(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
02 February 2023 (02.02.2023)

WIPO | PCT

(10) International Publication Number
WO 2023/009972 A1

(51) International Patent Classification:
E06B 9/24 (2006.01) *H02M 3/158* (2006.01)
G05B 15/02 (2006.01) *H04L 67/125* (2022.01)
G02F 1/163 (2006.01)

(21) International Application Number:
PCT/US2022/074066

(22) International Filing Date:
22 July 2022 (22.07.2022)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
63/203,543 27 July 2021 (27.07.2021) US

(71) Applicant: **VIEW, INC.** [US/US]; 195 South Milpitas Boulevard, Milpitas, California 95035 (US).

(72) Inventors: **MARTINSON, Robert Michael**; 195 South Milpitas Boulevard, Milpitas, California 95035 (US). **GOMEZ-MARTINEZ, Feliciano Vicente**; 195 South Milpitas Boulevard, Milpitas, California 95035 (US). **GOPINATHANASARI, Sajith Kamalnath**; 195 South Milpitas Boulevard, Milpitas, California 95035 (US). **TRIKHA, Nitesh**; 195 South Milpitas Boulevard, Milpitas, California 95035 (US). **BROWN, Stephen Clark**; 195 South Milpitas Boulevard, Milpitas, California 95035 (US). **PRADHAN, Anshu Ajit**; 195 South Milpitas Boulevard, Milpitas, California 95035 (US). **KAILASAM, Sridhar Karthik**; 195 South Milpitas Boulevard, Milpitas, California 95035 (US).

(74) Agent: **SRINIVASAN, Arthi G.** et al.; WEAVER AUSTIN VILLENEUVE & SAMPSON LLP, P. O. BOX 70250, OAKLAND, California 94612-0250 (US).

(54) Title: CONTROLLERS FOR OPTICALLY SWITCHABLE DEVICES

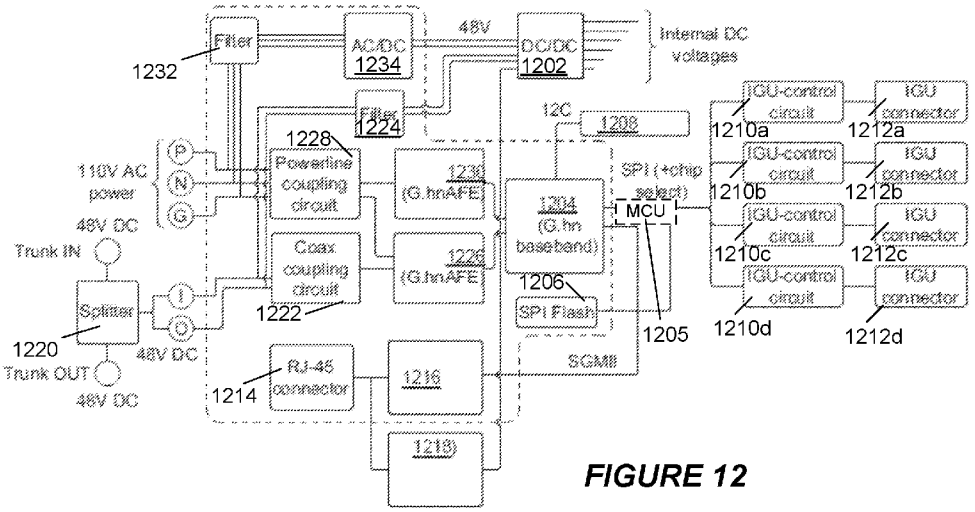


FIGURE 12

(57) Abstract: Window controller systems and methods are disclosed herein. In some embodiments, a window controller system for controlling multiple optically switchable devices comprises a printed circuit board comprising a first plurality of footprints to which a first plurality of components is mounted and a second plurality of footprints, wherein a subset of the second plurality of footprints is populated by a second plurality of components. The first plurality of components may comprise: a plurality of insulated glass unit (IGU) controllers, each configured to control an IGU of a corresponding plurality of IGUs operatively coupled to the window controller system; and a processing unit configured to control each of the plurality of IGU controllers. The second plurality of components may be selected based on a cable type and/or a protocol type used to provide power and data signals to the printed circuit board.

WO 2023/009972 A1

[Continued on next page]

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

(84) **Designated States** *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— *with international search report (Art. 21(3))*

— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

CONTROLLERS FOR OPTICALLY SWITCHABLE DEVICES

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to U.S. Provisional Application 63/203,543, filed on July 27, 2021 (Attorney Docket No. VIEWP132P), which is hereby incorporated by reference in its entirety for all purposes. This application claims priority as a Continuation-in-Part to U.S. Patent Application No. 17/654,458 filed on March 11, 2022 (Attorney Docket No. VIEWP083C2), which is a continuation of U.S. Patent Application No. 16/568,639 filed on September 12, 2019 (Attorney Docket No. VIEWP083C1), which is a continuation of U.S. Patent Application No. 15/334,835 filed on October 26, 2016 (Attorney Docket No. VIEWP083), which claims priority to U.S. Provisional Patent Application No. 62/248,181, filed October 29, 2015 (Attorney Docket No. VIEWP083P. U.S. Patent Application No. 15/334,835 is Continuation-In-Part of 14/998,019, filed October 6, 2015 (Attorney Docket No. VIEWP081). This application claims priority as a Continuation-in-Part to International Patent Application Serial No. PCT/US2021/017946 filed on February 12, 2021 (Attorney Docket No. VIEWP072X5WO), which claims priority to U.S. Patent Provisional Patent Application No. 62/977,001 filed on February 14, 2020 (Attorney Docket No. VIEWP072X5P), U.S. Provisional Patent Application No. 62/978,755 filed on February 19, 2020 (Attorney Docket No. VIEWP072X5P2), U.S. Provisional Patent Application No. 63/027,452 filed on May 20, 2020 (Attorney Docket No. VIEWP072X5P3), and U.S. Provisional Patent Application No. 63/146,365 filed on February 5, 2021 (Attorney Docket No. VIEWP072X5P4). This application claims priority as a Continuation-in-Part to U.S. Patent Application No. 16/949,855, filed on November 17, 2020 (Attorney Docket No. VIEWP041C3), which is a continuation of U.S. Patent Application No. 16/948,341, filed on September 14, 2020 (Attorney Docket No. VIEWP041C2US), which is a continuation of U.S. Patent Application No. 15/882,719, filed on January 29, 2018 (Attorney Docket No. VIEWP041C1US), which is a continuation of U.S. Patent Application No. 13/449,248, filed on April 17, 2012 (Attorney Docket No. VIEWP041US). This application claims priority as a Continuation-in-Part to U.S. Patent Application No. 17/247,662, filed on December 18, 2020 (Attorney Docket No. VIEWP008X1C4), which is a continuation of U.S. Patent Application No. 16/386,094, filed on April 16, 2019 (Attorney Docket No. VIEWP008X1C1US), which is a continuation of 16/253,971, filed on January 22, 2019 (Attorney Docket No. VIEWP008X1D1US), which is a division of U.S. Patent Application No. 14/951,410, filed on November 24, 2015 (Attorney Docket No. VIEWP008X1US), which claims the benefit of priority to U.S. Provisional Patent Application No. 62/085,179, filed on November 26, 2014 (Attorney Docket No.

VIEWP008X1PUS). Each of these applications is incorporated herein by reference in its entirety for all purposes.

TECHNICAL FIELD

[0002] This disclosure relates generally to optically-switchable devices, and more particularly, to controllers for optically-switchable devices.

BACKGROUND

[0003] The development and deployment of optically-switchable windows have increased as considerations of energy efficiency and system integration gain momentum. Electrochromic windows are a promising class of optically-switchable windows. Electrochromism is a phenomenon in which a material exhibits a reversible electrochemically-mediated change in one or more optical properties when stimulated to a different electronic state. Electrochromic materials and the devices made from them may be incorporated into, for example, windows for home, commercial, or other use. The color, tint, transmittance, absorbance, or reflectance of electrochromic windows can be changed by inducing a change in the electrochromic material, for example, by applying a voltage across the electrochromic material. Such capabilities can allow for control over the intensities of various wavelengths of light that may pass through the window. One area of relatively recent interest is in intelligent control systems and algorithms for driving optical transitions in optically-switchable windows to provide desirable lighting conditions while reducing the power consumption of such devices and improving the efficiency of systems with which they are integrated.

SUMMARY

[0004] Controllers for optically switchable windows are disclosed herein.

[0005] According to some implementations, a window controller system for controlling multiple optically switchable devices comprises a printed circuit board comprising a first plurality of footprints to which a first plurality of components is mounted and a second plurality of footprints, wherein a subset of the second plurality of footprints is populated by a second plurality of components. The first plurality of components comprises: a plurality of insulated glass unit (IGU) controllers, each configured to control an IGU of a corresponding plurality of IGUs operatively coupled to the window controller system; and a processing unit configured to control each of the plurality of IGU controllers. The second plurality of components is selected based on a cable type and/or a protocol type used to provide power and data signals to the printed circuit board.

[0006] In some examples, the protocol type comprises at least one of: a home networking protocol, or a power over ethernet (PoE) protocol. In some examples, the home networking protocol is a G.hn protocol.

[0007] In some examples, the cable type comprises at least one of: a coaxial cable, a powerline cable, or a twisted conductor pair cable.

[0008] In some examples, the protocol type comprises G.hn, and wherein the second plurality of components comprise one of: (i) a first subset of components configured to utilize a G.hn coaxial input, or (ii) a second subset of components configured to utilize a G.hn powerline input. In some examples, the first subset of components configured to utilize the G.hn coaxial input comprise a splitter configured to operatively couple to a trunk line, and a coaxial coupling circuit configured to provide power from the trunk line to an analog front end component. In some examples, the analog front end component is operatively coupled to the processing unit. In some examples, the second subset of components configured to utilize the G.hn powerline input comprise a powerline coupling circuit configured to operatively couple to an AC power outlet, and an analog front end component. In some examples, the analog front end component is operatively coupled to the processing unit. In some examples, the second set of components further comprise an A-to-D converter (ADC) configured to convert AC signals from the AC power outlet to DC signals.

[0009] In some examples, the protocol type comprises G.hn, and wherein the second plurality of components comprise a third subset of components to be utilized with either a G.hn coaxial input or a G.hn powerline input.

[0010] In some examples, the first plurality of components further comprises a memory device. In some examples, the memory device comprises a flash memory device. In some examples, the memory device is configured to store an indication of the cable type and/or the protocol type used to provide the power and the data signals.

[0011] In some examples, the processing unit is configured to characterize leakage current from one or more IGUs of the corresponding plurality of IGUs.

[0012] In some examples, the processing unit is further configured to communicate with a computing device. In some examples, the computing device is a cloud device remote from a premises on which the window controller system is installed. In some examples, the processing unit is further configured to: detect an anomalous event based on the characterized leakage current; and responsive to detecting the anomalous event, cause data associated with the characterized leakage current to be transmitted to the computing device. In some examples, the computing device is a cloud computing device. In some examples, the computing device is housed in a control panel installed in a facility in which the window controller system is installed.

[0013] In some examples, the window controller system further comprises a memory device configured to store window zone information associated with the plurality of IGUs and/or tint

schedule information associated with the plurality of IGUs. In some examples, the processing unit is further configured to: receive the window zone information and/or the tint schedule information from a computing device; and cause the window zone information and/or the tint schedule information to be stored in the memory.

[0014] In some examples, the processing unit is configured to: receive an override to a tint command for at least one IGU of the plurality of IGUs; and cause the at least one IGU to transition to a tint state based on the received override. In some examples, the override is received via an application executing on a mobile device. In some examples, the override is received via a local network to which the processing unit is communicatively coupled. In some examples, the override is based on information obtained by a sky sensor device configured to measure solar radiation external to a facility in which the window controller system is installed. In some examples, the override is received via a local network to which the sky sensor device and the processing unit are communicatively coupled.

[0015] In some examples, the processing unit is configured to cause each IGU controller to apply a drive voltage having a ramp rate to the corresponding IGU, wherein a first ramp rate associated with a first IGU of the plurality of IGUs is different than a second ramp rate associated with a second IGU of the plurality of IGUs. In some examples, the drive voltage applied to the first IGU is at a time point that is offset from application of the drive voltage to the second IGU.

[0016] In some examples, the processing unit is configured to communicate with at least one computing device configured to control building operations of a facility the window controller system is disposed in. In some examples, the at least one computing device is an edge server. In some examples, the processing unit is configured to enter a sleep mode responsive to a signal from the at least one computing device.

[0017] In some examples, the first plurality of components further comprises at least one component configured to provide one or more wireless communications channels. In some examples, the one or more wireless communications channels utilize at least one of a Wi-Fi communication protocol or a BLUETOOTH communication protocol. In some examples, the processing unit is configured to communicate using the Wi-Fi communication protocol with a cloud device. In some examples, the processing unit is configured to receive one or more updates from the cloud device using the Wi-Fi communication protocol. In some examples, the processing unit is configured to participate in a mesh network with at least one other processing unit associated with at least one other window controller system.

[0018] In some examples, the window controller system further comprises a second processing unit operatively coupled between the processing unit and each of the plurality of IGU

controllers, wherein the second processing unit is configured to control each of the plurality of IGU controllers.

[0019] According to some implementations, a window controller system for controlling multiple optically switchable windows comprises: an enclosure, wherein the enclosure comprises a 3-gang or a 4-gang outlet box and a faceplate. The window controller system may comprise a printed circuit board comprising a plurality of components, the plurality of components comprising: a plurality of insulated glass unit (IGU) controllers, each configured to control an IGU of a corresponding plurality of IGUs operatively coupled to the window controller system; and a processing unit configured to control each of the plurality of IGU controllers.

[0020] In some examples, the window controller system further comprises a heat sink disposed in the enclosure.

[0021] In some examples, the face plate is affixed to the 3-gang or 4-gang outlet box via a plurality of screws.

[0022] In some examples, at least one edge of the face plate is angled.

[0023] In some examples, the face plate comprises a plurality of air vents.

[0024] In some examples, the window controller system further comprises a cooling fan disposed within the enclosure.

[0025] In some examples, the window controller system further comprises a plurality of light-emitting diode (LED) indicators.

[0026] In some examples, the window controller system further comprises a second printed circuit board configured for providing power and/or data via a powerline system. In some examples, the second printed circuit board comprises an AC power connector. In some examples, the second printed circuit board is disposed above the printed circuit board and closer to the face plate than the printed circuit board. In some examples, the printed circuit board comprises a notch configured to pass through a powerline cable through the notch to the second printed circuit board.

[0027] According to some implementations, a method for controlling tinting of multiple optically switchable devices comprises (a) applying, via a window controller, a drive voltage to two or more optically switchable devices in connection with a tint transition from a first tint state to a second tint state, wherein applying the drive voltage to the two or more optically switchable device comprises providing a voltage from a power supply common to the two or more optically switchable devices. The method may involve (b) monitoring, for each of the two or more optically switchable devices, parameters indicative of a state of the tint transition. The method may involve (c) repeating (a) and (b) until it is determined that one or more optically switchable devices have completed the tint transition to the second tint state. The method may involve (d)

in response to determining at (c) that the one or more optically switchable devices have completed the tint transition to the second state, applying hold voltages to the one or more optically switchable devices while concurrently continuing to apply the drive voltage to a remaining one or more optically switchable devices of the two or more optically switchable devices that have not completed the tint transition. The method may involve (e) repeating (b) – (d) until the two or more optically switchable devices have completed the tint transition.

[0028] In some examples, monitoring the parameters indicative of the state of the tint transition comprises measuring, for each of the two or more optically switchable devices separately, an open circuit voltage (Voc) and/or a charge delivered. In some examples, determining that the one or more optically switchable devices have completed the tint transition to the second tint state comprises (i) determining that the Voc measured for the one or more optically switchable devices meets or exceeds a target Voc and/or (ii) determining that the charge delivered to the one or more optically switchable devices meets or exceeds a target charge to be delivered to the one or more optically switchable devices.

[0029] In some examples, the method may further involve: (f) measuring, for an optically switchable device for which the hold voltage is being applied, a hold state Voc; (g) determining whether the hold state Voc is less than a predetermined Voc threshold; and (h) in response to determining that the hold state Voc is less than the predetermined Voc threshold, applying a voltage pulse to the optically switchable device.

BRIEF DESCRIPTION OF THE DRAWINGS

[0030] Figure 1 shows a cross-sectional side view of an example electrochromic window 100 in accordance with some implementations.

[0031] Figure 2 illustrates an example control profile in accordance with some implementations.

[0032] Figure 3 shows a block diagram of an example network system operable to control a plurality of IGUs in accordance with some implementations.

[0033] Figure 4 shows a block diagram of an example master controller (MC) in accordance with some implementations.

[0034] Figure 5 shows a block diagram of an example network controller (NC) in accordance with some implementations.

[0035] Figure 6 shows a circuit schematic diagram of an example window controller (WC) in accordance with some implementations.

[0036] Figure 7 shows a diagram of an example connection architecture for coupling a window controller to an IGU in accordance with some implementations.

[0037] Figure 8 shows a block diagram of example modules of a network controller in accordance with some implementations.

[0038] Figure 9 shows a flowchart of an example process for controlling multiple IGUs in accordance with some implementations.

[0039] Figure 10 shows a flowchart of an example process for monitoring drive voltages of multiple IGUs in accordance with some implementations.

[0040] Figure 11 shows a schematic diagram of an example system for controlling multiple IGUs in accordance with some implementations.

[0041] Figure 12 shows a schematic diagram of an example circuit board layout for an example system for controlling multiple IGUs in accordance with some implementations.

[0042] Figures 13A-13C illustrate various views of an example system, including an enclosure, for controlling multiple IGUs in accordance with some implementations.

[0043] Figure 14 illustrates an example control panel system for interfacing with one or more controllers for controlling multiple IGUs in accordance with some implementations.

[0044] Figure 15 illustrates an example block diagram of a system for controlling multiple IGUs in accordance with some implementations.

[0045] Figure 16 illustrates an example block diagram of a system for communicating with multiple IGUs using various wireless communication protocols in accordance with some implementations.

[0046] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0047] The following detailed description is directed to specific example implementations for purposes of disclosing the subject matter. Although the disclosed implementations are described in sufficient detail to enable those of ordinary skill in the art to practice the disclosed subject matter, this disclosure is not limited to particular features of the specific example implementations described herein. On the contrary, the concepts and teachings disclosed herein can be implemented and applied in a multitude of different forms and ways without departing from their spirit and scope. For example, while the disclosed implementations focus on electrochromic windows (also referred to as smart windows), some of the systems, devices and methods disclosed herein can be made, applied or used without undue experimentation to incorporate, or while incorporating, other types of optically-switchable devices. Some other types of optically-switchable devices include liquid crystal devices, suspended particle devices, and even micro-blinds, among others. For example, some or all of such other optically-switchable devices can be powered, driven or otherwise controlled or integrated with one or

more of the disclosed implementations of controllers described herein. Additionally, in the following description, the phrases "operable to," "adapted to," "configured to," "designed to," "programmed to," or "capable of" may be used interchangeably where appropriate.

Example Electrochromic Window Architecture

[0048] Figure 1 shows a cross-sectional side view of an example electrochromic window 100 in accordance with some implementations. An electrochromic window is one type of optically-switchable window that includes an electrochromic device (ECD) used to provide tinting or coloring. The example electrochromic window 100 can be manufactured, configured or otherwise provided as an insulated glass unit (IGU) and will hereinafter also be referred to as IGU 100. This convention is generally used, for example, because it is common and because it can be desirable to have IGUs serve as the fundamental constructs for holding electrochromic panes (also referred to as "lites") when provided for installation in a building. An IGU lite or pane may be a single substrate or a multi-substrate construct, such as a laminate of two substrates. IGUs, especially those having double- or triple-pane configurations, can provide a number of advantages over single pane configurations; for example, multi-pane configurations can provide enhanced thermal insulation, noise insulation, environmental protection and/or durability when compared with single-pane configurations. A multi-pane configuration also can provide increased protection for an ECD, for example, because the electrochromic films, as well as associated layers and conductive interconnects, can be formed on an interior surface of the multi-pane IGU and be protected by an inert gas fill in the interior volume, 108, of the IGU.

