to confirm that the ring is now closed, and prior to moving on to the next device with **818**. The RTG **502** may then repeat the above process for the new reference device, and continue until each of the previously discovered devices **506-518** have been designated as the reference device (e.g. had a port disabled.). Upon completion of the process **800**, the ring topology table may include the following data as shown in Table 4 below, based on the system **500**.

TABLE 4

Ring Topology Table Including Complete Ring Order Count Data			
Ring Topology Table			
Ethernet Supervisor Port/Ring Device ID	Port to Previous Ring Device	Port to Next Ring Device	Ring Order Count
Supervisor Ring Port	N/A	3	N/A
Device C	2	1	4
Device F	1	2	1
Device B	1	2	5
Device G	2	1	0
Device D	1	2	3
Device A	1	2	6
Device E	1	2	2
Supervisor Ring Port	4	N/A	N/A

The RTG **502** may then order then devices **506-518** based on their respective ring order count. By ordering the devices **506-518** from highest ring order count to lowest ring order count, the devices **506-518** will be listed in the order they are installed in the ring in relation to the forwarding port of the ring supervisor **504**, as well as how the device **506-518** in the ring are connected to each other (e.g., the Port to Previous Ring Device for the first device in the ring topology table is connected to the ring manager forwarding port, the Port to Next Ring Device of the first device in the ring topology table is connected to the Port to Previous Ring Device of the second device in the ring topology table, and so on.). For the system **500** of FIG. **5**, the ordered ring topology table is shown below in Table 5.

TABLE 5

Ordered Ring Topology Table			
Ring Topology Table			
Ethernet Supervisor Port/Ring Device ID	Port to Previous Ring Device	Port to Next Ring Device	Ring Order Count
Supervisor Ring Port	N/A	3	N/A
Device A	1	2	6
Device B	1	2	5
Device C	2	1	4
Device D	1	2	3
Device E	1	2	2
Device F	1	2	1
Device G	2	1	0
Supervisor Ring Port	4	N/A	N/A

Turning now to FIG. **9**, a flow diagram illustrating a process **900** for determining a topology of an Ethernet ring is shown, according to some embodiments. The process **900** may be a combination of the processes **600**, **700** and **800** described above. In one embodiment, the process **900** operates at the application layer, allowing the topology of the Ethernet ring **520** to be determined by communicating directly with the devices in the Ethernet ring, rather than

relying on messaging type specific to the ring protocol to provide the topology information. In some embodiments, an RTG, such as RTG **502** described above, may be responsible for performing the various steps of the process **900**. In one example, the RTG may interface with a user, such as via a graphical user interface (GUI), to allow a user to both configure the RTG, but also to present the user with a graphical representation of the generated ring topology, in some embodiments. In one embodiment, when the RTG is initialized or launched, the RTG will initially generate the topology of the configured Ethernet ring as described below. Further, the RTG may continue to monitor the status of the ring to determine if/when the topology needs to be regenerated. For example, the RTG may regenerate the ring topology whenever the ring status transitions from open to closed to account for the addition/deletion of devices to the Ethernet ring as well as any cabling changes within the Ethernet ring (e.g. if the cables going into a particular device's Ethernet ports were swapped.)

At process block **902**, all the devices within the Ethernet ring are determined. In one embodiment, an RTG, such as RTG **502** may execute a process to discover all of the devices within the Ethernet ring. For example, the RTG may use a process such as process **600** described above to discover all the devices in the Ethernet ring. The RTG may also determine other information about the devices, such as IP addresses of the devices, device types, associated other devices, associated sub-networks, or the like.