[0049] Figure 1 more particularly shows an example implementation of an IGU 100 that includes a first pane 104 having a first surface S1 and a second surface S2. In some implementations, the first surface S1 of the first pane 104 faces an exterior environment, such as an outdoors or outside environment. The IGU 100 also includes a second pane 106 having a first surface S3 and a second surface S4. In some implementations, the second surface S4 of the second pane 106 faces an interior environment, such as an inside environment of a home, building or vehicle, or a room or compartment within a home, building or vehicle.

[0050] In some implementations, each of the first and the second panes 104 and 106 are transparent or translucent—at least to light in the visible spectrum. For example, each of the panes 104 and 106 can be formed of a glass material and especially an architectural glass or other shatter-resistant glass material such as, for example, a silicon oxide (SO_x) -based glass material. As a more specific example, each of the first and the second panes 104 and 106 can be a soda-lime glass substrate or float glass substrate. Such glass substrates can be composed of, for example, approximately 75% silica (SiO_2) as well as Na_2O, CaO, and several minor additives. However, each of the first and the second panes 104 and 106 can be formed of any

material having suitable optical, electrical, thermal, and mechanical properties. For example, other suitable substrates that can be used as one or both of the first and the second panes 104 and 106 can include other glass materials as well as plastic, semi-plastic and thermoplastic materials (for example, poly(methyl methacrylate), polystyrene, polycarbonate, allyl diglycol carbonate, SAN (styrene acrylonitrile copolymer), poly(4-methyl-1-pentene), polyester, polyamide), or mirror materials. In some implementations, each of the first and the second panes 104 and 106 can be strengthened, for example, by tempering, heating, or chemically strengthening.

[0051] Generally, each of the first and the second panes 104 and 106, as well as the IGU 100 as a whole, is a rectangular solid. However, in some other implementations other shapes are possible and may be desired (for example, circular, elliptical, triangular, curvilinear, convex or concave shapes). In some specific implementations, a length "L" of each of the first and the second panes 104 and 106 can be in the range of approximately 20 inches (in.) to approximately 10 feet (ft.), a width "W" of each of the first and the second panes 104 and 106 can be in the range of approximately 20 in. to approximately 10 ft., and a thickness "T" of each of the first and the second panes 104 and 106 can be in the range of approximately 0.3 millimeter (mm) to approximately 10 mm (although other lengths, widths or thicknesses, both smaller and larger, are possible and may be desirable based on the needs of a particular user, manager, administrator, builder, architect or owner). In examples where thickness T of substrate 104 is less than 3 mm, typically the substrate is laminated to an additional substrate which is thicker and thus protects the thin substrate 104. Additionally, while the IGU 100 includes two panes (104 and 106), in some other implementations, an IGU can include three or more panes. Furthermore, in some implementations, one or more of the panes can itself be a laminate structure of two, three, or more layers or sub-panes.

[0052] The first and second panes 104 and 106 are spaced apart from one another by a spacer 118, which is typically a frame structure, to form an interior volume 108. In some implementations, the interior volume is filled with Argon (Ar), although in some other implementations, the interior volume 108 can be filled with another gas, such as another noble gas (for example, krypton (Kr) or xenon (Xn)), another (non-noble) gas, or a mixture of gases (for example, air). Filling the interior volume 108 with a gas such as Ar, Kr, or Xn can reduce conductive heat transfer through the IGU 100 because of the low thermal conductivity of these gases as well as improve acoustic insulation due to their increased atomic weights. In some other implementations, the interior volume 108 can be evacuated of air or other gas. Spacer 118 generally determines the height "C" of the interior volume 108; that is, the spacing between the first and the second panes 104 and 106. In Figure 1, the thickness of the ECD, sealant 120/122 and bus bars 126/128 is not to scale; these components are generally very thin but are

exaggerated here for ease of illustration only. In some implementations, the spacing "C" between the first and the second panes 104 and 106 is in the range of approximately 6 mm to approximately 30 mm. The width "D" of spacer 118 can be in the range of approximately 5 mm to approximately 15 mm (although other widths are possible and may be desirable).

[0053] Although not shown in the cross-sectional view, spacer 118 is generally a frame structure formed around all sides of the IGU 100 (for example, top, bottom, left and right sides of the IGU 100). For example, spacer 118 can be formed of a foam or plastic material. However, in some other implementations, spacers can be formed of metal or other conductive material, for example, a metal tube or channel structure having at least 3 sides, two sides for sealing to each of the substrates and one side to support and separate the lites and as a surface on which to apply a sealant, 124. A first primary seal 120 adheres and hermetically seals spacer 118 and the second surface S2 of the first pane 104. A second primary seal 122 adheres and hermetically seals spacer 118 and the first surface S3 of the second pane 106. In some implementations, each of the primary seals 120 and 122 can be formed of an adhesive sealant such as, for example, polyisobutylene (PIB). In some implementations, IGU 100 further includes secondary seal 124 that hermetically seals a border around the entire IGU 100 outside of spacer 118. To this end, spacer 118 can be inset from the edges of the first and the second panes 104 and 106 by a distance "E." The distance "E" can be in the range of approximately 4 mm to approximately 8 mm (although other distances are possible and may be desirable). In some implementations, secondary seal 124 can be formed of an adhesive sealant such as, for example, a polymeric material that resists water and that adds structural support to the assembly, such as silicone, polyurethane and similar structural sealants that form a water tight seal.

[0054] In the particular configuration and form factor depicted in Figure 1, the ECD coating on surface S2 of substrate 104 extends about its entire perimeter to and under spacer 118. This configuration is functionally desirable as it protects the edge of the ECD within the primary sealant 120 and aesthetically desirable because within the inner perimeter of spacer 118 there is a monolithic ECD without any bus bars or scribe lines. Such configurations are described in more detail in U.S. Patent No. 8,164,818, issued April 24, 2012 and titled ELECTROCHROMIC WINDOW FABRICATION METHODS (Attorney Docket No. VIEWP006), U.S. Patent Application No. 13/456,056 filed April 25, 2012 and titled ELECTROCHROMIC WINDOW FABRICATION METHODS (Attorney Docket No. VIEWP006X1), PCT Patent Application No. PCT/US2012/068817 filed December 10, 2012 and titled THIN-FILM DEVICES AND FABRICATION (Attorney Docket No. VIEWP036WO), U.S. Patent Application No. 14/362,863 filed June 4, 2014 and titled THIN-FILM DEVICES AND FABRICATION (Attorney Docket No. VIEWP036US), and PCT Patent Application No. PCT/US2014/073081,

filed December 13, 2014 and titled THIN-FILM DEVICES AND FABRICATION (Attorney Docket No. VIEWP036X1WO), all of which are hereby incorporated by reference in their entireties and for all purposes.

[0055] In the implementation shown in Figure 1, an ECD 110 is formed on the second surface S2 of the first pane 104. In some other implementations, ECD 110 can be formed on another suitable surface, for example, the first surface S1 of the first pane 104, the first surface S3 of the second pane 106 or the second surface S4 of the second pane 106. The ECD 110 includes an electrochromic ("EC") stack 112, which itself may include one or more layers. For example, the EC stack 112 can include an electrochromic layer, an ion-conducting layer, and a counter electrode layer. In some implementations, the electrochromic layer is formed of one or more inorganic solid materials. The electrochromic layer can include or be formed of one or more of a number of electrochromic materials, including electrochemically-cathodic or electrochemically-anodic materials. For example, metal oxides suitable for use as the electrochromic layer can include tungsten oxide (WO_3) and doped formulations thereof. In some implementations, the electrochromic layer can have a thickness in the range of approximately 0.05 μm to approximately 1 μm.

[0056] In some implementations, the counter electrode layer is formed of an inorganic solid material. The counter electrode layer can generally include one or more of a number of materials or material layers that can serve as a reservoir of ions when the EC device 110 is in, for example, the transparent state. In certain implementations, the counter electrode not only serves as an ion storage layer but also colors anodically. For example, suitable materials for the counter electrode layer include nickel oxide (NiO) and nickel tungsten oxide (NiWO), as well as doped forms thereof, such as nickel tungsten tantalum oxide, nickel tungsten tin oxide, nickel vanadium oxide, nickel chromium oxide, nickel aluminum oxide, nickel manganese oxide, nickel magnesium oxide, nickel tantalum oxide, nickel tin oxide as non-limiting examples. In some implementations, the counter electrode layer can have a thickness in the range of approximately 0.05 μm to approximately 1 μm.

[0057] The ion-conducting layer serves as a medium through which ions are transported (for example, in the manner of an electrolyte) when the EC stack 112 transitions between optical states. In some implementations, the ion-conducting layer is highly conductive to the relevant ions for the electrochromic and the counter electrode layers, but also has sufficiently low electron conductivity such that negligible electron transfer (electrical shorting) occurs during normal operation. A thin ion-conducting layer with high ionic conductivity enables fast ion conduction and consequently fast switching for high performance EC devices 110. In some implementations, the ion-conducting layer can have a thickness in the range of approximately 1

nm to approximately 500 nm, more generally in the range of about 5 nm to about 100 nm thick. In some implementations, the ion-conducting layer also is an inorganic solid. For example, the ion-conducting layer can be formed from one or more silicates, silicon oxides (including silicon-aluminum-oxide), tungsten oxides (including lithium tungstate), tantalum oxides, niobium oxides, lithium oxide and borates. These materials also can be doped with different dopants, including lithium; for example, lithium-doped silicon oxides include lithium silicon-aluminum-oxide, lithium phosphorous oxynitride (LiPON) and the like.

[0058] In some other implementations, the electrochromic layer and the counter electrode layer are formed immediately adjacent one another, sometimes in direct contact, without an ion-conducting layer in between and then an ion conductor material formed *in situ* between the electrochromic and counter electrode layers. A further description of suitable devices is found in U.S. Patent No. 8,764,950, titled ELECTROCHROMIC DEVICES, by Wang et al., issued July 1, 2014 and U.S. Patent No. 9,261,751, titled ELECTROCHROMIC DEVICES, by Pradhan et al., issued February 16, 2016, each of which is hereby incorporated by reference in its entirety and for all purposes. In some implementations, the EC stack 112 also can include one or more additional layers such as one or more passive layers. For example, passive layers can be used to improve certain optical properties, to provide moisture or to provide scratch resistance. These or other passive layers also can serve to hermetically seal the EC stack 112. Additionally, various layers, including conducting layers (such as the first and the second TCO layers 114 and 116 described below), can be treated with anti-reflective or protective oxide or nitride layers.

[0059] The selection or design of the electrochromic and counter electrode materials generally governs the possible optical transitions. During operation, in response to a voltage generated across the thickness of the EC stack 112 (for example, between the first and the second TCO layers 114 and 116), the electrochromic layer transfers or exchanges ions to or from the counter electrode layer to drive the electrochromic layer to the desired optical state. In some implementations, to cause the EC stack 112 to transition to a transparent state, a positive voltage is applied across the EC stack 112 (for example, such that the electrochromic layer is more positive than the counter electrode layer). In some such implementations, in response to the application of the positive voltage, the available ions in the stack reside primarily in the counter electrode layer. When the magnitude of the potential across the EC stack 112 is reduced or when the polarity of the potential is reversed, ions are transported back across the ion conducting layer to the electrochromic layer causing the electrochromic material to transition to an opaque state (or to a "more tinted," "darker" or "less transparent" state). Conversely, in some other implementations using electrochromic layers having different properties, to cause the EC stack 112 to transition to an opaque state, a negative voltage can be applied to the electrochromic layer

relative to the counter electrode layer. In such implementations, when the magnitude of the potential across the EC stack 112 is reduced or its polarity reversed, the ions are transported back across the ion conducting layer to the electrochromic layer causing the electrochromic material to transition to a clear or "bleached" state (or to a "less tinted", "lighter" or "more transparent" state).

[0060] In some implementations, the transfer or exchange of ions to or from the counter electrode layer also results in an optical transition in the counter electrode layer. For example, in some implementations the electrochromic and counter electrode layers are complementary coloring layers. More specifically, in some such implementations, when or after ions are transferred into the counter electrode layer, the counter electrode layer becomes more transparent, and similarly, when or after the ions are transferred out of the electrochromic layer, the electrochromic layer becomes more transparent. Conversely, when the polarity is switched, or the potential is reduced, and the ions are transferred from the counter electrode layer into the electrochromic layer, both the counter electrode layer and the electrochromic layer become less transparent.

[0061] In one more specific example, responsive to the application of an appropriate electric potential across a thickness of EC stack 112, the counter electrode layer transfers all or a portion of the ions it holds to the electrochromic layer causing the optical transition in the electrochromic layer. In some such implementations, for example, when the counter electrode layer is formed from NiWO, the counter electrode layer also optically transitions with the loss of ions it has transferred to the electrochromic layer. When charge is removed from a counter electrode layer made of NiWO (that is, ions are transported from the counter electrode layer to the electrochromic layer), the counter electrode layer will transition in the opposite direction.

[0062] Generally, the transition of the electrochromic layer from one optical state to another optical state can be caused by reversible ion insertion into the electrochromic material (for example, by way of intercalation) and a corresponding injection of charge-balancing electrons. In some instances, some fraction of the ions responsible for the optical transition is irreversibly bound up in the electrochromic material. Some or all of the irreversibly bound ions can be used to compensate for "blind charge" in the material. In some implementations, suitable ions include lithium ions (Li^+) and hydrogen ions (H^+) (i.e., protons). In some other implementations, other ions can be suitable. Intercalation of lithium ions, for example, into tungsten oxide (WO_{3-y} ($0 < y \leq \sim 0.3$)) causes the tungsten oxide to change from a transparent state to a blue state.

[0063] The description below generally focuses on tinting transitions. One example of a tinting transition is a transition from a transparent (or "translucent," "bleached" or "least tinted") state to an opaque (or "fully darkened" or "fully tinted") state. Another example of a tinting

transition is the reverse—a transition from an opaque state to a transparent state. Other examples of tinting transitions include transitions to and from various intermediate tint states, for example, a transition from a less tinted, lighter or more transparent state to a more tinted, darker or less transparent state, and vice versa. Each of such tint states, and the tinting transitions between them, may be characterized or described in terms of percent transmission. For example, a tinting transition can be described as being from a current percent transmission (% T) to a target % T. Conversely, in some other instances, each of the tint states and the tinting transitions between them may be characterized or described in terms of percent tinting; for example, a transition from a current percent tinting to a target percent tinting.

[0064] However, although the following description generally focuses on tint states and tinting transitions between tint states, other optical states and optical transitions also are achievable in various implementations. As such, where appropriate and unless otherwise indicated, references to tint states or tinting transitions also are intended to encompass other optical states and optical transitions. In other words, optical states and optical state transitions also will be referred to herein as tint states and tint state transitions, respectively, but this is not intended to limit the optical states and state transitions achievable by the IGUs 302. For example, such other optical states and state transitions can include states and state transitions associated with various colors, intensities of color (for example, from lighter blue to darker blue and vice versa), reflectivity (for example, from less reflective to more reflective and vice versa), polarization (for example, from less polarization to more polarization and vice versa), and scattering density (for example, from less scattering to more scattering and vice versa), among others. Similarly, references to devices, control algorithms or processes for controlling tint states, including causing tinting transitions and maintaining tint states, also are intended to encompass such other optical transitions and optical states. Additionally, controlling the voltage, current or other electrical characteristics provided to an optically-switchable device, and the functions or operations associated with such controlling, also may be described hereinafter as "driving" the device or the respective IGU, whether or not the driving involves a tint state transition or the maintaining of a current tint state.

[0065] The ECD 110 generally includes first and second conducting (or "conductive") layers. For example, the ECD 110 can include a first transparent conductive oxide (TCO) layer 114 adjacent a first surface of the EC stack 112 and a second TCO layer 116 adjacent a second surface of the EC stack 112. In some implementations, the first TCO layer 114 can be formed on the second surface S2, the EC stack 112 can be formed on the first TCO layer 114, and the second TCO layer 116 can be formed on the EC stack 112. In some implementations, the first and the second TCO layers 114 and 116 can each be formed of one or more metal oxides

including metal oxides doped with one or more metals. For example, some suitable metal oxides and doped metal oxides can include indium oxide, indium tin oxide (ITO), doped indium oxide, tin oxide, doped tin oxide, fluorinated tin oxide, zinc oxide, aluminum zinc oxide, doped zinc oxide, ruthenium oxide and doped ruthenium oxide, among others. While such materials are referred to as TCOs in this document, the term encompasses non-oxides as well as oxides that are transparent and electrically conductive such as certain thin film metals and certain non-metallic materials such as conductive metal nitrides and composite conductors, among other suitable materials. In some implementations, the first and the second TCO layers 114 and 116 are substantially transparent at least in the range of wavelengths where electrochromism is exhibited by the EC stack 112. In some implementations, the first and the second TCO layers 114 and 116 can each be deposited by physical vapor deposition (PVD) processes including, for example, sputtering. In some implementations, the first and the second TCO layers 114 and 116 can each have a thickness in the range of approximately 0.01 microns (μm) to approximately 1 μm. A transparent conductive material typically has an electronic conductivity significantly greater than that of the electrochromic material or the counter electrode material.

[0066] The first and the second TCO layers 114 and 116 serve to distribute electrical charge across respective first and second surfaces of the EC stack 112 to apply an electrical potential (voltage) across the thickness of the EC stack 112. For example, a first applied voltage can be applied to a first one of the TCO layers and a second applied voltage can be applied to a second one of the TCO layers. In some implementations, a first busbar 126 distributes the first applied voltage to the first TCO layer 114 and a second busbar 128 distributes the second applied voltage to the second TCO layer 116. In some other implementations, one of the first and the second busbars 126 and 128 can ground the respective one of the first and the second TCO layers 114 and 116. In other implementations the load can be floated with respect to the two TCOs. In various implementations, to modify one or more optical properties of the EC stack 112, and thus cause an optical transition, a controller can alter one or both of the first and second applied voltages to bring about a change in one or both of the magnitude and the polarity of the effective voltage applied across the EC stack 112. Desirably, the first and the second TCO layers 114 and 116 serve to uniformly distribute electrical charge over respective surfaces of the EC stack 112 with relatively little Ohmic potential drop from the outer regions of the respective surfaces to the inner regions of the surfaces. As such, it is generally desirable to minimize the sheet resistance of the first and the second TCO layers 114 and 116. In other words, it is generally desirable that each of the first and the second TCO layers 114 and 116 behaves as a substantially equipotential layer across all portions of the respective layer. In this way, the first

and the second TCO layers 114 and 116 can uniformly apply an electric potential across a thickness of the EC stack 112 to effect a uniform optical transition of the EC stack 112.

[0067] In some implementations, each of the first and the second busbars 126 and 128 is printed, patterned, or otherwise formed such that it is oriented along a length of the first pane 104 along at least one border of the EC stack 112. For example, each of the first and the second busbars 126 and 128 can be formed by depositing a conductive ink, such as a silver ink, in the form of a line. In some implementations, each of the first and the second busbars 126 and 128 extends along the entire length (or nearly the entire length) of the first pane 104, and in some implementations, along more than one edge of the EC stack 112.