At process block **904**, the RTG may determine a port configuration of the discovered device. In some embodiments, the RTG may determine the port configurations using a process similar to process **700** described above. For example, the RTG may broadcast a ring port request to the devices on the Ethernet ring. In some embodiments, the RTG first ensures that one of the ports of an Ethernet ring supervisor associated with the Ethernet ring is disabled, ensuring that any message broadcasted by the RTG is only provided to the Ethernet ring via a single port. By determining the port configurations of the discovered devices, the RTG can determine a relationship of the Ethernet ports of the devices in regards to each other, and an Ethernet ring supervisor, such as ring supervisor **504**. The RTG may then populate a data structure with the individual devices within the Ethernet ring, as well as the relationships of their individual ports. For example, the RTG may populate the data structure with data such as the "port to previous ring device," indicating which port the device received the broadcasted message from the RTG.

At process block **906**, the RTG may transmit an instruction to a first device within the Ethernet ring and instruct the first device to disable one of the Ethernet ports of the device. The device with the disabled port may be referred to as the reference device. The RTG may then determine if the Ethernet ring is open at process block **908**. In some embodiments, the RTG may verify that the Ethernet ring is open by querying the Ethernet ring supervisor, which can provide a status of the Ethernet ring (e.g. whether the Ethernet ring is open or closed). In other embodiments, the RTG may broadcast a message to the devices on the Ethernet ring and determine if a response is issued from all of the devices, or a portion of the devices, indicating whether the ring is open or closed.

Once the device determines that the Ethernet ring is open, the RTG broadcasts a ring port request to the devices in the Ethernet ring at process block **910**. The ring port request may include a message ID indicating that the message is a ring port request. The ring port request may be a standard

request in Ethernet ring configurations. At process block 912, the devices receiving the ring port request transmit a ring port response to the RTG. The ring port response may include a message ID, a device ID of the device transmitting the ring port response, a result and an Ethernet port message. The device ID may be a unique identifier for the device. For example, the device ID may be a BACnet OID, a logical name, a media access control (MAC) address, or some other existing ID. The result message may indicate whether the ring port response is successful or unsuccessful. Finally, the Ethernet port message may indicate on which port of the device the ring port request message was received. In one embodiment, the Ethernet port message may only be included in the ring port response where the result message indicates that the ring port response was successful. The RTG, having received the ring port responses from each of the devices receiving the ring port request increments a count associated with each device from which the ring port response is received. In one embodiment, the RTG may ignore the ring port response received from the reference device (e.g. the device with the disabled port.).

Once all of the ring port responses have been received, the RTG enables the previously disabled port on the reference device at process block 914. The RTG may further query the ring supervisor to confirm the ring is closed. The RTG then determines if all of the previously discovered devices within the Ethernet ring have had a port disabled (e.g. has each device been the reference device) at process block 916. If the RTG determines that not all of the previous devices have had an Ethernet port disabled, the process 900 returns to process block 906. If the RTG does determine that all of the devices have had their Ethernet ports disabled, the RTG may determine a topology of the Ethernet ring (e.g. determine the order of the discovered devices within the ring) at process block 918. In one embodiment, the RTG may determine the topology of the devices in the Ethernet ring by ordering the devices based on the number of ring port responses received from each device. The RTG may then determine that the device with the highest number of associated port responses is the device closest to the forwarding port (e.g. active port) of the Ethernet ring supervisor). Therefore, the device with the lowest number of associated port responses is the device further away from the forwarding port of the Ethernet ring supervisory. Accordingly, the order of the devices, based on the number of received port responses, will correspond to the position of the device within the Ethernet ring with respect to the forwarding port of the Ethernet ring supervisor.

In some embodiments, the RTG may utilize the generated topology to be able to identify the location of one or more breaks in an Ethernet ring. For example, whenever a ring device's Ethernet port goes out of service (OOS), the device could send a message to the RTG reporting the OOS port. The RTG could then map the port to the current ring topology and mark the port to indicate the OOS status of the port. In other embodiments, failure or break events should be generated in pairs (i.e., if one end of a connection between two devices within the Ethernet ring goes OOS, then the other should as well, such that the RTG should mark both ends of the connection OOS. If another break in the ring is detected, the connectivity to one or more devices in the ring will be lost. The RTG could determine the devices which have no connectivity based on the most recent ring topology and indicate to a user which devices have lost connection.