[0068] In some implementations, the first TCO layer 114, the EC stack 112 and the second TCO layer 116 do not extend to the edges of the first pane 104. For example, a laser edge delete (LED) or other operation can be used to remove portions of the first TCO layer 114, the EC stack 112 and the second TCO layer 116 such that these layers are separated or inset from the respective edges of the first pane 104 by a distance "G," which can be in the range of approximately 8 mm to approximately 10 mm (although other distances are possible and may be desirable). Additionally, in some implementations, an edge portion of the EC stack 112 and the second TCO layer 116 along one side of the first pane 104 is removed to enable the first busbar 126 to be formed on the first TCO layer 114 to enable conductive coupling between the first busbar 126 and the first TCO layer 114. The second busbar 128 is formed on the second TCO layer 116 to enable conductive coupling between the second busbar 128 and the second TCO layer 116. In some implementations, the first and the second busbars 126 and 128 are formed in a region between spacer 118 and the first pane 104 as shown in Figure 1. For example, each of the first and the second busbars 126 and 128 can be inset from an inner edge of spacer 118 by at least a distance "F," which can be in the range of approximately 2 mm to approximately 3 mm (although other distances are possible and may be desirable). This arrangement can be advantageous for a number of reasons including, for example, to hide the busbars from view.

[0069] As noted above, the usage of the IGU convention is for convenience only. Indeed, in some implementations the basic unit of an electrochromic window can be defined as a pane or substrate of transparent material, upon which an ECD is formed or otherwise arranged, and to which associated electrical connections are coupled (to drive the ECD). As such, references to an IGU in the following description do not necessarily include all of the components described with reference to the IGU 100 of Figure 1.

Example Control Profile for Driving Optical Transitions

[0070] Figure 2 illustrates an example control profile 200 in accordance with some implementations. The control profile 200 can be used to drive a transition in an optically-switchable device, such as the ECD 110 described above. In some implementations, a window controller can be used to generate and apply the control profile 200 to drive an ECD from a first optical state (for example, a transparent state or a first intermediate state) to a second optical state (for example, a fully tinted state or a more tinted intermediate state). To drive the ECD in the reverse direction—from a more tinted state to a less tinted state—the window controller can apply a similar but inverted profile. For example, the control profile for driving the ECD from the second optical state to the first optical state can be a mirror image of the voltage control profile depicted in Figure 2. In some other implementations, the control profiles for tinting and lightening can be asymmetric. For example, transitioning from a first more tinted state to a second less tinted state can in some instances require more time than the reverse; that is, transitioning from the second less tinted state to the first more tinted state. In some other instances, the reverse may be true; that is, transitioning from the second less tinted state to the first more tinted state can require more time. In other words, by virtue of the device architecture and materials, bleaching or lightening is not necessarily simply the reverse of coloring or tinting. Indeed, ECDs often behave differently for each transition due to differences in driving forces for ion intercalation and deintercalation to and from the electrochromic materials.

[0071] In some implementations, the control profile 200 is a voltage control profile implemented by varying a voltage provided to the ECD. For example, the solid line in Figure 2 represents an effective voltage V_{Eff} applied across the ECD over the course of a tinting transition and a subsequent maintenance period. In other words, the solid line can represent the relative difference in the electrical voltages V_{App1} and V_{App2} applied to the two conducting layers of the ECD (for example, the first and the second TCO layers 114 and 116 of the ECD 110). The dotted line in Figure 2 represents a corresponding current (I) through the device. In the illustrated example, the voltage control profile 200 includes four stages: a ramp-to-drive stage 202 that initiates the transition, a drive stage that continues to drive the transition, a ramp-to-hold stage, and subsequent hold stage.

[0072] The ramp-to-drive stage 202 is characterized by the application of a voltage ramp that increases in magnitude from an initial value at time t_0 to a maximum driving value of V_{Drive} at time t_1. In some implementations, the ramp-to-drive stage 202 can be defined by three drive parameters known or set by the window controller: the initial voltage at t_0 (the current voltage across the ECD at the start of the transition), the magnitude of V_{Drive} (governing the ending optical state), and the time duration during which the ramp is applied (dictating the speed of the

transition). Additionally or alternatively, the window controller also can set a target ramp rate, a maximum ramp rate or a type of ramp (for example, a linear ramp, a second degree ramp or an n^{th}-degree ramp). In some applications, the ramp rate can be limited to avoid damaging the ECD.

[0073] The drive stage 204 is characterized by the application of a constant voltage V_{Drive} starting at time t_1 and ending at time t_2, at which point the ending optical state is reached (or approximately reached). The ramp-to-hold stage 206 is characterized by the application of a voltage ramp that decreases in magnitude from the drive value V_{Drive} at time t_2 to a minimum holding value of V_{Hold} at time t_3. In some implementations, the ramp-to-hold stage 206 can be defined by three drive parameters known or set by the window controller: the drive voltage V_{Drive}, the holding voltage V_{Hold}, and the time duration during which the ramp is applied. Additionally or alternatively, the window controller also can set a ramp rate or a type of ramp (for example, a linear ramp, a second degree ramp or an n^{th}-degree ramp).

[0074] The hold stage 208 is characterized by the application of a constant voltage V_{Hold} starting at time t_3. The holding voltage V_{Hold} is used to maintain the ECD at the ending optical state. As such, the duration of the application of the holding voltage V_{hold} may be concomitant with the duration of time that the ECD is to be held in the ending optical state. For example, because of non-idealities associated with the ECD, a leakage current I_{Leak} can result in the slow drainage of electrical charge from the ECD. Such a drainage of electrical charge can result in a corresponding reversal of ions across the ECD, and consequently, a slow reversal of the optical transition. In such applications, the holding voltage V_{Hold} can be continuously applied to counter or prevent the leakage current. In some other implementations, the holding voltage V_{Hold} can be applied periodically to "refresh" the desired optical state, or in other words, to bring the ECD back to the desired optical state.

[0075] The voltage control profile 200 illustrated and described with reference to Figure 2 is only one example of a voltage control profile suitable for some implementations. However, many other profiles may be desirable or suitable in such implementations or in various other implementations or applications. These other profiles also can readily be achieved using the controllers and optically-switchable devices disclosed herein. For example, in some implementations, a current profile can be applied instead of a voltage profile. In some such instances, a current control profile similar to that of the current density shown in Figure 2 can be applied. In some other implementations, a control profile can have more than four stages. For example, a voltage control profile can include one or more overdrive stages. In one example implementation, the voltage ramp applied during the first stage 202 can increase in magnitude beyond the drive voltage V_{Drive} to an overdrive voltage V_{OD}. In some such implementations, the

first stage 202 can be followed by a ramp stage 203 during which the applied voltage decreases from the overdrive voltage V_{OD} to the drive voltage V_{Drive}. In some other such implementations, the overdrive voltage V_{OD} can be applied for a relatively short time duration before the ramp back down to the drive voltage V_{Drive}.

[0076] Additionally, in some implementations, the applied voltage or current profiles can be interrupted for relatively short durations of time to provide open circuit conditions across the device. While such open circuit conditions are in effect, an actual voltage or other electrical characteristics can be measured, detected or otherwise determined to monitor how far along an optical transition has progressed, and in some instances, to determine whether changes in the profile are desirable. Such open circuit conditions also can be provided during a hold stage to determine whether a holding voltage V_{Hold} should be applied or whether a magnitude of the holding voltage V_{Hold} should be changed. Additional information related to driving and monitoring an optical transition is provided in PCT Patent Application No. PCT/US14/43514 filed June 20, 2014 and titled CONTROLLING TRANSITIONS IN OPTICALLY SWITCHABLE DEVICES, which is hereby incorporated by reference in its entirety and for all purposes.

Example Controller Network Architecture

[0077] In many instances, optically-switchable windows can form or occupy substantial portions of a building envelope. For example, the optically-switchable windows can form substantial portions of the walls, facades and even roofs of a corporate office building, other commercial building or a residential building. In various implementations, a distributed network of controllers can be used to control the optically-switchable windows. Figure 3 shows a block diagram of an example network system, 300, operable to control a plurality of IGUs 302 in accordance with some implementations. For example, each of the IGUs 302 can be the same or similar to the IGU 100 described above with reference to Figure 1. One primary function of the network system 300 is controlling the optical states of the ECDs (or other optically-switchable devices) within the IGUs 302. In some implementations, one or more of the windows 302 can be multi-zoned windows, for example, where each window includes two or more independently controllable ECDs or zones. In various implementations, the network system 300 is operable to control the electrical characteristics of the power signals provided to the IGUs 302. For example, the network system 300 can generate and communicate tinting instructions (also referred to herein as "tint commands") to control voltages applied to the ECDs within the IGUs 302.

[0078] In some implementations, another function of the network system 300 is to acquire status information from the IGUs 302 (hereinafter "information" is used interchangeably with

"data"). For example, the status information for a given IGU can include an identification of, or information about, a current tint state of the ECD(s) within the IGU. The network system 300 also can be operable to acquire data from various sensors, such as temperature sensors, photosensors (also referred to herein as light sensors), humidity sensors, air flow sensors, or occupancy sensors, whether integrated on or within the IGUs 302 or located at various other positions in, on or around the building.

[0079] The network system 300 can include any suitable number of distributed controllers having various capabilities or functions. In some implementations, the functions and arrangements of the various controllers are defined hierarchically. For example, the network system 300 includes a plurality of distributed window controllers (WCs) 304, a plurality of network controllers (NCs) 306, and a master controller (MC) 308. In some implementations, the MC 308 can communicate with and control tens or hundreds of NCs 306. In various implementations, the MC 308 issues high level instructions to the NCs 306 over one or more wired or wireless links 316 (hereinafter collectively referred to as "link 316"). The instructions can include, for example, tint commands for causing transitions in the optical states of the IGUs 302 controlled by the respective NCs 306. Each NC 306 can, in turn, communicate with and control a number of WCs 304 over one or more wired or wireless links 314 (hereinafter collectively referred to as "link 314"). For example, each NC 306 can control tens or hundreds of the WCs 304. Each WC 304 can, in turn, communicate with, drive or otherwise control one or more respective IGUs 302 over one or more wired or wireless links 312 (hereinafter collectively referred to as "link 312").

[0080] The MC 308 can issue communications including tint commands, status request commands, data (for example, sensor data) request commands or other instructions. In some implementations, the MC 308 can issue such communications periodically, at certain predefined times of day (which may change based on the day of week or year), or based on the detection of particular events, conditions or combinations of events or conditions (for example, as determined by acquired sensor data or based on the receipt of a request initiated by a user or by an application or a combination of such sensor data and such a request). In some implementations, when the MC 308 determines to cause a tint state change in a set of one or more IGUs 302, the MC 308 generates or selects a tint value corresponding to the desired tint state. In some implementations, the set of IGUs 302 is associated with a first protocol identifier (ID) (for example, a BACnet ID). The MC 308 then generates and transmits a communication—referred to herein as a "primary tint command"— including the tint value and the first protocol ID over the link 316 via a first communication protocol (for example, a BACnet compatible protocol). In some implementations, the MC 308 addresses the primary tint command to the particular NC

306 that controls the particular one or more WCs 304 that, in turn, control the set of IGUs 302 to be transitioned.

[0081] The NC 306 receives the primary tint command including the tint value and the first protocol ID and maps the first protocol ID to one or more second protocol IDs. In some implementations, each of the second protocol IDs identifies a corresponding one of the WCs 304. The NC 306 subsequently transmits a secondary tint command including the tint value to each of the identified WCs 304 over the link 314 via a second communication protocol. In some implementations, each of the WCs 304 that receives the secondary tint command then selects a voltage or current profile from an internal memory based on the tint value to drive its respectively connected IGUs 302 to a tint state consistent with the tint value. Each of the WCs 304 then generates and provides voltage or current signals over the link 312 to its respectively connected IGUs 302 to apply the voltage or current profile.

[0082] In some implementations, the various IGUs 302 can be advantageously grouped into zones of EC windows, each of which zones includes a subset of the IGUs 302. In some implementations, each zone of IGUs 302 is controlled by one or more respective NCs 306 and one or more respective WCs 304 controlled by these NCs 306. In some more specific implementations, each zone can be controlled by a single NC 306 and two or more WCs 304 controlled by the single NC 306. Said another way, a zone can represent a logical grouping of the IGUs 302. For example, each zone may correspond to a set of IGUs 302 in a specific location or area of the building that are driven together based on their location. As a more specific example, consider a building having four faces or sides: a North face, a South face, an East Face and a West Face. Consider also that the building has ten floors. In such a didactic example, each zone can correspond to the set of electrochromic windows 100 on a particular floor and on a particular one of the four faces. Additionally or alternatively, each zone may correspond to a set of IGUs 302 that share one or more physical characteristics (for example, device parameters such as size or age). In some other implementations, a zone of IGUs 302 can be grouped based on one or more non-physical characteristics such as, for example, a security designation or a business hierarchy (for example, IGUs 302 bounding managers' offices can be grouped in one or more zones while IGUs 302 bounding non-managers' offices can be grouped in one or more different zones).

[0083] In some such implementations, each NC 306 can address all of the IGUs 302 in each of one or more respective zones. For example, the MC 308 can issue a primary tint command to the NC 306 that controls a target zone. The primary tint command can include an abstract identification of the target zone (hereinafter also referred to as a "zone ID"). In some such implementations, the zone ID can be a first protocol ID such as that just described in the

example above. In such cases, the NC 306 receives the primary tint command including the tint value and the zone ID and maps the zone ID to the second protocol IDs associated with the WCs 304 within the zone. In some other implementations, the zone ID can be a higher level abstraction than the first protocol IDs. In such cases, the NC 306 can first map the zone ID to one or more first protocol IDs, and subsequently map the first protocol IDs to the second protocol IDs.

User or Third Party Interaction with Network

[0084] In some implementations, the MC 308 is coupled to one or more outward-facing networks, 310, (hereinafter collectively referred to as "the outward-facing network 310") via one or more wired or wireless links 318 (hereinafter "link 318"). In some such implementations, the MC 308 can communicate acquired status information or sensor data to remote computers, mobile devices, servers, databases in or accessible by the outward-facing network 310. In some implementations, various applications, including third party applications or cloud-based applications, executing within such remote devices can access data from or provide data to the MC 308. In some implementations, authorized users or applications can communicate requests to modify the tint states of various IGUs 302 to the MC 308 via the network 310. In some implementations, the MC 308 can first determine whether to grant the request (for example, based on power considerations or based on whether the user has the appropriate authorization) prior to issuing a tint command. The MC 308 can then calculate, determine, select or otherwise generate a tint value and transmit the tint value in a primary tint command to cause the tint state transitions in the associated IGUs 302.

[0085] For example, a user can submit such a request from a computing device, such as a desktop computer, laptop computer, tablet computer or mobile device (for example, a smartphone). In some such implementations, the user's computing device can execute a client-side application that is capable of communicating with the MC 308, and in some instances, with a master controller application executing within the MC 308. In some other implementations, the client-side application can communicate with a separate application, in the same or a different physical device or system as the MC 308, which then communicates with the master controller application to effect the desired tint state modifications. In some implementations, the master controller application or other separate application can be used to authenticate the user to authorize requests submitted by the user. In some implementations, the user can select the IGUs 302 to be tinted, and inform the MC 308 of the selections, by entering a room number via the client-side application.

[0086] Additionally or alternatively, in some implementations, a user's mobile device or other computing device can communicate wirelessly with various WCs 304. For example, a

client-side application executing within a user's mobile device can transmit wireless communications including tint state control signals to a WC 304 to control the tint states of the respective IGUs 302 connected to the WC 304. For example, the user can use the client-side application to maintain or modify the tint states of the IGUs 302 adjoining a room occupied by the user (or to be occupied by the user or others at a future time). Such wireless communications can be generated, formatted or transmitted using various wireless network topologies and protocols (described in more detail below with reference to the WC 600 of Figure 6).

[0087] In some such implementations, the control signals sent to the respective WC 304 from the user's mobile device (or other computing device) can override a tint value previously received by the WC 304 from the respective NC 306. In other words, the WC 304 can provide the applied voltages to the IGUs 302 based on the control signals from the user's computing device rather than based on the tint value. For example, a control algorithm or rule set stored in and executed by the WC 304 can dictate that one or more control signals from an authorized user's computing device take precedence over a tint value received from the NC 306. In some other instances, such as in high demand cases, control signals such as a tint value from the NC 306 may take precedence over any control signals received by the WC 304 from a user's computing device. In some other instances, a control algorithm or rule set may dictate that tint overrides from only certain users or groups or classes of users may take precedence based on permissions granted to such users, as well as in some instances, other factors including time of day or the location of the IGUs 302.

[0088] In some implementations, based on the receipt of a control signal from an authorized user's computing device, the MC 308 can use information about a combination of known parameters to calculate, determine, select or otherwise generate a tint value that provides lighting conditions desirable for a typical user, while in some instances also using power efficiently. In some other implementations, the MC 308 can determine the tint value based on preset preferences defined by or for the particular user that requested the tint state change via the computing device. For example, the user may be required to enter a password or otherwise login or obtain authorization to request a tint state change. In such instances, the MC 308 can determine the identity of the user based on a password, a security token or based on an identifier of the particular mobile device or other computing device. After determining the user's identity, the MC 308 can then retrieve preset preferences for the user, and use the preset preferences alone or in combination with other parameters (such as power considerations or information from various sensors) to generate and transmit a tint value for use in tinting the respective IGUs 302.

Wall Devices

[0089] In some implementations, the network system 300 also can include wall switches, dimmers or other tint-state-controlling devices. A wall switch generally refers to an electromechanical interface connected to a WC. The wall switch can convey a tint command to the WC, which can then convey the tint command to the NC. Such devices also are hereinafter collectively referred to as "wall devices," although such devices need not be limited to wall-mounted implementations (for example, such devices also can be located on a ceiling or floor, or integrated on or within a desk or a conference table). For example, some or all of the offices, conference rooms or other rooms of the building can include such a wall device for use in controlling the tint states of the adjoining IGUs 302. For example, the IGUs 302 adjoining a particular room can be grouped into a zone. Each of the wall devices can be operated by an end user (for example, an occupant of the respective room) to control the tint state or other functions or parameters of the IGUs 302 that adjoin the room. For example, at certain times of the day, the adjoining IGUs 302 may be tinted to a dark state to reduce the amount of light energy entering the room from the outside (for example, to reduce AC cooling requirements). Now suppose that a user desires to use the room. In various implementations, the user can operate the wall device to communicate control signals to cause a tint state transition from the dark state to a lighter tint state.

[0090] In some implementations, each wall device can include one or more switches, buttons, dimmers, dials or other physical user interface controls enabling the user to select a particular tint state or to increase or decrease a current tinting level of the IGUs 302 adjoining the room. Additionally or alternatively, the wall device can include a display having a touchscreen interface enabling the user to select a particular tint state (for example, by selecting a virtual button, selecting from a dropdown menu or by entering a tint level or tinting percentage) or to modify the tint state (for example, by selecting a "darken" virtual button, a "lighten" virtual button, or by turning a virtual dial or sliding a virtual bar). In some other implementations, the wall device can include a docking interface enabling a user to physically and communicatively dock a portable device such as a smartphone, multimedia device, tablet computer or other portable computing device (for example, an IPHONE, IPOD or IPAD produced by Apple, Inc. of Cupertino, CA). In such implementations, the user can control the tinting levels via input to the portable device, which is then received by the wall device through the docking interface and subsequently communicated to the MC 308, NC 306 or WC 304. In such implementations, the portable device may include an application for communicating with an API presented by the wall device.

[0091] For example, the wall device can transmit a request for a tint state change to the MC 308. In some implementations, the MC 308 can first determine whether to grant the request (for example, based on power considerations or based on whether the user has the appropriate authorizations/permissions). The MC 308 can then calculate, determine, select or otherwise generate a tint value and transmit the tint value in a primary tint command to cause the tint state transitions in the adjoining IGUs 302. In some such implementations, each wall device can be connected with the MC 308 via one or more wired links (for example, over communication lines such as CAN or Ethernet compliant lines or over power lines using power line communication techniques). In some other implementations, each wall device can be connected with the MC 308 via one or more wireless links. In some other implementations, the wall device can be connected (via one or more wired or wireless connections) with an outward-facing network 310 such as a customer-facing network, which then communicates with the MC 308 via link 318.

[0092] In some implementations, the MC 308 can identify the IGUs 302 associated with the wall device based on previously programmed or discovered information associating the wall device with the IGUs 302. In some implementations, a control algorithm or rule set stored in and executed by the MC 308 can dictate that one or more control signals from a wall device take precedence over a tint value previously generated by the MC 308. In some other instances, such as in times of high demand (for example, high power demand), a control algorithm or rule set stored in and executed by the MC 308 can dictate that the tint value previously generated by the MC 308 takes precedence over any control signals received from a wall device.

[0093] In some other implementations or instances, based on the receipt of a tint-state-change request or control signal from a wall device, the MC 308 can use information about a combination of known parameters to generate a tint value that provides lighting conditions desirable for a typical user, while in some instances also using power efficiently. In some other implementations, the MC 308 can generate the tint value based on preset preferences defined by or for the particular user that requested the tint state change via the wall device. For example, the user may be required to enter a password into the wall device or to use a security token or security fob such as the IBUTTON or other 1-Wire device to gain access to the wall device. In such instances, the MC 308 can determine the identity of the user, based on the password, security token or security fob, retrieve preset preferences for the user, and use the preset preferences alone or in combination with other parameters (such as power considerations or information from various sensors) to calculate, determine, select or otherwise generate a tint value for the respective IGUs 302.

[0094] In some other implementations, the wall device can transmit a tint state change request to the appropriate NC 306, which then communicates the request, or a communication based on

the request, to the MC 308. For example, each wall device can be connected with a corresponding NC 306 via one or more wired links such as those just described for the MC 308 or via a wireless link (such as those described below). In some other implementations, the wall device can transmit a request to the appropriate NC 306, which then itself determines whether to override a primary tint command previously received from the MC 308 or a primary or secondary tint command previously generated by the NC 306 (as described below, the NC 306 can in some implementations generate tint commands without first receiving a tint command from an MC 308). In some other implementations, the wall device can communicate requests or control signals directly to the WC 304 that controls the adjoining IGUs 302. For example, each wall device can be connected with a corresponding WC 304 via one or more wired links such as those just described for the MC 308 or via a wireless link (such as those described below with reference to the WC 600 of Figure 6).

[0095] In some specific implementations, the NC 306 or the MC 308 determines whether the control signals from the wall device should take priority over a tint value previously generated by the NC 306 or the MC 308. As described above, in some implementations, the wall device can communicate directly with the NC 306. However, in some other implementations, the wall device can communicate requests directly to the MC 308 or directly to a WC 304, which then communicates the request to the NC 306. In still other implementations, the wall device can communicate requests to a customer-facing network (such as a network managed by the owners or operators of the building), which then passes the requests (or requests based therefrom) to the NC 306 either directly or indirectly by way of the MC 308. In some implementations, a control algorithm or rule set stored in and executed by the NC 306 or the MC 308 can dictate that one or more control signals from a wall device take precedence over a tint value previously generated by the NC 306 or the MC 308. In some other instances, such as in times of high demand (for example, high power demand), a control algorithm or rule set stored in and executed by the NC 306 or the MC 308 can dictate that the tint value previously generated by the NC 306 or the MC 308 takes precedence over any control signals received from a wall device.

[0096] As described above with reference to the MC 308, in some other implementations, based on the receipt of a tint-state-change request or control signal from a wall device, the NC 306 can use information about a combination of known parameters to generate a tint value that provides lighting conditions desirable for a typical user, while in some instances also using power efficiently. In some other implementations, the NC 306 or the MC 308 can generate the tint value based on preset preferences defined by or for the particular user that requested the tint state change via the wall device. As described above with reference to the MC 308, the user may be required to enter a password into the wall device or to use a security token or security

fob such as the IBUTTON or other 1-Wire device to gain access to the wall device. In such instances, the NC 306 can communicate with the MC 308 to determine the identity of the user, or the MC 308 can alone determine the identity of the user, based on the password, security token or security fob, retrieve preset preferences for the user, and use the preset preferences alone or in combination with other parameters (such as power considerations or information from various sensors) to calculate, determine, select or otherwise generate a tint value for the respective IGUs 302.

[0097] In some implementations, the MC 308 is coupled to an external database (or "data store" or "data warehouse") 320. In some implementations, the database 320 can be a local database coupled with the MC 308 via a wired hardware link 322. In some other implementations, the database 320 can be a remote database or a cloud-based database accessible by the MC 308 via an internal private network or over the outward-facing network 310. In some implementations, other computing devices, systems or servers also can have access to read the data stored in the database 320, for example, over the outward-facing network 310. Additionally, in some implementations, one or more control applications or third party applications also can have access to read the data stored in the database via the outward-facing network 310. In some cases, the MC 308 stores in the database 320 a record of all tint commands including the corresponding tint values issued by the MC 308. The MC 308 also can collect status and sensor data and store it in the database 320. In such instances, the WCs 304 can collect the sensor data and status data from the IGUs 302 and communicate the sensor data and status data to the respective NCs 306 over link 314 for communication to the MC 308 over link 316. Additionally or alternatively, the NCs 306 or the MC 308 themselves also can be connected to various sensors such as light, temperature or occupancy sensors within the building as well as light or temperature sensors positioned on, around or otherwise external to the building (for example, on a roof of the building). In some implementations the NCs 306 or the WCs 304 also can transmit status or sensor data directly to the database 320 for storage.

Integration with Other Systems or Services

[0098] In some implementations, the network system 300 also can be designed to function in conjunction with modern heating, ventilation, and air conditioning (HVAC) systems, interior lighting systems, security systems or power systems as an integrated and efficient energy control system for an entire building or a campus of buildings. Some implementations of the network system 300 are suited for integration with a building management system (BMS), 324. A BMS is broadly a computer-based control system that can be installed in a building to monitor and control the building's mechanical and electrical equipment such as HVAC systems (including furnaces or other heaters, air conditioners, blowers and vents), lighting systems, power systems,

elevators, fire systems, and security systems. The BMS can include hardware and associated firmware and software for maintaining conditions in the building according to preferences set by the occupants or by a building manager or other administrator. The software can be based on, for example, internet protocols or open standards. A BMS can typically be used in large buildings where it functions to control the environment within the building. For example, the BMS can control lighting, temperature, carbon dioxide levels, and humidity within the building. To control the building environment, the BMS can turn on and off various mechanical and electrical devices according to rules or in response to conditions. Such rules and conditions can be selected or specified by a building manager or administrator, for example. One function of a BMS can be to maintain a comfortable environment for the occupants of a building while minimizing heating and cooling energy losses and costs. In some implementations, the BMS can be configured not only to monitor and control, but also to optimize the synergy between various systems, for example, to conserve energy and lower building operation costs.

[0099] Additionally or alternatively, some implementations of the network system 300 are suited for integration with a smart thermostat service, alert service (for example, fire detection), security service or other appliance automation service. On example of a home automation service is NEST®, made by Nest Labs of Palo Alto, California, (NEST® is a registered trademark of Google, Inc. of Mountain View, California). As used herein, references to a BMS can in some implementations also encompass, or be replaced with, such other automation services.

[0100] In some implementations, the MC 308 and a separate automation service, such as a BMS 324, can communicate via an application programming interface (API). For example, the API can execute in conjunction with a master controller application (or platform) within the MC 308, or in conjunction with a building management application (or platform) within the BMS 324. The MC 308 and the BMS 324 can communicate over one or more wired links 326 or via the outward-facing network 310. In some instances, the BMS 324 can communicate instructions for controlling the IGUs 302 to the MC 308, which then generates and transmits primary tint commands to the appropriate NCs 306. In some implementations, the NCs 306 or the WCs 304 also can communicate directly with the BMS 324 (whether through a wired/hardware link or wirelessly through a wireless data link). In some implementations, the BMS 324 also can receive data, such as sensor data, status data and associated timestamp data, collected by one or more of the MC 308, the NCs 306 and the WCs 304. For example, the MC 308 can publish such data over the network 310. In some other implementations in which such data is stored in a database 320, the BMS 324 can have access to some or all of the data stored in the database 320.

Example Master Controller

[0101] Figure 4 shows a block diagram of an example master controller (MC) 400 in accordance with some implementations. For example, the MC 400 of Figure 4 can be used to implement the MC 308 described above with reference to the network system 300 of Figure 3. As used herein, references to "the MC 400" also encompass the MC 308, and vice versa; in other words, the two references may be used interchangeably. The MC 400 can be implemented in or as one or more computers, computing devices or computer systems (herein used interchangeably where appropriate unless otherwise indicated). Additionally, reference to "the MC 400" collectively refers to any suitable combination of hardware, firmware and software for implementing the functions, operations, processes or capabilities described. For example, the MC 400 can refer to a computer that implements a master controller application (also referred to herein as a "program" or a "task").

[0102] As shown in Figure 4, the MC 400 generally includes one or more processors 402 (also collectively referred to hereinafter as "the processor 402"). Processor 402 can be or can include a central processing unit (CPU), such as a single core or a multi-core processor. The processor 402 can additionally include a digital signal processor (DSP) or a network processor in some implementations. In some implementations, the processor 402 also can include one or more application-specific integrated circuits (ASICs). The processor 402 is coupled with a primary memory 404, a secondary memory 406, an inward-facing network interface 408 and an outward-facing network interface 410. The primary memory 404 can include one or more high-speed memory devices such as, for example, one or more random-access memory (RAM) devices including dynamic-RAM (DRAM) devices. Such DRAM devices can include, for example, synchronous DRAM (SDRAM) devices and double data rate SDRAM (DDR SDRAM) devices (including DDR2 SDRAM, DDR3 SDRAM, and DDR4 SDRAM), thyristor RAM (T-RAM), and zero-capacitor (Z-RAM®), among other suitable memory devices.

[0103] The secondary memory 406 can include one or more hard disk drives (HDDs) or one or more solid-state drives (SSDs). In some implementations, the memory 406 can store processor-executable code (or "programming instructions") for implementing a multi-tasking operating system such as, for example, an operating system based on a Linux® kernel. In some other implementations, the operating system can be a UNIX®- or Unix-like-based operating system, a Microsoft Windows®-based operating system, or another suitable operating system. The memory 406 also can store code executable by the processor 402 to implement the master controller application described above, as well as code for implementing other applications or programs. The memory 406 also can store status information, sensor data or other data collected from network controllers, window controllers and various sensors.

[0104] In some implementations, the MC 400 is a "headless" system; that is, a computer that does not include a display monitor or other user input device. In some such implementations, an administrator or other authorized user can log in to or otherwise access the MC 400 from a remote computer or mobile computing device over a network (for example, the network 310) to access and retrieve information stored in the MC 400, to write or otherwise store data in the MC 400, and to control various functions, operations, processes or parameters implemented or used by the MC 400. In some other implementations, the MC 400 also can include a display monitor and a direct user input device (for example, one or more of a mouse, a keyboard and a touchscreen).

[0105] In various implementations, the inward-facing network interface 408 enables the MC 400 to communicate with various distributed controllers, and in some implementations, also with various sensors. The inward-facing network interface 408 can collectively refer to one or more wired network interfaces or one or more wireless network interfaces (including one or more radio transceivers). In the context of the network system 300 of Figure 3, the MC 400 can implement the MC 308 and the inward-facing network interface 408 can enable communication with the downstream NCs 306 over the link 316.

[0106] The outward-facing network interface 410 enables the MC 400 to communicate with various computers, mobile devices, servers, databases or cloud-based database systems over one or more networks. The outward-facing network interface 410 also can collectively refer to one or more wired network interfaces or one or more wireless network interfaces (including one or more radio transceivers). In the context of the network system 300 of Figure 3, the outward-facing network interface 410 can enable communication with various computers, mobile devices, servers, databases or cloud-based database systems accessible via the outward-facing network 310 over the link 318. As described above, in some implementations, the various applications, including third party applications or cloud-based applications, executing within such remote devices can access data from or provide data to the MC 400 or to the database 320 via the MC 400. In some implementations, the MC 400 includes one or more APIs for facilitating communication between the MC 400 and various third party applications. Some example implementations of APIs that the MC 400 can enable are described in PCT Patent Application No. PCT/US15/64555 (Attorney Docket No. VIEWP073WO) filed December 8, 2015 and titled MULTIPLE INTERACTING SYSTEMS AT A SITE, which is hereby incorporated by reference in its entirety and for all purposes. For example, such third party applications can include various monitoring services including thermostat services, alert services (for example, fire detection), security services or other appliance automation services. Additional examples of monitoring services and systems can be found in PCT Patent Application No.

PCT/US2015/019031 (Attorney Docket No. VIEWP061WO) filed March 5, 2015 and titled MONITORING SITES CONTAINING SWITCHABLE OPTICAL DEVICES AND CONTROLLERS, which is hereby incorporated by reference in its entirety and for all purposes.

[0107] In some implementations, one or both of the inward-facing network interface 408 and the outward-facing network interface 410 can include a BACnet compatible interface. BACnet is a communications protocol typically used in building automation and control networks and defined by the ASHRAE/ANSI 135 and ISO 16484-5 standards. The BACnet protocol broadly provides mechanisms for computerized building automation systems and devices to exchange information, regardless of the particular services they perform. For example, BACnet has traditionally been used to enable communication among heating, ventilating, and air-conditioning control (HVAC) systems, lighting control systems, access or security control systems, and fire detection systems as well as their associated equipment. In some other implementations, one or both of the inward-facing network interface 408 and the outward-facing network interface 410 can include an oBIX (Open Building Information Exchange) compatible interface or another RESTful Web Services-based interface. As such, while the following description is sometimes focused on BACnet implementations, in other implementations, other protocols compatible with oBIX or other RESTful Web Services can be used.

[0108] The BACnet protocol is generally based on a server-client architecture. In some implementations, as viewed from the outward-facing network 310, the MC 400 functions as a BACnet server. For example, the MC 400 can publish various information through the outward-facing network interface 410 over the network 310 to various authorized computers, mobile devices, servers or databases, or to various authorized applications executing on such devices. When viewed from the rest of the network system 300, the MC 400 can function as a client. In some such implementations, the NCs 306 function as BACnet servers collecting and storing status data, sensor data or other data acquired from the WCs 304, and publishing this acquired data such that it is accessible to the MC 400.

[0109] The MC 400 can communicate as a client to each of the NCs 306 using BACnet standard data types. Such BACnet data types can include analog values (AVs). In some such implementations, each NC 306 stores an array of AVs. The array of AVs can be organized by BACnet IDs. For example, each BACnet ID can be associated with at least two AVs; a first one of the AVs can be associated with a tint value set by the MC 400 and a second one of the AVs can be associated with a status indication value set (or received) from a respective WC 304. In some implementations, each BACnet ID can be associated with one or more WCs 304. For example, each of the WCs 304 can be identified by a second protocol ID such as a Controller

Area Network (CAN) vehicle bus standard ID (referred to hereinafter as a "CAN ID"). In such implementations, each BACnet ID can be associated with one or more CAN IDs in the NC 306.

[0110] In some implementations, when the MC 400 determines to tint one or more IGUs 302, the MC 400 writes a specific tint value to the AV in the NC 306 associated with the one or more respective WCs 304 that control the target IGUs 302. In some more specific implementations, the MC 400 generates a primary tint command including a BACnet ID associated with the WCs 304 that control the target IGUs 302. The primary tint command also can include a tint value for the target IGUs 302. The MC 400 can direct the transmission of the primary tint command through the inward-facing interface 408 and to the particular NC 306 using a network address of the NC 306. For example, the network address of the NC 306 can include an Internet Protocol (IP) address (for example, an IPv4 or IPv6 address) or a Media Access Control (MAC) address (for example, when communicating over an Ethernet link 316).

[0111] The MC 400 can calculate, determine, select or otherwise generate a tint value for one or more IGUs 302 based on a combination of parameters. For example, the combination of parameters can include time or calendar information such as the time of day, day of year or time of season. Additionally or alternatively, the combination of parameters can include solar calendar information such as, for example, the direction of the sun relative to the IGUs 302. In some instances, the direction of the sun relative to the IGUs 302 can be determined by the MC 400 based on time and calendar information together with information known about the geographical location of the building on the Earth and the direction that the IGUs face (for example, in a North-East-Down coordinate system). The combination of parameters also can include the outside temperature (external to the building), the inside temperature (within a room adjoining the target IGUs 302), or the temperature within the interior volume of the IGUs 302. The combination of parameters also can include information about the weather (for example, whether it is clear, sunny, overcast, cloudy, raining or snowing). Parameters such as the time of day, day of year, or direction of the sun can be programmed into and tracked by the MC 308. Parameters such as the outside temperature, inside temperature or IGU temperature can be obtained from sensors in, on or around the building or sensors integrated on or within the IGUs 302. Some information about the weather also can be obtained from such sensors. Additionally or alternatively, parameters such as the time of day, time of year, direction of the sun, or weather can be provided by, or determined based on information provided by, various applications including third party applications over the network 310. Additional examples of algorithms, routines, modules, or other means for generating tint values are described in U.S. Patent Application No. 13/722,969 (Attorney Docket No. VIEWP049) filed February 21, 2013 and titled CONTROL METHOD FOR TINTABLE WINDOWS, and in PCT Patent Application No.

PCT/2015/029675 (Attorney Docket No. VIEWP049X1WO) filed May 7, 2015 and titled CONTROL METHOD FOR TINTABLE WINDOWS, both of which are hereby incorporated by reference in their entireties and for all purposes.

[0112] Generally, each ECD within each IGU 302 is capable of being tinted, responsive to a suitable driving voltage applied across the EC stack, to virtually any tint state within a continuous tint spectrum defined by the material properties of the EC stack. However, in some implementations, the MC 400 is programmed to select a tint value from a finite number of discrete tint values. For example, the tint values can be specified as integer values. In some such implementations, the number of available discrete tint values can be 4, 8, 16, 32, 64, 128 or 256 or more. For example, a 2-bit binary number can be used to specify any one of four possible integer tint values, a 3-bit binary number can be used to specify any one of eight possible integer tint values, a 4-bit binary number can be used to specify any one of sixteen possible integer tint values, a 5-bit binary number can be used to specify any one of thirty-two possible integer tint values, and so on. Each tint value can be associated with a target tint level (for example, expressed as a percentage of maximum tint, maximum safe tint, or maximum desired or available tint). For didactic purposes, consider an example in which the MC 400 selects from among four available tint values: 0, 5, 10 and 15 (using a 4-bit or higher binary number). The tint values 0, 5, 10 and 15 can be respectively associated with target tint levels of 60%, 40%, 20% and 4%, or 60%, 30%, 10% and 1%, or another desired, advantageous, or suitable set of target tint levels.

Example Network Controller

[0113] Figure 5 shows a block diagram of an example network controller (NC) 500 in accordance with some implementations. For example, the NC 500 of Figure 5 can be used to implement the NC 306 described above with reference to the network system 300 of Figure 3. As used herein, references to "the NC 500" also encompass the NC 306, and vice versa; in other words, the two references may be used interchangeably. The NC 500 can be implemented in or as one or more network components, networking devices, computers, computing devices or computer systems (herein used interchangeably where appropriate unless otherwise indicated). Additionally, reference to "the NC 500" collectively refers to any suitable combination of hardware, firmware and software for implementing the functions, operations, processes or capabilities described. For example, the NC 500 can refer to a computer that implements a network controller application (also referred to herein as a "program" or a "task").