Configuration of Exemplary Embodiments

The construction and arrangement of the systems and methods as shown in the various exemplary embodiments are illustrative only. Although only a few embodiments have been described in detail in this disclosure, many modifications are possible (e.g., variations in sizes, dimensions, structures, shapes and proportions of the various elements, values of parameters, mounting arrangements, use of materials, colors, orientations, etc.). For example, the position of elements may be reversed or otherwise varied and the nature or number of discrete elements or positions may be altered or varied. Accordingly, all such modifications are intended to be included within the scope of the present disclosure. The order or sequence of any process or method steps may be varied or re-sequenced according to alternative embodiments. Other substitutions, modifications, changes, and omissions may be made in the design, operating conditions and arrangement of the exemplary embodiments without departing from the scope of the present disclosure.

The present disclosure contemplates methods, systems and program products on any machine-readable media for accomplishing various operations. The embodiments of the present disclosure may be implemented using existing computer processors, or by a special purpose computer processor for an appropriate system, incorporated for this or another purpose, or by a hardwired system. Embodiments within the scope of the present disclosure include program products comprising machine-readable media for carrying or having machine-executable instructions or data structures stored thereon. Such machine-readable media can be any available media that can be accessed by a general purpose or special purpose computer or other machine with a processor. By way of example, such machine-readable media can comprise RAM, ROM, EPROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to carry or store desired program code in the form of machine-executable instructions or data structures and which can be accessed by a general purpose or special purpose computer or other machine with a processor. Combinations of the above are also included within the scope of machine-readable media. Machine-executable instructions include, for example, instructions and data which cause a general purpose computer, special purpose computer, or special purpose processing machines to perform a certain function or group of functions.

Although the figures show a specific order of method steps, the order of the steps may differ from what is depicted. Also two or more steps may be performed concurrently or with partial concurrence. Such variation will depend on the software and hardware systems chosen and on designer choice. All such variations are within the scope of the disclosure. Likewise, software implementations could be accomplished with standard programming techniques with rule based logic and other logic to accomplish the various connection steps, processing steps, comparison steps and decision steps.

The invention claimed is:

1. A method of determining a topology of devices in an Ethernet ring, comprising:
 - broadcasting a device discovery command from a ring topology generator to the devices in the Ethernet ring;
 - receiving device discovery information from one or more of the devices in the Ethernet ring at the ring topology generator;

23

querying a ring supervisor to confirm that the Ethernet ring is closed and to determine which Ethernet port of the ring supervisor is connected to the Ethernet ring as a forwarding port;

the ring topology generator configured to, for each discovered device in the Ethernet ring:

- (i) disabling an Ethernet port of [one of the devices] *the discovered device* in the Ethernet ring;
- (ii) verifying that the Ethernet ring is open;
- (iii) broadcasting a ring port request onto the Ethernet ring via the ring supervisor;
- (iv) receiving a ring port response from each of the devices having received the ring port request;
- (v) incrementing a count associated with each [device based on receiving] *of the devices from which a ring port response [from the device] is received other than the discovered device*;
- (vi) re-enabling the Ethernet port of the [one of the devices] *discovered device* in the Ethernet ring; and
- (vii) repeating steps (i) through (vi) until each device in the Ethernet ring has had an Ethernet port disabled; and

determining the topology of the devices of the Ethernet ring by ordering the devices in the Ethernet ring based on a value of the count associated with each device, wherein the device having a highest count value being closest to the forwarding port of the ring supervisor, and the device having a lowest count value being furthest from the forwarding port of the ring supervisor.

2. The method of claim 1, wherein the ring topology generator does not increment a count for a ring port response from the *discovered* device that currently has the Ethernet port disabled.