[0114] As shown in Figure 5, the NC 500 generally includes one or more processors 502 (also collectively referred to hereinafter as "the processor 502"). In some implementations, the processor 502 can be implemented as a microcontroller or as one or more logic devices including

one or more application-specific integrated circuits (ASICs) or programmable logic devices (PLDs), such as field-programmable gate arrays (FPGAs) or complex programmable logic devices (CPLDs). If implemented in a PLD, the processor can be programmed into the PLD as an intellectual property (IP) block or permanently formed in the PLD as an embedded processor core. In some other implementations, the processor 502 can be or can include a central processing unit (CPU), such as a single core or a multi-core processor. The processor 502 is coupled with a primary memory 504, a secondary memory 506, a downstream network interface 508 and an upstream network interface 510. In some implementations, the primary memory 504 can be integrated with the processor 502, for example, as a system-on-chip (SOC) package, or in an embedded memory within a PLD itself. In some other implementations, the NC 500 alternatively or additionally can include one or more high-speed memory devices such as, for example, one or more RAM devices.

[0115] The secondary memory 506 can include one or more solid-state drives (SSDs) storing one or more lookup tables or arrays of values. In some implementations, the secondary memory 506 can store a lookup table that maps first protocol IDs (for example, BACnet IDs) received from the MC 400 to second protocol IDs (for example, CAN IDs) each identifying a respective one of the WCs 304, and vice versa. In some implementations, the secondary memory 506 can additionally or alternatively store one or more arrays or tables. In some implementations, such arrays or tables can be stored as comma-separated values (CSV) files or via another table-structured file format. For example, each row of the file can be identified by a timestamp corresponding to a transaction with a WC 304. Each row can include a tint value (C) for the IGUs 302 controlled by the WC 304 (for example, as set by the MC 400 in the primary tint command); a status value (S) for the IGUs 302 controlled by the WC 304; a set point voltage (for example, the effective applied voltage V_{Eff}) an actual voltage level V_{Act} measured, detected or otherwise determined across the ECDs within the IGUs 302; an actual current level I_{Act} measured, detected or otherwise determined through the ECDs within the IGUs 302; and various sensor data. In some implementations, each row of the CSV file can include such status information for each and all of the WCs 304 controlled by the NC 500. In some such implementations, each row also includes the CAN IDs or other IDs associated with each of the respective WC 304.

[0116] In some implementations in which the NC 500 is implemented in a computer that executes a network controller application, the secondary memory 506 also can store processor-executable code (or "programming instructions") for implementing a multi-tasking operating system such as, for example, an operating system based on a Linux® kernel. In some other implementations, the operating system can be a UNIX®- or Unix-like-based operating system, a

Microsoft Windows®-based operating system, or another suitable operating system. The memory 506 also can store code executable by the processor 502 to implement the network controller application described above, as well as code for implementing other applications or programs.

[0117] In various implementations, the downstream network interface 508 enables the NC 500 to communicate with distributed WCs 304, and in some implementations, also with various sensors. In the context of the network system 300 of Figure 3, the NC 500 can implement the NC 306 and the downstream network interface 508 can enable communication with the WCs 304 over the link 314. The downstream network interface 508 can collectively refer to one or more wired network interfaces or one or more wireless network interfaces (including one or more radio transceivers). In some implementations, the downstream interface 508 can include a CANbus interface enabling the NC 500 to distribute commands, requests or other instructions to various WCs 304, and to receive responses including status information from the WCs 304, according to a CANBus protocol (for example, via the CANopen communication protocol). In some implementations, a single CANbus interface can enable communication between the NC 500 and tens, hundreds or thousands of WCs 304. Additionally or alternatively, the downstream interface 508 can include one or more Universal Serial Bus (USB) interfaces (or "ports"). In some such implementations, to enable communication via a CANbus communication protocol, a USB-to-CAN adapter can be used to couple the USB port of the downstream interface 508 with CANbus-compatible cables. In some such implementations, to enable the NC 500 to control even more WCs 304, a USB hub (for example, having 2, 3, 4, 5 10 or more hub ports) can be plugged into the USB port of the downstream interface 508. A USB-to-CAN adapter can then be plugged into each hub port of the USB hub.

[0118] The upstream network interface 510 enables the NC 500 to communicate with the MC 400, and in some implementations, also with various other computers, servers or databases (including the database 320). The upstream network interface 510 also can collectively refer to one or more wired network interfaces or one or more wireless network interfaces (including one or more radio transceivers). In the context of the network system 300 of Figure 3, the upstream network interface 510 can enable communication with the MC 308 over the link 318. In some implementations, the upstream network interface 510 also can be coupled to communicate with applications, including third party applications and cloud-based applications, over the outward-facing network 310. For example, in implementations in which the NC 500 is implemented as a network controller application executing as a task within a computer, the network controller application can communicate directly with the outward-facing network 310 via the operating system and the upstream network interface 510. In some other implementations, the NC 500 may

be implemented as a task running on the MC 308 and managing the CANbus devices via the CANbus interface. In such implementations, in addition or as an alternative to TCP/IP or UDP/IP communications to the MC, the communications could be via UNIX Domain Sockets (UDS) or other communication methods like shared memory, or other non-IP communication methods.

[0119] In some implementations, the upstream interface 510 can include BACnet compatible interface, an oBIX compatible interface or another RESTful Web Services-based interface. As described above with reference to Figure 4, in some implementations the NC 500 functions as a BACnet server collecting and storing status data, sensor data or other data acquired from the WCs 304, and publishing this acquired data such that it is accessible to the MC 400. In some implementations, the NC 500 also can publish this acquired data over the network 310 directly; that is, without first passing the data to the MC 400. The NC 500 also functions in some respects similar to a router. For example, the NC 500 can function as a BACnet to CANBus gateway, receiving communications transmitted from the MC 400 according to the BACnet protocol, converting commands or messages from the BACnet protocol to a CANBus protocol (for example, the CANopen communication protocol), and distributing commands or other instructions to various WCs 304 according to the CANBus protocol.

[0120] BACnet is built over the user datagram protocol (UDP). In some other implementations, a non-broadcast-based communication protocol can be used for communication between the MC 400 and the NCs 500. For example, the transmission control protocol (TCP) can serve as the transport layer as opposed to UDP. In some such implementations, the MC 400 can communicate with the NCs 500 via an oBIX-compatible communication protocol. In some other implementations, the MC 400 can communicate with the NCs 500 via a WebSocket-compatible communication protocol. Such TCP protocols also can allow the NCs 500 to communicate directly with one another.

[0121] In various implementations, the NC 500 can be configured to perform protocol translation (or "conversion") between one or more upstream protocols and one or more downstream protocols. As described above, the NC 500 can perform translation from BACnet to CANopen, and vice versa. As another example, the NC 500 can receive upstream communications from the MC 400 via an oBIX protocol and translate the communications into CANopen or other CAN-compatible protocols for transmission to the downstream WCs 304, and vice versa. In some wireless implementations, the NC 500 or the MC 400 also can translate various wireless protocols including, for example, protocols based on the IEEE 802.11 standard (for example, WiFi), protocols based on the IEEE 802.15.4 standard (for example, ZigBee, 6LoWPAN, ISA100.11a, WirelessHART or MiWi), protocols based on the Bluetooth standard

(including the Classic Bluetooth, Bluetooth high speed and Bluetooth low energy protocols and including the Bluetooth v4.0, v4.1 and v4.2 versions), or protocols based on the EnOcean standard (ISO/IEC 14543-3-10). For example, the NC 500 can receive upstream communications from the MC 400 via an oBIX protocol and translate the communications into WiFi or 6LowPAN for transmission to the downstream WCs 304, and vice versa. As another example, the NC 500 can receive upstream communications from the MC 400 via WiFi or 6LowPAN and translate the communications into CANopen for transmission to the downstream WCs 304, and vice versa. In some other examples, the MC 400 rather than the NC 500 handles such translations for transmission to downstream WCs 304.

[0122] As described above with reference to Figure 4, when the MC 400 determines to tint one or more IGUs 302, the MC 400 can write a specific tint value to the AV in the NC 500 associated with the one or more respective WCs 304 that control the target IGUs 302. In some implementations, to do so, the MC 400 generates a primary tint command communication including a BACnet ID associated with the WCs 304 that control the target IGUs 302. The primary tint command also can include a tint value for the target IGUs 302. The MC 400 can direct the transmission of the primary tint command to the NC 500 using a network address such as, for example, an IP address or a MAC address. Responsive to receiving such a primary tint command from the MC 400 through the upstream interface 510, the NC 500 can unpackage the communication, map the BACnet ID (or other first protocol ID) in the primary tint command to one or more CAN IDs (or other second protocol IDs), and write the tint value from the primary tint command to a first one of the respective AVs associated with each of the CAN IDs.

[0123] In some implementations, the NC 500 then generates a secondary tint command for each of the WCs 304 identified by the CAN IDs. Each secondary tint command can be addressed to a respective one of the WCs 304 by way of the respective CAN ID. Each secondary tint command also can include the tint value extracted from the primary tint command. The NC 500 transmits the secondary tint commands to the target WCs 304 through the downstream interface 508 via a second communication protocol (for example, via the CANOpen protocol). In some implementations, when a WC 304 receives such a secondary tint command, the WC 304 transmits a status value back to the NC 500 indicating a status of the WC 304. For example, the tint status value can represent a "tinting status" or "transition status" indicating that the WC is in the process of tinting the target IGUs 302, an "active" or "completed" status indicating that the target IGUs 302 are at the target tint state or that the transition has been finished, or an "error status" indicating an error. After the status value has been stored in the NC 500, the NC 500 can publish the status information or otherwise make the status information accessible to the MC 400 or to various other authorized computers or

applications. In some other implementations, the MC 400 can request status information for a particular WC 304 from the NC 500 based on intelligence, a scheduling policy, or a user override. For example, the intelligence can be within the MC 400 or within a BMS. A scheduling policy can be stored in the MC 400, another storage location within the network system 300, or within a cloud-based system.

Integrated Master Controller and Network Controller

[0124] As described above, in some implementations the MC 400 and the NC 500 can be implemented as a master controller application and a network controller application, respectively, executing within respective physical computers or other hardware devices. In some alternative implementations, each of the master controller application and the network controller application can be implemented within the same physical hardware. For example, each of the master controller application and the network controller application can be implemented as a separate task executing within a single computer device that includes a multi-tasking operating system such as, for example, an operating system based on a Linux® kernel or another suitable operating system.

[0125] In some such integrated implementations, the master controller application and the network controller application can communicate via an application programming interface (API). In some particular implementations, the master controller and network controller applications can communicate over a loopback interface. By way of reference, a loopback interface is a virtual network interface, implemented through an operating system, which enables communication between applications executing within the same device. A loopback interface is typically identified by an IP address (often in the 127.0.0.0/8 address block in IPv4, or the 0:0:0:0:0:0:0:1 address (also expressed as ::1) in IPv6). For example, the master controller application and the network controller application can each be programmed to send communications targeted to one another to the IP address of the loopback interface. In this way, when the master controller application sends a communication to the network controller application, or vice versa, the communication does not need to leave the computer.

[0126] In implementations in which the MC 400 and the NC 500 are implemented as master controller and network controller applications, respectively, there are generally no restrictions limiting the available protocols suitable for use in communication between the two applications. This generally holds true regardless of whether the master controller application and the network controller application are executing as tasks within the same or different physical computers. For example, there is no need to use a broadcast communication protocol, such as BACnet, which limits communication to one network segment as defined by a switch or router boundary.

For example, the oBIX communication protocol can be used in some implementations for communication between the MC 400 and the NCs 500.

[0127] In the context of the network system 300, each of the NCs 500 can be implemented as an instance of a network controller application executing as a task within a respective physical computer. In some implementations, at least one of the computers executing an instance of the network controller application also executes an instance of a master controller application to implement the MC 400. For example, while only one instance of the master controller application may be actively executing in the network system 300 at any given time, two or more of the computers that execute instances of network controller application can have an instance of the master controller application installed. In this way, redundancy is added such that the computer currently executing the master controller application is no longer a single point of failure of the entire system 300. For example, if the computer executing the master controller application fails or if that particular instance of the master controller application otherwise stops functioning, another one of the computers having an instance of the master network application installed can begin executing the master controller application to take over for the other failed instance. In some other applications, more than one instance of the master controller application may be executing concurrently. For example, the functions, processes or operations of the master controller application can be distributed to two (or more) instances of the master controller application.

Example Window Controller

[0128] Figure 6 shows a circuit schematic diagram of an example window controller (WC) 600 in accordance with some implementations. For example, the WC 600 of Figure 6 can be used to implement each one of the WCs 304 described above with reference to the network system 300 of Figure 3. As used herein, references to "the WC 600" also encompass the WC 304, and vice versa; in other words, the two references may be used interchangeably. As described above, the WC 600 is generally operable and adapted to drive optical state transitions in, or to maintain the optical states of, one or more coupled optically-switchable devices such as the ECDs 110 described above with reference to Figure 1. In some implementations, the one or more ECDs coupled with the WC 600 are configured within respective IGUs 602 (such as the IGU 100 described above with reference to Figure 1). The WC 600 also is operable to communicate with the coupled IGUs 602, for example, to read data from or to transfer data to the IGUs 602.

[0129] The WC 600 broadly includes a processing unit 604. The WC 600 also broadly includes a power circuit 606, a drive circuit 608 and a feedback circuit 610 (each of which is delineated with a heavy dashed line and gray shading). In the illustrated implementation, the

WC 600 additionally includes a communications circuit 612. Each of the drive circuit 608, the power circuit 606, the feedback circuit 610 and the communications circuit 612 can include a number of individual circuit components including integrated circuits (ICs). Each of the various components described in more detail below may be described as being "a part of" a respective one of the aforementioned circuits 606, 608, 610 and 612. However, the groupings of components into respective ones of the circuits 606, 608, 610 and 612 are in name only and for purposes of convenience in facilitating the disclosure of the described implementations. As such, the functions, capabilities and limitations of the various described components are not intended to be defined by the respective grouping; rather, the functions, abilities and limitations of each of the individual components are defined only by those of the components themselves, and by their integration with other components to which they are electrically connected or coupled.

[0130] The WC 600 includes a first upstream interface (or set of interfaces) 614 for coupling to an upstream set of cables 616. For example, the upstream set of cables 616 can implement the link 314 described above with reference to the network system 300 Figure 3. In some implementations, the upstream set of cables 616 includes at least four lines: two power distribution lines and two communication lines. In some five-line implementations, the upstream set of cables 616 additionally includes a system ground line, such as a building ground or Earth ground (for practical purposes an absolute ground from which all other voltages in the building can be measured). The upstream interface 614 can include a corresponding number of pins (not shown)—one pin to couple each of the lines in the upstream set of cables 616 into the WC 600. For example, a first one of the pins can couple a first one of the power distribution lines from the upstream set of cables 616 to a first power supply line 622 within the WC 600. A second one of the pins can couple a second one of the power distribution lines (for example, a power supply return) from the upstream set of cables 616 to a second power supply line 624 within the WC 600. A third one of the pins can couple a first one of the communication lines from the upstream set of cables 616 to a first communication line 626 within the WC 600. A fourth one of the pins can couple a second one of the communication lines from the upstream set of cables 616 to a second communication line 628 within the WC 600. In implementations that include a system ground line, a fifth one of the pins can couple the system ground line from the upstream set of cables 616 to a system ground line 630 within the WC 600.

[0131] The two power distribution lines in the upstream set of cables 616 can be implemented as two separate cables or configured together as, for example, a twisted pair cable. The first power line 622 carries a first supply voltage V_{Sup1} and the second power line 624 is a power supply return. In some implementations, the first supply voltage V_{Sup1} is a DC voltage having a

value in the range of approximately 5 Volts (V) to 42 V, and in one example application, a value of 24 V (although higher voltages may be desirable and are possible in other implementations). In some other implementations, the first supply voltage V_{Sup1} can be a pulsed voltage power signal. As described above, the second one of the power lines 624 can be a power supply return, also referred to as a signal ground (or "common ground"). In other words, the voltage V_{Sup2} on the second one of the power lines can be a reference voltage, for example, a ground. In such implementations, it is the voltage difference between the first supply voltage V_{Sup1} and the second supply voltage V_{Sup2} that is the voltage of interest, as opposed to the actual values of the individual voltages V_{Sup1} and V_{Sup2} relative to the system ground. For example, the value of the difference between V_{Sup1} and V_{Sup2} can be in the range of approximately 5 V to 42 V, and in one example application, 24 V. In implementations that include a system ground line, the system ground line can be implemented as a single cable or configured with the two power distribution lines described above as a 3-wire cable.

[0132] The two communication lines in the upstream set of cables 616 also can be implemented as two separate cables or configured together as a twisted pair cable. In some other implementations, the two communication lines can be bundled with the two power distribution lines just described as a 4-wire cable, or bundled with the two power distribution lines and the system ground line as a 5-wire cable. As described above, pins or other interconnects within the upstream interface 614 electrically connect the first and the second communication lines in the upstream set of cables 616 with the first and the second communication lines 626 and 628, respectively, in the WC 600. The first and the second communication lines 626 and 628, also referred to herein collectively as a communication bus 632, can carry first and second data signals $Data_1$ and $Data_2$, respectively.

[0133] At different times or stages throughout an optical transition cycle or at other times, the data signals $Data_1$ and $Data_2$ can be communicating information to the WC 600 from an upstream network controller (such as the NC 306 or NC 400) or communicating information to the network controller from the WC 600. As an example of a downstream communication, the data signals $Data_1$ and $Data_2$ can include a tint command or other instructions (for example, such as the secondary tint command described above) sent from a network controller to the WC 600. As an example of an upstream communication, the data signals $Data_1$ and $Data_2$ can include status information (such as a current tint status) or sensor data to be sent to the network controller. In some implementations, the signals $Data_1$ and $Data_2$ are complementary signals, for example, forming a differential pair of signals (also referred to herein collectively as a differential signal).

[0134] In some implementations, the communication bus 632 is designed, deployed and otherwise configured in accordance with the Controller Area Network (CAN) vehicle bus standard. In terms of the Open Systems Interconnection (OSI) model, the physical (PHY) layer can be implemented according to the ISO 11898-2 CAN standard, and the data link layer can be implemented according to the ISO 11898-1 CAN standard. In some such implementations, the first data signal $Data_1$ can refer to the high CAN signal (the "CANH signal" as it is typically referred to in the CAN protocol), while the second data signal $Data_2$ can refer to the low CAN signal (the "CANL signal"). In some implementations, the WC 600 communicates with the upstream network controller over the communication bus 632 (and the coupled communication lines in the upstream set of cables 616) according to the CANopen communication protocol. In terms of the OSI model, the CANopen communication protocol implements the network layer and other layers above the network layer (for example, the transport layer, the session layer, the presentation layer and the application layer). According to the CAN protocol, it is the difference between the CANH and CANL signal values that determines the value of the bit being communicated by the differential pair.

[0135] In some implementations, the upstream set of cables 616 is directly connected with the upstream network controller. In some other implementations, the upstream set of cables 616 includes a set of droplines connected to (for example, tapped off of) a trunk line that contains corresponding power distribution and communication lines. In some such latter implementations, each of a plurality of WCs 600 can be connected to the same trunk line via a corresponding set of droplines. In some such implementations, each of the plurality of WCs 600 coupled to the same trunk line can be in communication with the same network controller via the communication lines within the trunk line. In some implementations, the power distribution lines that power the WCs 600 also can be coupled to the same network controller to power the network controller. In some other implementations, a different set of power distribution lines can power the network controller. In either case, the power distribution lines that power the WCs 600 can terminate at a power control panel or other power insertion point.

[0136] The WC 600 also includes a second downstream interface (or set of interfaces) 618 for coupling to a downstream set of cables 620. For example, the downstream set of cables 620 can implement the link 312 described above with reference to the network system 300 Figure 3. In some implementations, the downstream set of cables 620 also includes at least four lines: two power distribution lines and two communication lines. The downstream interface 618 also can include a corresponding number of pins (not shown)—one pin to couple each of the lines in the downstream set of cables 620 into the WC 600. For example, a first one of the pins can couple a first one of the power distribution lines 633 from the downstream set of cables 620 to a first

power drive line 634 within the WC 600. A second one of the pins can couple a second one of the power distribution lines 635 from the downstream set of cables 620 to a second power drive line 636 within the WC 600. A third one of the pins can couple a first one of the communication lines 637 from the downstream set of cables 620 to a first communication line 638 within the WC 600. A fourth one of the pins can couple a second one of the communication lines 639 from the downstream set of cables 620 to a second communication line 640 within the WC 600. In implementations that include a fifth line, a fifth one of the pins can couple the fifth line 641 from the downstream set of cables 620 to a fifth line 642 within the WC 600.

[0137] The two power distribution lines 633 and 635 in the downstream set of cables 620 can be implemented as two separate cables or configured together as, for example, a twisted pair cable. In some implementations, the first power distribution line 633 carries a first applied voltage V_{App1} and the second power distribution line 635 carries a second applied voltage V_{App2}. In some implementations, the first and the second applied voltages V_{App1} and V_{App2} are, for all intents and purposes, DC voltage signals. In some other implementations, the first and the second applied voltages V_{App1} and V_{App2} can be pulsed voltage signals (for example, pulse-width modulated (PWM) signals). In some implementations, the first applied voltage V_{App1} can have a value in the range of approximately 0 V to 10 V, and in some specific applications, in the range of approximately 0 V to 5 V. In some implementations, the second applied voltage V_{App2} can have a value in the range of approximately 0 V to -10 V, and in some specific applications, in the range of approximately 0 V to -5 V. In some other implementations, the second power distribution line 635 in the downstream set of cables 620 can be a power supply return, also referred to as a signal ground or common ground. In other words, the voltage V_{App2} on the second power distribution line can be a reference voltage, for example, a floating ground.

[0138] The first and the second power distribution lines 633 and 635 in the downstream set of cables 620 are provided to each of the one or more IGUs 602 controlled by the WC 600. More specifically, the first and the second power distribution lines 633 and 635 are electrically connected to (or coupled with) the busbars and conductive layers that power the electrochromic states and state transitions of the respective ECDs (such as, for example, the first and second busbars 126 and 128 and the first and second TCO layers 114 and 116 in the IGU 100 of Figure 1). In some implementations, it is the voltage difference between the first applied voltage V_{App1} and the second applied voltage V_{App2} that is the voltage of interest, as opposed to the actual values of the individual voltages V_{App1} and V_{App2} relative to a system ground. For example, the value of the difference between V_{App1} and V_{App2}—referred to herein as the "effective applied voltage" V_{Eff} or simply as the applied voltage V_{Eff}—can be in the range of approximately -10 V to

10 V in some applications, and in some specific applications in the range of approximately -5 V to 5 V, depending on various device parameters and drive parameters.

[0139] The two communication lines 637 and 639 in the downstream set of cables 620 also can be implemented as two separate cables or configured together as a twisted pair cable. In some other implementations, the two communication lines 637 and 639 can be bundled with the two power distribution lines 633 and 635 just described as a 4-wire cable, or bundled with the two power distribution lines and the fifth line as a 5-wire cable. As described above, pins or other interconnects within the downstream interface 618 electrically connect the first and the second communication lines 637 and 639 in the downstream set of cables 620 with the first and the second communication lines 638 and 640 within the WC 600. The first and the second communication lines 638 and 640, also referred to herein collectively as a communication bus 644, can carry data signals $Data_3$ and $Data_4$, respectively.

[0140] At different times or stages throughout a transition cycle or at other times, the data signals $Data_3$ and $Data_4$ can be communicating information to one or more connected IGUs 602 from the WC 600 or communicating information to the WC 600 from one or more of the IGUs 602. As an example of a downstream communication, the data signals $Data_3$ and $Data_4$ can include a status request command or other instructions to be sent to one or more of the IGUs 602. As an example of an upstream communication, the data signals $Data_3$ and $Data_4$ can include status information (such as a current tint status) or sensor data sent from one or more of the IGUs 602 to the WC 600. In some implementations, the communication bus 644 is designed, deployed and otherwise configured in accordance with the 1-Wire device communications bus system protocol. In such 1-Wire implementations, the communication line 638 is a data line and the data signal $Data_3$ conveys the data to be communicated, while the communication line 640 is a signal ground line and the data signal $Data_4$ provides a reference voltage, such as a signal ground, relative to which the data signal $Data_3$ is measured or compared to recover the data of interest.

Example Connection Architecture

[0141] In some implementations, the downstream set of cables 620 is directly connected with a single IGU 602. In some other implementations, the downstream set of cables 620 includes a junction that connects the downstream set of cables 620 to two or more IGUs 602 via corresponding sets of cables. Figure 7 shows a diagram of an example connection architecture 700 for coupling a window controller to an IGU in accordance with some implementations. In the illustrated implementation, the connection architecture 700 couples the WC 600 to an IGU 602 that includes an ECD 746 (only an end portion of the IGU 602 and ECD 746 are shown). While only one IGU 602 is shown, as described above, the connection architecture 700 can

couple the WC 600 to multiple IGUs 602. To facilitate such multi-IGU implementations, the downstream set of cables 620 can connect the WC 600 with a junction 748. In some implementations, the junction 748 electrically couples each of the lines 633, 635, 637, 639 and 641 within the downstream set of cables 620 to corresponding lines 734, 736, 738, 740 and 742 in each of multiple secondary sets of cables 750_1–750_N. In this way, a single WC 600 can provide power to multiple IGUs 602.

[0142] In the illustrated diagrammatic implementation, the IGU 602 includes a plug-in component 752 that facilitates the connection of the downstream set of cables 620, or more particularly the secondary set of cables 750_1, with the IGU 602 and the ECD 746 within it. In some implementations, the plug-in component 752 is readily insertable and removable from the IGU 602 (for example, for ease of manufacture, maintenance, or replacement). As shown, the plug-in component 752 includes an interface 754 (which can be similar to the interface 618 of the WC 600) for receiving the power distribution lines 734 and 736, the communication lines 738 and 740 and the fifth line 742 (in implementations that include a fifth line). In some implementations, the ends of the lines 734, 736, 738, 740 and 742 can include connectors that are adapted to be inserted within corresponding connection receivers within the interface 754. The plug-in component 752 serves to electrically couple power distribution lines 734 and 736 with bus bars 758 and 760, respectively. Bus bars 758 and 760 are, in turn, electrically connected to respective conducting layers on either side of the EC stack of the ECD 746.

[0143] The plug-in component 752 includes a communication module 756 that is connected to transmit and receive data to and from the WC 600 over the communication lines 738 and 740. In some implementations, the communication module 756 can be implemented as a single chip. In some such implementations, the communication module 756 can be implemented as a 1-Wire chip that includes a non-volatile memory such as, for example, EEPROM (E^2PROM), Flash or other suitable solid state memory. Each communication module 756 also can include various processing, controller and logic functionalities, authentication capabilities, or other functionalities or capabilities. When implemented as a 1-Wire chip, each communication module 756 can be identified with a unique 1-Wire ID (for example, a 48-bit serial number). One example of such a 1-Wire chip suitable for use in some implementations is the DS28EC20, 20Kb 1-wire EPROM chip provided by Maxim Integrated Products, Inc. of San Jose, CA. In some other implementations, the communication module 756 can include a memory chip (including non-volatile memory and memory controller functionality) and a separate ID chip storing the unique ID (for example, the 1-Wire ID). Some examples of functions and hardware that can be associated with such a 1-Wire chip are described in U.S. Patent Application No. 13/049,756 (Attorney Docket No. VIEWP007) filed March 16, 2011 and titled

MULTIPURPOSE CONTROLLER FOR MULTISTATE WINDOWS, which is hereby incorporated by reference in its entirety and for all purposes.

[0144] In some implementations, various device or drive parameters for the particular ECD 746 are programmed into and stored within the memory component within the communication module 756 (for example, during or at the end of manufacturing or fabrication of the ECD or IGU or at a later time during or after installation). For example, such pre-programmed device parameters for the ECD 746 can include a length, width, thickness, cross-sectional area, shape, age, model number, version number, or number of previous optical transitions of or associated with the respective ECD 746 (or of a pane on which the ECD is formed or otherwise arranged). Pre-programmed drive parameters can include, for example, a ramp-to-drive rate, a drive voltage, a drive voltage duration, a ramp-to-hold rate and a holding voltage for each possible combination of current tint state and target tint state. In some implementations, the processing unit 604 reads the device parameters and drive parameters prior to the start of each tint state transition. Additionally or alternatively, in some implementations, the processing unit 604 reads the device and drive parameters when the respective IGU 602 is powered on and commissioned. The processing unit 604 can additionally or alternatively read the device and drive parameters periodically, such as daily.

[0145] In some other implementations, a surface of the communication module 756 can additionally or alternatively have an identifier (ID) scribed or etched on it. For example, the ID can be scribed or etched on the communication module 756 during or after production of the ECD. In some implementations, the ID is a lite ID of the lite (pane) on which the ECD is formed. Additionally or alternatively, the ID can include an IGU ID of the associated IGU 302. In some implementations, the WC 304 will then read this information optically or electronically after it is connected to the ECD. In some such implementations, the WC 304 can retrieve parameters such as the length, width, thickness, cross-sectional area, shape, age, model number, version number etc. from the MC 308. For example, the MC 400 can previously be programmed to store such parameters. In some other implementations, the MC 400 can retrieve such parameters from the producer of the ECD/IGU through an external communication interface (for example, the interface 410) either in advance or in response to a request for such parameters or related information by the WC 304 or NC 306.

[0146] The number and size of the IGUs 602 that each WC 600 can drive is generally limited by the load on the WC 600. The load is typically defined by the voltage, current, or power requirements necessary to cause the desired optical transitions in the IGUs 602 driven by the WC 600 within a desired timeframe. Because the maximum load that a given WC 600 can drive is generally limited by the capabilities and safe operating ranges of the electrical components

within the WC 600, or by the power carrying limitations of the power drive lines 634 and 636 or the power distribution lines 633 and 635, there can be a tradeoff between acceptable transition time and the number and size of the ECDs driven by each WC 600.

[0147] The power requirements necessary to cause the desired optical transitions in the IGUs 602 driven by a given WC 600 within a desired timeframe are, in turn, a function of the surface area of the connected IGUs 602, and more particularly, the surface area of the ECDs within the IGUs 602. This relationship can be nonlinear; that is, the power requirements can increase nonlinearly with the surface area of the ECDs. The nonlinear relationship can exist, at least in part, because the sheet resistances of the conductive layers (such as the first and second TCO layers 114 and 116 of the IGU 100) used to deliver the applied voltages to the electrochromic stack of the ECD increase nonlinearly with distance across the length and width of the respective conductive layers. For example, it can take more power to drive a single 50 ft^2 ECD than to drive two 25 ft^2 ECDs. System- or building-wide power considerations also may require that the power available to each WC 600 be limited to less than that which the WC 600 is capable of handling and providing to the connected IGUs 602.

[0148] In some implementations, such as that described with reference to the connection architecture 700 of Figure 7, each of the IGUs 602 connected with the WC 600 can include its own respective plug-in component 752 and communication module 756. Each communication module 756 can include a respective 1-Wire chip storing device parameters for the respective ECD. In some implementations, each of the parallel-connected IGUs 602 receives the same voltages V_{App1} and V_{App2}. In some such implementations, it can generally be desirable or preferable for each of the IGUs 602 connected with a single WC 600 to have the same or similar device parameters (such as surface area) so that each of the respective ECDs behaves the same or similarly responsive to the voltages V_{App1} and V_{App2}. For example, it is generally desirable that each of the IGUs 602 connected with a given WC 600 have the same tint whether during a transition or during a holding period between transitions. However, in implementations in which the IGUs 602 have different device parameters, the processing unit 604 can compare or otherwise integrate the device parameters from each of the connected IGUs 602 to generate a command signal V_{Drive} that results in a best or least harmful effective applied voltage V_{Eff}, for example, a voltage that is maintained within a safe but effective range for all of the connected IGUs 602.

[0149] In some other implementations, there can be a one-to-one relationship between the number of WCs 600 and IGUs 602; that is, each IGU 602 can be driven and otherwise controlled by a respective dedicated WC 600. In some such integrated implementations, the WC 600 can be located within the IGU 602, for example, within a housing having a thin form factor within

the interior volume of the IGU. In some other implementations, the WC 600 can be located adjacent the IGU 602, for example, hidden by a frame or mullion that supports the IGU 602. In some other implementations, the WC 600 can be located at an interior lower boundary or at an interior corner of the IGU 602 where it is less visible or noticeable but still accessible to an installer or technician. For example, such latter implementations can be useful for applications in which easier access to the WC 600 is desirable (for example, to replace, repair or map the WC 600).

[0150] Additionally, such implementations also can be desirable where the WC 600 can include an energy storage device (for example, a rechargeable battery, battery pack or supercapacitor), that is also readily replaceable by a technician. For example, the IGU can include a docking module that the battery can plug into. In such case, the docking module can be electrically connected to the WC 600 rather than the battery directly. In implementations in which the WC 600 is integrated with the IGU 602, the WC 600 itself can include a docking module that the battery can plug into. In implementations in which the WC 600 is integrated with the IGU 602, the IGU 602 can still include a plug-in component 752 that connects with the WC 600. In some other integrated implementations, the WC 600 can be directly connected to the busbars of the associated ECD. In such latter integrated implementations, the communication module storing the device parameters of the ECD can be located within the WC 600, for example, in a non-volatile memory within the WC 600. More examples of the use of integrated window controllers and energy storage devices are described in U.S. Patent Application No. 14/951,410 (Attorney Docket No. VIEWP008X1US) filed November 24, 2015 and titled SELF-CONTAINED EC IGU, and PCT Patent Application No. PCT/US16/41176 (Attorney Docket No. VIEWP080WO) filed July 6, 2016 and titled POWER MANAGEMENT FOR ELECTROCHROMIC WINDOW NETWORKS, both of which are hereby incorporated by reference in their entireties and for all purposes.

Processing unit 604

[0151] At a high level, the processing unit 604 functions to communicate with the upstream network controller and to control the tint states of the IGUs 602 connected with the WC 600. One primary function of the processing unit 604 is to generate a command signal V_{DCmnd}. As will be described in more detail below, the command signal V_{DCmnd} is provided to the drive circuit 608 for generating the applied voltage signals V_{App1} and V_{App2}, which are output from the WC 600 for driving one or more IGUs 602 controlled by the WC 600. In various implementations the processing unit 604 can generate the command signal V_{DCmnd} based on a number of different device parameters, drive parameters, input values, algorithms or instructions. For example, the processing unit 604 can generate the command signal V_{DCmnd} based on a tint

command received from the upstream network controller. As described above, the tint command can include a tint value corresponding to a target tint state for the IGUs 602 controlled by the WC 600.

[0152] In some implementations, responsive to receiving a tint command, the processing unit 604 initiates a tinting transition in one or more of the IGUs 602 controlled by the WC 600. In some implementations, the processing unit 604 calculates, selects, determines or otherwise generates the command signal V_{DCmnd} based on drive parameters including the current tint state of an IGU 602 to be transitioned and the target tint state of the IGU 602 (based on the tint value in the tint command). The processing unit 604 also can generate the command signal V_{DCmnd} based on other drive parameters, for example, a ramp-to-drive rate, a drive voltage, a drive voltage duration, a ramp-to-hold rate and a holding voltage for each possible combination of current tint state and target tint state. Other drive parameters can include parameters based on current or recent sensor data, for example, an indoor temperature, an outdoor temperature, a temperature within the interior volume of the IGU 602 (or of one or more of the panes), a light intensity in a room adjacent the IGU 602 and a light intensity outside of the IGU 602, among other suitable or desirable parameters. In some implementations, such sensor data can be provided to the WC 600 via the upstream network controller over communication lines 626 and 628. Additionally or alternatively, the sensor data can be received from sensors located within or on various portions of the IGU 602. In some such implementations, the sensors can be within or otherwise coupled with a communication module within the IGU 602 (such as the communication module 756). For example, multiple sensors including photosensors, temperature sensors or transmissivity sensors can be coupled via the same communication lines 739 and 741 shown in Figure 7 according to the 1-Wire communication protocol.

[0153] In some implementations, the processing unit 604 also can generate the command signal V_{DCmnd} based on the device parameters associated with the ECD within the IGU 602. As described above, the device parameters for the ECD can include a length, width, thickness, cross-sectional area, shape, age, model number, version number, or number of previous optical transitions of or associated with the respective ECD (or of a pane on which the ECD is formed or otherwise arranged). In some implementations, the processing unit 604 is configured to track the number of tinting transitions for each of the connected IGUs 602.

[0154] In some implementations, the processing unit 604 generates the command signal V_{DCmnd} based on a voltage control profile, for example, such as that described above with reference to Figure 2. For example, the processing unit 604 can use the drive parameters and device parameters to select a voltage control profile from a predefined set of voltage control profiles stored in a memory within or accessible by the processing unit 604. In some

implementations, each set of voltage control profiles is defined for a particular set of device parameters. In some implementations, each voltage control profile in a given set of voltage control profiles is defined for a particular combination of drive parameters. The processing unit 604 generates the command signal V_{DCmnd} such that the drive circuit 608 implements the selected voltage control profile. For example, the processing unit 604 adjusts the command signal V_{DCmnd} to cause the drive circuit 608 to, in turn, adjust the applied voltage signals V_{App1} and V_{App2}. More specifically, the drive circuit 608 adjusts the applied voltage signals V_{App1} and V_{App2} such that the effective voltage V_{Eff} applied across the ECD tracks the voltage levels indicated by the voltage control profile throughout the progression through the profile.

[0155] In some implementations, the processing unit 604 also can modify the command signal V_{DCmnd} dynamically (whether during a transition or during a holding period after a transition) based on sensor data. As described above, such sensor data can be received from various sensors within or otherwise integrated with the connected IGUs 602 or from other external sensors. In some such implementations, the processing unit 604 can include intelligence (for example, in the form of programming instructions including rules or algorithms), that enable the processing unit 604 to determine how to modify the command signal V_{DCmnd} based on the sensor data. In some other implementations, the sensor data received by the WC 600 from such sensors can be communicated to the network controller, and in some instances from the network controller to the master controller. In such implementations, the network controller or the master controller can revise the tint value for the IGUs 602 based on the sensor data and transmit a revised tint command to the WC 600. Additionally or alternatively, the network controller or the master controller can receive sensor data from one or more other sensors external to the building, for example, one or more light sensors positioned on a roof top or a facade of the building. In some such implementations, the master controller or the network controller can generate or revise the tint value based on such sensor data.

[0156] In some implementations, the processing unit 604 also can generate or modify the drive signal V_{Drive} dynamically based on one or more feedback signals V_{Feed} received from the feedback circuit 610. For example, and as will be described in more detail below, the feedback circuit 610 can provide one or more voltage feedback signals V_{OC} based on actual voltage levels detected across the ECDs (for example, as measured during periodic open circuit instances), one or more current feedback signals V_{Cur} based on actual current levels detected through the ECDs, or based on one or more voltage compensation signals V_{Comp} associated with voltage drops detected or determined along the power transmission lines that provide the applied voltage signals V_{App1} and V_{App2} to the IGUs 602.