3. The method of claim 1, wherein the Ethernet ring is configured as a Media Redundancy Protocol (MRP) ring.

4. The method of claim 1, wherein the device discovery information includes a device identification (ID).

5. The method of claim 4, wherein the device ID is one or more of a media access control (MAC) address, a building automation control network (BACnet) object identification (OID), and a logical device name.

6. The method of claim 1, further comprising:

broadcasting a ring port query message based on receiving the device discovery information from the one or more devices in the Ethernet ring;

receiving a ring port query response from the one or more devices in the Ethernet ring; and

wherein the ring topology generator is configured to orient the Ethernet ports of the one or more devices in the Ethernet ring based on the received ring port query responses.

7. The method of claim 6, wherein the ring port query response comprises a message identification (ID), a [Device] *device* identification (ID), and an Ethernet port on which the ring port query message was received.

8. A system for determining a topology of an Ethernet ring, the system comprising:

a plurality of Ethernet devices configured in the Ethernet ring;

a ring supervisor, the ring supervisor integral to the Ethernet ring and configured to manage data flow within the Ethernet ring; and

a ring topology generator configured to:

broadcast a device discovery command from a ring topology generator to the devices in the Ethernet ring;

24

receive device discovery information from one or more of the plurality of devices in the Ethernet ring, the received device discovery information indicating a discovered device;

query the ring supervisor to confirm that the Ethernet ring is closed and to determine which Ethernet port of the ring supervisor is connected to the Ethernet ring as a forwarding port;

execute *the* following instructions for each [of the] discovered [devices] *device*:

- (i) disable an Ethernet port of [one of] *the* discovered [devices] *device* in the Ethernet ring;
- (ii) verify that the Ethernet ring is open;
- (iii) broadcast a ring port request onto the Ethernet ring via the ring supervisor;
- (iv) receive a ring port response from each of the [discovered] devices having received the ring port request;
- (v) increment a count associated with each [device based on receiving] *of the devices from which a ring port [request from the device] response is received other than the discovered device*;
- (vi) re-enable the Ethernet port of [the one of] the discovered [devices] *device*; and
- (vii) repeat steps (i) through (vi) until each device in the Ethernet ring has had an Ethernet port disabled; and

determine the topology of the devices on the Ethernet ring by ordering the devices in the Ethernet ring based on a value of the count associated with each device, wherein the device having a highest count value is closest to the forwarding port of the ring supervisor, and the device having a lowest count value is the device furthest from the forwarding port of the ring supervisor.

9. The system of claim 8, wherein the ring topology generator does not increment a count for a ring port response from the *discovered* device that currently has the Ethernet port disabled.

10. The system of claim 8, wherein the Ethernet ring is configured as one of a Media Redundancy Protocol (MRP) ring and a Spanning Tree Protocol (STP) ring.

11. The system of claim 8, wherein the device discovery information comprises a device identification (ID), the device ID comprising one or more of a media access control (MAC) address, a building automation control network (BACnet) object identification (OID), or a logical device name.

12. The system of claim 8, wherein the ring topology generator is further configured to:

broadcast a ring port query message based on receiving the device discovery information from the one or more devices in the Ethernet ring;

receive a ring port query response from the one or more devices in the Ethernet ring; and

orient the Ethernet ports of the one or more devices in the Ethernet ring based on the received ring port query responses.

13. The system of claim 12, wherein the ring port query response comprises a message identification (ID), a [Device] *device* identification (ID), and an Ethernet port on which the ring port query message was received.

14. A method of determining a topology of devices in an Ethernet ring, comprising:

discovering a number and identification of devices within the Ethernet ring;

for each *discovered* device [discovered] in the Ethernet ring;

25

(i) disabling an Ethernet port of [one of the devices] *the discovered device* in the Ethernet ring;

(ii) verifying that the Ethernet ring is open;

(iii) broadcasting a ring port request onto the Ethernet ring from a ring supervisor;

(iv) receiving a ring port response from each of the devices having received the ring port request;

(v) incrementing a count associated with each [device based on receiving] *of the devices from which a ring port response [from the device] is received other than the discovered device;*

(vi) re-enabling the previously disabled Ethernet port of the [one of the devices] *discovered device* in the Ethernet ring; and

(vii) repeating steps (i) through (vi) until each device in the Ethernet ring has had an Ethernet port disabled; and

determining the topology of the devices in the Ethernet ring by ordering the devices in the Ethernet ring based on a value of the count associated with each device, wherein the device having a highest count value being closest to the forwarding port of the ring supervisor, and the device having a lowest count value being furthest from the forwarding port of the ring supervisor.

15. The method of claim 14, wherein discovering the number and [identity] *identification* of devices on the Ethernet ring comprises:

broadcasting a device discovery command from a ring topology generator to the devices in the Ethernet ring; and

receiving device discovery information from the devices in the Ethernet ring at the ring topology generator.

16. The method of claim 15, wherein the device discovery information includes a device identification (ID), the Device identification (ID) comprising one or more of a media access control (MAC) address, a building automation control network (BACnet) object identification (OID), and a logical device name.

17. The method of claim 15, further comprising:

broadcasting a ring port query message to the Ethernet ring from the ring topology generator based on receiving the device discovery information from the one or more devices in the Ethernet ring;

receiving a ring port query response from the one or more devices in the Ethernet ring; and

wherein the ring topology generator is configured to orient the Ethernet ports of the one or more devices in the Ethernet ring based on the received port query responses.

18. The method of claim 17, wherein the ring port response comprises a message identification (ID), a [Device] *device identification (ID)*, and an Ethernet port identification on which the ring port query message was received.

19. The method of claim 14, wherein the count for the ring port response received from the *discovered device* that currently has the Ethernet port disabled is not incremented while the Ethernet port is disabled.

20. The method of claim 14, wherein the Ethernet ring is configured as a Media Redundancy Protocol (MRP) ring.

21. *A method of determining a topology of devices in an Ethernet ring or a daisy chain, comprising:*

broadcasting a device discovery command from a ring topology generator to the devices in the Ethernet ring or the daisy chain;

receiving device discovery information from one or more of the devices in the Ethernet ring or the daisy chain at the ring topology generator;

26

iteratively performing a plurality of steps for each discovered device in the Ethernet ring or the daisy chain until the plurality of steps have been repeated for each discovered device, the plurality of steps comprising:

(i) disabling one or more ports on the discovered device;

(ii) broadcasting a ring port request to the devices in the Ethernet ring or daisy chain;

(iii) receiving a ring port response from each of the devices having received the ring port request; and

(iv) re-enabling the one or more ports on the discovered device; and

determining the topology of the devices of the Ethernet ring or the daisy chain based on the ring port responses.

22. *The method of claim 21, wherein the ring topology generator does not receive a ring port response from the discovered device that currently has the one or more ports disabled.*

23. *The method of claim 21, wherein the Ethernet ring is configured as a Media Redundancy Protocol (MRP) ring.*

24. *The method of claim 21, wherein the device discovery information includes a device identification (ID).*

25. *The method of claim 24, wherein the device ID is one or more of a media access control (MAC) address, a building automation control network (BACnet) object identification (OID), and a logical device name.*

26. *The method of claim 21, further comprising:*

broadcasting a ring port query message based on receiving the device discovery information from the one or more devices in the Ethernet ring or the daisy chain; and

receiving a ring port query response from the one or more devices in the Ethernet ring or the daisy chain;

wherein the ring topology generator is configured to orient the ports of the one or more devices in the Ethernet ring or the daisy chain based on the received ring port query responses.