[0157] Generally, the processing unit 604 can be implemented with any suitable processor or logic device, including combinations of such devices, capable of performing the functions or processes described herein. In some implementations, the processing unit 604 is a microcontroller (also referred to as a microcontroller unit (MCU)). In some more specific applications, the processing unit 604 can be a microcontroller particularly designed for embedded applications. In some implementations, the processing unit 604 includes a processor core (for example, a 200 MHz processor core or other suitable processor core) as well as a program memory (for example, a 2018 KB or other suitable non-volatile memory), a random-access memory (RAM) (for example, a 512 KB or other suitable RAM), and various I/O interfaces. The program memory can include, for example, code executable by the processor core to implement the functions, operations or processes of the processing unit 604.

[0158] In some implementations, the RAM can store status information for the IGUs 602 controlled by the WC 600. The RAM also can store the device parameters for the ECDs within the IGUs 602. In some other implementations, the processing unit 604 can store such status information or device parameters in another memory device (for example, a Flash memory device) external to the processing unit 604 but also within the WC 600. In some specific implementations, the I/O interfaces of the processing unit 604 include one or more CAN interfaces, one or more synchronous serial interfaces (for example, 4-wire Serial Peripheral Interface (SPI) interfaces), and one or more Inter-Integrated Circuit (I^2C) interfaces. One example of such a controller suitable for use in some implementations is the PIC32MZ2048ECH064 controller provided by Microchip Technology Inc. of Chandler, AZ.

[0159] In the implementation illustrated in Figure 6, the WC 600 additionally includes a data bus transceiver 664. The data bus transceiver 664 is coupled with the upstream interface 614 via the communication bus 632. The data bus transceiver 664 also is coupled with the processing unit 604 via a communication bus 666. As described above, in some implementations, the communication bus 632 is designed, deployed and otherwise configured in accordance with the CAN bus standard, which is a differential bus standard. In some implementations, the communication bus 666 also conforms to the CAN bus standard and includes a differential pair of lines for transferring a differential pair of signals. As such, the data bus transceiver 664 can include two sets of differential ports; a first set for coupling with the communication bus 632 and a second set for coupling with the communication bus 666, which in turn is coupled with a CAN interface of the processing unit 604.

[0160] In various implementations, the data bus transceiver 664 is configured to receive data from a network controller (such as the NC 500) via the communication bus 632, process the data, and transmit the processed data to the processing unit 604 via the communication bus 666.

Similarly, the data bus transceiver 664 is configured to receive data from the processing unit 604 via the communication bus 666, process the data, and transmit the processed data over the communication bus 632 to the interface 614 and ultimately over the upstream set of cables 616 to the network controller. In some such implementations, processing the data includes converting or translating the data from a first protocol to a second protocol (for example, from a CAN protocol (such as CANopen) to a protocol readable by the processing unit 604 and vice versa). One example of such a data bus transceiver suitable for use in some implementations is the SN65HVD1050 data bus transceiver provided by Texas Instruments Inc. of Dallas, TX. In some other implementations, the processing unit 604 can include an integrated data bus transceiver or otherwise include functionalities of the data bus transceiver 664 rendering the inclusion of the external data bus transceiver 664 unnecessary.

Power Circuit

[0161] At a high level, the power circuit 606 is operable to receive power from the power supply lines 622 and 624 and to provide power to various components of the WC 600 including the processing unit 604, the drive circuit 608, the feedback circuit 610 and the communications circuit 612. As described above, the first power supply line 622 receives a supply voltage V_{Sup1}, for example, a DC voltage having a value in the range of approximately 5 V to 42 V (relative to the supply voltage V_{Sup2}), and in one example application, a value of 24 V (although higher voltages may be desirable and are possible in other implementations). As is also described above, the second power supply line 624 can be a power supply return. For example, the voltage V_{Sup2} on the second power supply line 624 can be a reference voltage, for example, a floating ground.

[0162] The power circuit 606 includes at least one down converter (also referred to herein as a "buck converter") for stepping down the supply voltage V_{Sup1}. In the illustrated implementation, the power circuit 606 includes two down converters: a first relatively low power (LP) down converter 668 and a second relatively high power (HP) down converter 670. The LP down converter 668 functions to step down the supply voltage V_{Sup1} to a first down-converted voltage V_{Dwn1}. In some implementations, the down-converted voltage V_{Dwn1} can have a value in the range of approximately 0 to 5 V, and in one example application, a value of approximately 3.3 V. The down-converted voltage V_{Dwn1} is provided to the processing unit 604 for powering the processing unit 604. One example of an LP down converter suitable for use in some implementations is the TPS54240 2.5 Ampere (Amp) DC-DC step-down converter provided by Texas Instruments Inc. of Dallas, TX.

[0163] The HP down converter 670 functions to step down the supply voltage V_{Sup1} to a second down-converted voltage V_{Dwn2}. One example of an HP down converter suitable for use in

some implementations is the TPS54561 5 Amp DC-DC step-down converter provided by Texas Instruments Inc. of Dallas, TX. In some implementations, the down-converted voltage V_{Dwn2} can have a value in the range of approximately 6V to 24V, and in one example application, a value of approximately 6 V. The down-converted voltage V_{Dwn2} is provided to the voltage regulator 680, described below with reference to the drive circuit 608. In some implementations, the down-converted voltage V_{Dwn2} also is provided to the rest of the components within the WC 600 that require power to perform their respective functions (although these connections are not shown in order to avoid over complicating the illustration and to avoid obscuring the other components and connections).

[0164] In some implementations, the HP down converter 670 provides the down-converted voltage V_{Dwn2} only when enabled (or instructed) to do so, for example, when or while the processing unit 604 asserts an enable signal *En*. In some implementations, the enable signal *En* is provided to the HP down converter 670 via a Serial Peripheral Interface (SPI) interface bus 686. Although the SPI interface bus 686 may be described herein in the singular form, the SPI bus 686 may collectively refer to two or more SPI buses, each of which can be used to communicate with a respective component of the WC 600. In some implementations, the processing unit asserts the enable signal *En* only when the WC 600 is in an "active mode," as opposed to a "sleep mode."

[0165] In some implementations, the power circuit 606 further includes or is coupled with an energy storage device (or "energy well") 672 such as, for example, a capacitive storage device such as a rechargeable battery (or set of batteries) or a supercapacitor. For example, one example of a supercapacitor suitable for use in some implementations can have a capacitance C_S of at least 400 Farads at 0.4 watt hours (Wh). In some implementations, the energy storage device 672 can be charged by a charger 674. In some such implementations, the charger 674 can be powered by the supply voltage V_{Sup1}. One example of such a charger suitable for use in some implementations is the LT3741 constant-current, constant-voltage, step-down controller provided by Linear Technology Corp. of Milpitas, CA. In some implementations, the charger 674 also is configured to provide power stored in the energy storage device 672 to the power supply line 622.

[0166] In some implementations, the charger 674 can alternatively or additionally be powered by one or more photovoltaic (or "solar") cells. For example, such photovoltaic (PV) cells can be integrated onto or into the IGUs 602, such as on one or more panes of the IGUs, controlled by the WC 600. In some such implementations, the power received via the PV cell can be regulated by a voltage regulator 676 prior to being provided to the charger 674 and ultimately the energy storage device 672. For example, the voltage regulator 676 can serve to step up or step down the

voltage of the power received from the PV cells. The voltage regulator 676 also can generally be used to regulate the power provided by the PV cells as such power fluctuates throughout a day, for example, to maintain the voltage of the power at a fixed level. In some implementations, when the power stored in the energy storage device 672 is desired or needed, it gets released via the charger 674. In some implementations, to prevent back drive (that is, to ensure that power from the energy storage device 672 or the PV cells does not flow upstream over the upstream set of cables 616), the power circuit 606 can additionally include an asymmetric conductor 678, for example, a low loss semiconductor diode such as a Schottky junction diode or a p-n junction diode. The use of such a diode 678 can be especially advantageous in implementations in which one or more of the supply voltages V_{Sup1} and V_{Sup2} are pulsed. More examples of the use of integrated PV cells are described in U.S. Patent Application No. 14/951,410 (Attorney Docket No. VIEWP008X1) filed November 24, 2015 and titled SELF-CONTAINED EC IGU, which is hereby incorporated by reference in its entirety and for all purposes.

[0167] The integration of energy storage devices can be advantageous for a number of reasons, whether such devices are included within respective WCs 600 (like the energy storage device 672) or are otherwise distributed throughout a network system (such as the network system 300). For example, the power circuit 606 within each WC 600 can supplement or augment the power provided by the respective power supply lines 622 and 624 with power drawn from the energy storage device 672. Additionally or alternatively, energy storage devices external to the WCs 600 can provide power directly to the power distribution lines that distribute power throughout the network system to supply the WCs 600. Such implementations can be especially advantageous in high demand instances in which many IGUs 602 are to be transitioned concurrently. In times of lower demand, the normal power supply (for example, the power supply provided by a building source) can recharge the energy storage devices. More examples of the use of energy storage devices are described in U.S. Patent Application No. 14/951,410 (Attorney Docket No. VIEWP008X1) filed November 24, 2015 and titled SELF-CONTAINED EC IGU, and PCT Patent Application No. PCT/US16/41176 (Attorney Docket No. VIEWP080WO) filed July 6, 2016 and titled POWER MANAGEMENT FOR ELECTROCHROMIC WINDOW NETWORKS, both of which are hereby incorporated by reference in their entireties and for all purposes.

[0168] Additionally or alternatively, in some implementations, the transitions of the IGUs 602 can be staggered. For example, the MC 400 or the NC 500 can issue tint commands for subsets of the WCs 600 at different times so as to keep the total power consumed by the network system (or a portion of the network system) at any given time under a desirable, safe, permitted

or maximum limit. In some other implementations, the WCs 304 can be programmed via various parameters received from the MC 400 or NC 500 to delay their transitions. For example, the secondary tint command issued by the NC 500 also can include a delay value that informs the WC 400 to begin a tint change after the time associated with the delay value has lapsed. As another example, the secondary tint command issued by the NC 500 also can include a time value that informs the WC 400 to begin a tint change when a time associated with the time value has been reached. In these latter two examples, the NC 500 can issue tint commands to the WCs 304 approximately simultaneously or contemporaneously while ensuring that staggering of the transitions is still achieved.

Drive Circuit

[0169] At a high level, the drive circuit 608 is generally operable to receive the command signal V_{DCmnd} from the processing unit 604 and to provide the applied voltage signals V_{App1} and V_{App2} for driving the connected IGUs 602 based on the command signal V_{DCmnd}. The drive circuit 608 includes a voltage regulator 680 that receives the down-converted voltage V_{Dwn2} from the HP down converter 670 in the power circuit 606. The voltage regulator 680 regulates, adjusts or otherwise transforms the voltage V_{Dwn2} to provide (or "generate") first and second regulated voltage signals V_{P1} and V_{P2} based on the command signal V_{DCmnd}. In some implementations, the voltage regulator 680 is a buck-boost converter; that is, the voltage regulator 680 can be capable of functioning as a down converter to step down the voltage V_{Dwn2} as well as an up converter to step up the input voltage V_{Dwn2}. Whether the voltage regulator 680 behaves as a down converter or as an up converter is dependent on the command signal V_{DCmnd}, as is the magnitude of the down conversion or up conversion, respectively. In some more specific implementations, the voltage regulator 680 is a synchronous buck-boost DC-DC converter. In some such implementations, the regulated voltage signals V_{P1} and V_{P2} are effectively fixed-amplitude DC signals from the perspective of the IGUs 602, and in particular, the ECDs within the IGUs 602.

[0170] As described in more detail above, the processing unit 604 can generate the command signal V_{DCmnd} based on a number of different parameters, input values, algorithms or instructions. In some implementations, the processing unit 604 generates the command signal V_{DCmnd} in the form of a digital voltage signal. In some such implementations, the drive circuit 608 can additionally include a digital-to-analog converter (DAC) 682 for converting the digital command signal V_{DCmnd} to an analog command voltage signal V_{ACmnd}. In some implementations, the DAC 682 can be external to the processing unit 604, while in some other implementations, the DAC 682 is internal to the processing unit 604. In such implementations, the voltage regulator 680 more specifically generates the regulated voltage signals V_{P1} and V_{P2} based on the command

voltage signal V_{ACmnd}. One example of a DAC suitable for use in some implementations is the AD5683R DAC by Analog Devices Inc. of Norwood, MA.

[0171] In some specific implementations, the regulated voltage signals V_{P1} and V_{P2} are rectangular wave (or "pulsed") DC signals, for example, pulse-width modulated (PWM) voltage signals. In some such implementations, the voltage regulator 680 includes an H-bridge circuit to generate the regulated voltage signals V_{P1} and V_{P2}. In some such implementations, each of the regulated voltage signals V_{P1} and V_{P2} has the same frequency. In other words, the period from the start of a current pulse to the start of the next pulse in each of the regulated voltage signals V_{P1} and V_{P2} has the same time duration. In some implementations, the voltage regulator 680 is operable to modify the duty cycles of the respective voltage signals V_{P1} and V_{P2} such that the respective duty cycles are not equal. In this way, while the amplitude (or "magnitude") of the pulses (or "on" durations) of the first regulated voltage signal V_{P1} can be equal to the magnitude of the pulses of the second regulated voltage signal V_{P2}, each of the first and the second regulated voltage signals V_{P1} and V_{P2} can have a different effective DC voltage magnitude from the perspective of the corresponding busbars and conducting layers of the ECDs in the IGUs 602. However, in some other implementations, the voltage regulator 680 can additionally or alternatively modify the respective magnitudes of the pulses of the voltage signals V_{P1} and V_{P2}.

[0172] For example, consider an application in which each of the pulses of each of the regulated voltage signals V_{P1} and V_{P2} has an amplitude of 5 V, but in which the first voltage signal V_{P1} has a 60% duty cycle while the second voltage signal V_{P2} has a 40% duty cycle. In such an application, the effective DC voltage provided by each of the regulated voltage signals V_{P1} and V_{P2} can be approximated as the product of the respective pulse amplitude and the fraction of the duty cycle occupied the respective pulses. For example, the effective DC voltage provided by the first voltage signal V_{P1} can be approximated as 3 V (the product of 5 V and 0.6) while the effective voltage provided by the second voltage signal V_{P2} can be approximated as 2 V (the product of 5 V and 0.4). In some implementations, the duty cycle of first voltage signal V_{P1} is complementary to the duty cycle of the second voltage signal V_{P2}. For example, as in the case of the example just provided, if the first voltage signal V_{P1} has a duty cycle of X%, the duty cycle of the second voltage signal V_{P2} can be Y%, where Y% = 100% - X%. In some such implementations, the "on" durations of the first voltage signal V_{P1} can coincide with the "off" durations of the second voltage signal V_{P2}, and similarly, the "off" durations of the first voltage signal V_{P1} can coincide with the "on" durations of the second voltage signal V_{P2}. In some other implementations, the duty cycles do not necessarily have to be complementary; for example, the first voltage signal V_{P1} can have a duty cycle of 50% while the second voltage signal V_{P2} can have a duty cycle of 15%.

[0173] As described above, in some implementations, the regulated voltage signals V_{P1} and V_{P2} are effectively fixed-amplitude DC signals from the perspective of the IGUs 602, and in particular, the ECDs within the IGUs 602. To further such implementations, the voltage regulator 680 also can include one or more electronic filters, and in particular, one or more passive filter components such as one or more inductors. Such filters or filter components can smooth out the regulated voltage signals V_{P1} and V_{P2} prior to their provision to ensure that the regulated voltage signals V_{P1} and V_{P2} are effectively fixed-amplitude DC signals. To further facilitate the smoothing of the regulated voltage signals V_{P1} and V_{P2}, the frequency of the pulses in the voltage signals V_{P1} and V_{P2} can be greater than or equal to 1 kilohertz (kHz) in some implementations. For example, as one of ordinary skill in the art will appreciate, the greater the frequency of the voltage oscillations applied to a conductor, the less able the electric charge in the conductor is able to react to the voltage oscillations. Additionally, the greater the inductance of an inductor, the more smoothing out of the voltage oscillations that are provided through the inductor.

[0174] In some implementations, the voltage regulator 680 can advantageously be capable of operating in a burst mode to reduce the power consumption of the WC 600 over time. In the burst mode of operation, the voltage regulator 680 automatically enters and exits the burst mode to minimize the power consumption of the voltage regulator 680. One example of such a voltage regulator suitable for use in some implementations is the LTC3112 15 V, 2.5 Amp Synchronous Buck-Boost DC/DC Converter provided by Linear Technology Corp. of Milpitas, CA.

[0175] In some implementations, the regulated voltage signals V_{P1} and V_{P2} are the applied voltage signals V_{App1} and V_{App2}, respectively. In some such implementations, the difference between the regulated voltage signals V_{P1} and V_{P2} is the effective voltage V_{Eff}. In some implementations, to effect a lightening tinting transition, the processing unit 604 generates the command signal V_{DCmnd} such that the voltage regulator 680 provides a positive effective voltage V_{Eff}, while to effect a darkening tinting transition, the processing unit 604 generates the command signal V_{DCmnd} such that the voltage regulator 680 provides a negative effective voltage V_{Eff}. Conversely, in some other implementations involving different electrochromic layers or counter electrode layers, a darkening tinting transition is achieved by providing a positive effective voltage V_{Eff} while a lightening tinting transition is achieved by providing a negative effective voltage V_{Eff}.

[0176] Either way, the voltage regulator 680 can provide a positive effective voltage V_{Eff} by increasing the duty cycle of the first voltage signal V_{P1} or by decreasing the duty cycle of the second voltage signal V_{P2} such that the duty cycle of the first voltage signal V_{P1} is greater than the duty cycle of the second voltage signal V_{P2}, and consequently, the effective DC voltage of

the first applied voltage signal V_{App1} is greater than the effective DC voltage of the second applied voltage signal V_{App2}. Similarly, the voltage regulator 680 can provide a negative effective voltage V_{Eff} by decreasing the duty cycle of the first voltage signal V_{P1} or by increasing the duty cycle of the second voltage signal V_{P2} such that the duty cycle of the first voltage signal V_{P1} is less than the duty cycle of the second voltage signal V_{P2}, and consequently, the effective DC voltage of the first applied voltage signal V_{App1} is less than the effective DC voltage of the second applied voltage signal V_{App2}.

[0177] In some other implementations, including that illustrated in Figure 6, the drive circuit 608 additionally includes a polarity switch 684. The polarity switch 684 receives the two regulated voltage signals V_{P1} and V_{P2} from the voltage regulator 680 and outputs the applied voltage signals V_{App1} and V_{App2} that are provided to the power lines 634 and 636, respectively. The polarity switch 482 can be used to switch the polarity of the effective voltage V_{Eff} from positive to negative, and vice versa. Again, in some implementations, the voltage regulator 680 can increase the magnitude of V_{P1} relative to V_{P2}, and thus increase the magnitude of V_{Eff}, by increasing the duty cycle of the first voltage signal V_{P1} or by decreasing the duty cycle of the second voltage signal V_{P2}. Similarly, the voltage regulator 680 can decrease the magnitude of V_{P1} relative to V_{P2}, and thus decrease the magnitude of V_{Eff}, by decreasing the duty cycle of the first voltage signal V_{P1} or by increasing the duty cycle of the second voltage signal V_{P2}.