27. *The method of claim 26, wherein the ring port query response comprises a message identification (ID), a device identification (ID), and an Ethernet port on which the ring port query message was received.*

28. *A system for determining a topology of an Ethernet ring or a daisy chain, the system comprising:*

a plurality of devices configured in the Ethernet ring or the daisy chain; and

a ring topology generator configured to:

broadcast a device discovery command from the ring topology generator to the devices in the Ethernet ring or the daisy chain;

receive device discovery information from one or more of the devices in the Ethernet ring or the daisy chain at the ring topology generator;

execute instructions for each discovered device in the Ethernet ring or daisy chain, the instructions causing the ring topology generator to:

(i) disable one or more ports on the discovered device;

(ii) broadcast a ring port request to the devices in the Ethernet ring or daisy chain; and

(iii) receive a ring port response from each of the devices having received the ring port request; and

(iv) re-enable the one or more ports on the discovered device; and

(v) repeat steps (i) through (iv) until each of the devices in the Ethernet ring or daisy chain has had one or more ports disabled and re-enabled,

27

wherein the plurality of devices remain connected to the ring topology generator during steps (i) through (v) when arranged in the Ethernet ring; and

determine the topology of the devices of the Ethernet ring or the daisy chain based on the ring port responses.

29. The system of claim 28, wherein the ring topology generator does not receive a ring port response from the discovered device that currently has the one or more ports disabled.

30. The system of claim 28, wherein the Ethernet ring is configured as one of a Media Redundancy Protocol (MRP) ring and a Spanning Tree Protocol (STP) ring.

31. The system of claim 28, wherein the device discovery information comprises a device identification (ID), the device ID comprising one or more of a media access control (MAC) address, a building automation control network (BACnet) object identification (OD), or a logical device name.

32. The system of claim 28, wherein the ring topology generator is further configured to:

broadcast a ring port query message based on receiving the device discovery information from the one or more devices in the Ethernet ring or the daisy chain;

receive a ring port query response from the one or more devices in the Ethernet ring or the daisy chain; and

orient the ports of the one or more devices in the Ethernet ring or the daisy chain based on the received ring port query responses.

33. The system of claim 32, wherein the ring port query response comprises a message identification (ID), a device identification (ID), and an Ethernet port on which the ring port query message was received.

34. A method of determining a topology of devices in an Ethernet ring or a daisy chain, comprising:

discovering a number and identification of devices within the Ethernet ring or the daisy chain;

initiating a procedure to determine the topology of the devices of the Ethernet ring or the daisy chain, the procedure comprising:

28

(i) disabling one or more ports on a reference device in the Ethernet ring or daisy chain;

(ii) broadcasting a ring port request to the devices in the Ethernet ring or daisy chain;

(iii) receiving a ring port response from each of the devices having received the ring port request;

(iv) re-enabling the one or more ports on the reference device; and

(v) repeating steps (i) through (iv) using a different device in the Ethernet ring or daisy chain as the reference device until each device in the Ethernet ring or daisy chain has been used as the reference device; and

determining the topology of the devices of the Ethernet ring or the daisy chain based on the ring port responses.

35. The method of claim 34, further comprising disabling a port on each of the devices and verifying that the Ethernet ring or the daisy chain is open.

36. The method of claim 34, wherein discovering the number and identification of devices on the Ethernet ring or the daisy chain comprises:

broadcasting a device discovery command from a ring topology generator to the devices in the Ethernet ring or the daisy chain; and

receiving device discovery information from the devices in the Ethernet ring or the daisy chain at the ring topology generator.

37. The method of claim 36, wherein the device discovery information includes a device identification (ID).

38. The method of claim 37, wherein the device identification (ID) comprises one or more of a media access control (MAC) address, a building automation control network (BACnet) object identification (OD), and a logical device name.

39. The method of claim 34, wherein the ring port response comprises a message identification (ID), a device identification (ID), and an Ethernet port or daisy chain identification on which the ring port query message was received.

40. The method of claim 34, wherein the Ethernet ring is configured as a Media Redundancy Protocol (MRP) ring.

* * * * *