[0178] In some other implementations, the second voltage V_{P2} can be a signal ground. In such implementations, the second voltage V_{P2} can remain fixed or floating during transitions as well as during times between transitions. In such implementations, the voltage regulator 680 can increase or decrease the magnitude of V_{P1}, and thus the magnitude of V_{Eff}, by increasing or decreasing the duty cycle of the first voltage signal V_{P1}. In some other such implementations, the voltage regulator 680 can increase or decrease the magnitude of V_{P1}, and thus the magnitude of V_{Eff}, by directly increasing or decreasing the amplitude of the first voltage signal V_{P1} with or without also adjusting the duty cycle of the first voltage signal V_{P1}. Indeed, in such latter implementations, the first voltage signal V_{P1} can be an actual fixed DC signal rather than a pulsed signal.

[0179] In implementations that include a polarity switch 684, the second voltage signal V_{P2} can be a signal ground and the first voltage signal V_{P1} can always be a positive voltage relative to the second voltage signal V_{P2}. In such implementations, the polarity switch 684 can include two configurations (for example, two electrical configurations or two mechanical configurations). The processing unit 604 can control which of the configurations the polarity switch 684 is in via a control signal V_{Polar} provided, for example, over the SPI bus 686. For example, the processing unit 604 can select the first configuration when implementing a lightening transition and the

second configuration when implementing a darkening transition. For example, while the polarity switch 684 is in the first configuration, the polarity switch can output a positive first applied voltage signal V_{App1} relative to the second applied voltage signal V_{App2}. Conversely, while the polarity switch 684 is in the second configuration, the polarity switch can output a negative first applied voltage signal V_{App1} relative to the second applied voltage signal V_{App2}.

[0180] In some implementations, while in the first configuration, the polarity switch 684 passes the first voltage signal V_{P1} (or a buffered version thereof) as the first applied voltage signal V_{App1} and passes the second voltage signal V_{P2} (or a grounded version thereof) as the second applied voltage signal V_{App2}, resulting in a positive effective voltage V_{Eff}. In some implementations, while in the second configuration, the polarity switch 684 passes the first voltage signal V_{P1} (or a buffered version thereof) as the second applied voltage signal V_{App2} and passes the second voltage signal V_{P2} (or a grounded version thereof) as the first applied voltage signal V_{App2}, resulting in a negative effective voltage V_{Eff}. In some implementations, the polarity switch 684 can include an H-bridge circuit. Depending on the value of V_{Polar}, the H-bridge circuit can function in the first configuration or the second configuration. One example of a polarity switch suitable for use in some implementations is the IRF7301 HEXFET Power MOSFET provided by International Rectifier Corp. of San Jose, CA.

[0181] In some implementations, when switching from a positive voltage V_{Eff} to a negative voltage V_{Eff}, or vice versa, the polarity switch 684 can be configured to switch from a first conducting mode, to a high impedance mode and then to a second conducting mode, or vice versa. For didactic purposes, consider an example in which the first regulated voltage V_{P1} is at a positive hold value and in which the polarity switch 684 is in the first configuration. As described above, in some implementations the polarity switch 684 passes V_{P1} (or a buffered version thereof) as the first applied voltage V_{App1} resulting in a first applied voltage V_{App1} that also is at the positive hold value. To simplify the illustration, also assume that V_{P2} and V_{App2} are both signal grounds. The result would be an effective applied voltage V_{Eff} at the positive hold value. Now consider that the processing unit 604 is initiating a tinting transition that will result in an end state in which the effective applied voltage V_{Eff} is at a negative hold value. In some implementations, to effect the tinting transition, the processing unit 604 adjusts the command signal V_{DCmnd} to cause the voltage regulator 680 to lower the magnitude of the voltage V_{P1} based on a negative ramp-to-drive profile. In some implementations, as the magnitude of the voltage V_{P1} reaches a threshold value close to zero (for example, 10 millivolts (mV)), the processing unit 604 changes the polarity switching signal V_{Polar} from a first value to a second value to cause the polarity switch 684 to switch from a positive conducting mode (the first configuration described above) to a high impedance mode.

[0182] While in the high impedance mode the polarity switch 684 does not pass V_{P1}. Instead, the polarity switch 684 can output values of V_{App1} (or V_{App2}) based on predefined calculations or estimations. Meanwhile, the voltage regulator 680 continues to decrease the magnitude of V_{P1} to zero. When the magnitude of V_{P1} reaches zero, the voltage regulator 680 begins increasing the magnitude of V_{P1} up to the magnitude of the negative drive value. When the magnitude of V_{P1} reaches a threshold value (for example, 10 mV), the processing unit 604 then changes the polarity switching signal V_{Polar} from the second value to a third value to cause the polarity switch 684 to switch from the high impedance mode to a negative conducting mode (the second configuration described above). As described above, in some such implementations, the polarity switch 684 passes V_{P1} as the second applied voltage V_{App2}, while the first applied voltage V_{App1} is a signal ground. To summarize, while the magnitude of V_{P1} is greater than or equal to a threshold voltage (for example, 10 mV) the polarity switch 684 passes the regulated voltage V_{P1} as either the first applied voltage V_{App1} or the second applied voltage V_{App2}, depending on whether the polarity switch 684 is in the positive conducting mode (first configuration) or the negative conducting mode (second configuration), respectively. As such, the effective applied voltage V_{Eff} is dictated by the magnitude of V_{P1} and the polarity configuration of the polarity switch 684 while the value of V_{Eff} is less than or equal to -10mV or greater than or equal to +10mV. But while the polarity switch 684 is in the high impedance mode, in the range when $-10mV<V_{Eff}<10mV$, the value of V_{Eff}, and more generally the values of V_{App1} and V_{App2}, are determined based on predefined calculations or estimations.

Feedback Circuit

[0183] As described above, in some implementations the processing unit 604 can modify the command signal V_{DCmnd} during operation (for example, during a tinting transition or during times between tinting transitions) based on one or more feedback signals V_{Feed}. In some implementations, a feedback signal V_{Feed} is based on one or more voltage feedback signals V_{OC}, which are in turn based on actual voltage levels detected across the ECDs of the connected IGUs. Such voltage feedback signals V_{OC} can be measured during periodic open circuit conditions (during or in between transitions) while the applied voltages V_{App1} and V_{App2} are turned off for brief durations of time. For example, an open-circuit voltage feedback signal V_{OC} can be measured using a differential amplifier 688 having a first input connected with power line 634, a second input connected with power line 636, and an output connected with an analog-to-digital converter (ADC) 692. The ADC 692 can be internal or external with respect to the processing unit 604. One example of a differential amplifier suitable for use in some implementations is the low power, adjustable gain, precision LT1991 provided by Linear Technology Corp. of Milpitas, CA.

[0184] Additionally or alternatively, a second feedback signal V_{Feed} can be based on one or more current feedback signals V_{Cur}, which are in turn based on actual current levels detected through the ECDs. Such current feedback signals V_{Cur} can be measured using an operational amplifier 690 having a first input connected with a first input terminal of a resistor 691, which is also connected to an output of the polarity switch 684. A second input of the operational amplifier 690 can be connected with a second terminal of the resistor 691, which is also connected to a node at the second supply voltage V_{Sup2}. The output of the operational amplifier 690 can be connected with the ADC 692. One example of an operational amplifier suitable for use in some implementations is the low noise, CMOS, precision AD8605 provided by Analog Devices Inc. of Norwood, MA. Because the resistance R_F of the resistor 691 is known, the actual current flowing out of the polarity switch 684 can be determined by processing unit 604 based on the voltage difference signal V_{Cur}.

[0185] In some implementations, the processing unit 604 also is configured to compensate for transmission losses resulting from the passage of the voltage signals V_{App1} and V_{App2} through the conducting power distribution lines 633 and 635. More specifically, the actual voltages provided to the busbars of a given IGU 602 can be less than the voltages V_{App1} and V_{App2} at the output of the WC 600. As such, the actual voltage V_{Act} applied across the ECD within the IGU 402 can be less than the difference between the voltages V_{App1} and V_{App2} at the output of the WC 600. For example, the resistances of the power distribution lines 634 and 636—diagrammatically represented as resistors each having resistance R_T—can result in significant voltage drops along the power distribution lines 634 and 636. The resistance of each power distribution line is, of course, directly proportional to the length of the power distribution line and inversely proportional to the cross-sectional area of the power distribution line. An expected voltage drop can thus be calculated based on knowledge of the length of the power distribution lines. However, this length information is not necessarily available. For example, installers may not record such length information during installation of the IGUs or may not record such information accurately, precisely or correctly. Additionally, in some legacy installations where existing wires are utilized, such length information may not be available.

[0186] If information about the lengths of the power distribution lines is available, this information can be used to create a lookup table, for example, that is stored in the memory chip within the plug-in component. This length information can then be read by the WC 600 upon power-up of the WC 600. In such implementations, the voltages V_{App1} and V_{App2} can be increased (for example, using firmware or software) to compensate for the estimated voltage drops along the respective power distribution lines 634 and 636. While such compensation schemes and algorithms can be to some extent effective, such schemes and algorithms cannot

precisely account for the dynamic changes in the resistances of the power distribution lines resulting from changes in the temperatures of the power distribution lines, which can change greatly in a given day based on use of the power distribution lines, based on the position of the sun as the Earth spins, based on the weather, and based on the season.

[0187] Additionally or alternatively, a third feedback signal V_{Feed} can be based on one or more voltage compensation signals V_{Comp}, which are in turn based on an actual voltage drop detected along at least one of the power distribution lines. For example, such voltage compensation signals V_{Comp} can be measured using a differential amplifier 694 having a first input connected with a one of the power distribution lines 634 or 634 in the WC 600, a second input connected with the fifth line 642 in the WC 600, and an output connected with the ADC 692. In some such implementations, such as that shown and described with reference to Figure 7, the plug-in component 752 includes a voltage compensation circuit 762. In one example implementation, the voltage compensation circuit 762 includes a conductor that provides a short between the fifth line 742 and the first or the second power distribution line 734 or 736, respectively, within the plug-in component 752. In such an implementation, the differential amplifier 694 detects the offset voltage V_{Comp}, which is proportional to the current I through the power distribution line between the WC 600 and the IGU 602, as well as the length of, and the cross-sectional area of, the power distribution line between the WC 600 and the IGU 602. The current I is determined by the processing unit 604 based on the signal V_{Cur} output from operational amplifier 690. In this way, the processing unit can increase or decrease the command voltage signal V_{DCmnd} to compensate for the static and dynamic voltage drops along the power distribution lines without having direct knowledge of the length or the cross-sectional area of the power distribution lines.

[0188] In one implementation, the resistance, R_T, of each power distribution line between the WC 600 and the IGU 602 is calculated by dividing V_{Comp} by I. This resistance information is then stored in a parameter table within the WC 600. V_{Comp} is then dynamically calculated as $2*R_T*V_{Cur}$. The voltage signals V_{App1} and V_{App2} can subsequently dynamically adjusted automatically using the calculated V_{Comp} amount to compensate for voltage drop in the lines 633 and 635. In another scenario, the voltage signals V_{App1} and V_{App2} are adjusted dynamically by $2*V_{Comp}$ to account for voltage drop in lines 633 and 635.

[0189] Voltage compensation also is described in more detail in U.S. Patent Application No. 13/449,248 (Attorney Docket No. VIEWP041) filed April 17, 2012 and titled CONTROLLER FOR OPTICALLY SWITCHABLE WINDOWS, and U.S. Patent Application No. 13/449,251 (Attorney Docket No. VIEWP042) filed April 17, 2012 and titled CONTROLLER FOR OPTICALLY SWITCHABLE WINDOWS, both of which are hereby incorporated by reference

in their entireties and for all purposes. In some other implementations, a voltage compensation circuit 762 can be connected to communication lines 739 and 741, which connect to the chip 756. In some other implementations, the voltage compensation circuit 762 can be directly coupled with the communication lines 637 and 639 via the interface 754 and the communication lines 738 and 740.

[0190] Each of the open-circuit voltage feedback signal V_{OC}, the current feedback signal V_{Cur} and the voltage compensation feedback signal V_{Comp} can be digitized by the ADC 692 and provided to the processing unit 604 as a feedback signal V_{Feed}. One example of an ADC suitable for use in some implementations is the low power AD7902 by Analog Devices Inc. of Norwood, MA. In some instances above, while the feedback signal V_{Feed} is referenced in the singular form, the feedback signal V_{Feed} can collectively refer to three (or more or less) individual feedback signals: a first one for the digitized open-circuit voltage signal V_{OC}, a second one for the digitized current signal V_{Cur} and a third one for the digitized voltage compensation signal V_{Comp}. The feedback signal V_{Feed} can be provided to the processing unit 604 via the SPI bus 686. The processing unit 604 can then use the feedback signal V_{Feed} to dynamically modify the command signal VD_{Cmnd} such that the actual value V_{Act} of the voltage applied across the ECD stack of the IGU 602 is approximately equal to the desired effective voltage V_{Eff}, and thus, such that the target tint state is reached.

[0191] For example, as the outside environment becomes brighter, the WC 600 can receive a tint command from the NC 500 to darken an IGU 602. However, in some implementations or instances, as the respective ECD becomes increasingly more tinted, the temperature of the ECD can rise significantly as a result of the increased photon absorption. Because the tinting of the ECD can be dependent on the temperature of the ECD, the tint state can change if the command signal V_{DCmnd} is not adjusted to compensate for the temperature change. In some implementations, rather than detecting the temperature fluctuation directly, the processing unit 604 can adjust the command signal V_{DCmnd} based on the actual voltage detected across the ECD or the actual current detected through the ECD, as determined via the feedback signals V_{OC} and V_{Cur}.

[0192] Additionally, as described above, each WC 600 can be connected to and power a plurality of IGUs 602. While the cross-sectional areas of the set of power distribution lines that connect a given WC 600 to each respective one of the plurality of connected IGUs 602 are generally the same, the lengths of each set of power distribution lines can be different based on the location of the respective IGU 602 relative to the WC 600. Thus, while the WC 600 provides the voltages V_{App1} and V_{App2} to the plurality of connected IGUs 602 via a common node (such as through the coupling connecter 748 described above with reference to Figure 7), the

values of the voltages V_{App1} and V_{App2} actually received by each of the plurality of IGUs 602 can be different based on the locations of the respective ones of the IGUs 402 relative to the WC 600. In some implementations, it can be desirable that the power distribution lines connecting each of the IGUs 602 to a given WC 600 have the same or similar length to reduce the disparities between the actual applied voltages received by the IGUs 602.

Communications Circuit

[0193] The communications circuit 612 is generally configured to enable communication between the processing unit 604 and various other components within or outside of the WC 600. For example, the communications circuit 612 can include a bridge device 696. In some implementations, the bridge device 696 enables the processing unit 696 to communicate and receive data signals $Data_3$ and $Data_4$ over communication lines 638 and 640 (collectively referred to as data bus 644), and corresponding communication lines 637 and 639. In some implementations, the bridge device 696 can be a 1-Wire bridge device configured to communicate according to the 1-Wire communications protocol. In some such implementations, the communication lines 639 and 640 can be signal grounds, while the communication lines 637 and 639, which carry the data signal $Data_3$, can provide both data and power to the chip 756 as well as to any number of 1-Wire-compatible sensors within the IGU 602. In some implementations, the chip 756 within the IGU 602 can be an intermediary for communications of data between the processing unit 604 and the sensors within the IGU 602. For example, the sensors can be connected to communication lines 739 and 741, which connect to the chip 756. In some other implementations, the sensors can be directly coupled with the communication lines 637 and 639 via the interface 754 and the communication lines 738 and 740. At other times, the data signal $Data_3$ can communicate sensor data back to the processing unit 604.

[0194] The bridge device 696 is configured to manage the communications to, from and among the 1-Wire devices. The processing unit 604 can communicate instructions to the bridge device 696, or receive data from the bridge device, via an I^2C bus 697. Although the I^2C bus 697 may be described herein in the singular form, the I^2C bus 697 may collectively refer to two or more I^2C buses, each of which can be used to communicate with a respective component of the WC 600. Thus, in some implementations, the bridge device 696 functions as an I^2C to 1-Wire bridge that interfaces directly to an I^2C host port of the I^2C master (the processing unit 604) to perform bidirectional protocol conversion between the processing unit 604 and the downstream 1-Wire slave devices including the chip 756 and any sensors on or within the IGU 602. One such bridge device suitable for use in some implementations is the DS2482 1-Wire Master device provided by Maxim Integrated Products, Inc. of San Jose, CA. In some other

implementations, the functions of the bridge device 696 can be integrated into the processing unit 604.

[0195] In some implementations, responsive to powering on or otherwise activating the processing unit 604, the processing unit 604 instructs, via the bridge device 696, the communication module 756 within the plug-in component 752 to transfer the device and drive parameters to the RAM or other memory device within the processing unit 604. Additionally or alternatively, the processing unit 604 can periodically poll for the communication module 756 via the bridge device 696. The communication module 756 can then respond to the poll by transferring the drive parameters to the RAM or other memory device within the WC 600 via the bridge device 696.

[0196] In some implementations, the communications circuit 612 also includes a radio transceiver 698. For example, the radio transceiver 698 can communicate with the processing unit 604 via the I^2C bus 697. The radio transceiver 698 can enable wireless communication between the processing unit 604 and other devices having such radio transceivers including, for example, other WCs 600, the NC 500, the IGUs 602 as well as mobile devices or other computing devices. While referred to herein in the singular form, the radio transceiver 698 can collectively refer to one or more radio transceivers each configured for wireless communication according to a different respective protocol. For example, some wireless network protocols suitable for use in some implementations can be based on the IEEE 802.11 standard, such as Wi-Fi (or "WiFi"). Additionally or alternatively, the radio transceiver 698 can be configured to communicate based on the IEEE 802.15.4 standard, which defines the physical layer and media access control for low-rate wireless personal area networks (LR-WPANs). For example, higher level protocols compatible with the IEEE 802.15.4 standard can be based on the ZigBee, 6LoWPAN, ISA100.11a, WirelessHART or MiWi specifications and standards. Additionally or alternatively, the radio transceiver 698 can be configured to communicate based on the Bluetooth standard (including the Classic Bluetooth, Bluetooth high speed and Bluetooth low energy protocols and including the Bluetooth v4.0, v4.1 and v4.2 versions). Additionally or alternatively, the radio transceiver 698 can be configured to communicate based on the EnOcean standard (ISO/IEC 14543-3-10).

[0197] As described above, wireless communication can take the place of communication over physical cables between the WC 600 and the NC 500. In some other implementations, both wired and wireless communications can be established between the WC 600 and the NC 500. In other words, at least two communication links of different types can be simultaneously maintained to send data between the WC and the MC. For instance, the WC can be in wired communication with the NC using CANbus for some less data intensive messaging such as WC

voltage data, current data and sensor data. At the same time, the WC can be in wireless communication with the NC via WiFi or other any wireless communication technique disclosed herein for more data intensive communications such as a video camera feed and/or an audio feed. When two or more communication links are maintained, one communication link can serve as a backup for the other in case of a disruption or other error condition. In some implementations, sensors and other devices can be in communication with the WC using a wireless link, a wired link or both. In some implementations, the distributed WCs 600 can form a mesh network for communicating various information to one another or to the MC 400, the NC 500 or to other devices, rendering physical communication lines between the various controllers of a network system such as network system 300 unnecessary. As also noted above, the WC 600 can communicate wirelessly with the IGUs 602 it controls. For example, the communication module 756 within each IGU 602 also can include a radio transceiver for communicating with the radio transceiver 698 and the processing unit 604 of the WC 600. In some implementations, wireless communication can take the place of communication over physical cables between the WC 600 and the IGU 602. For example, wireless communication can take the place of the 1-Wire communication bus 644, the communication lines 637 and 639, and the communication lines 738 and 740. Such wireless implementations can facilitate the manufacture and installation of self-contained IGUs, for example, IGUs that don't require the attachment of physical cables. In some such self-contained implementations, each IGU can include an energy storage device and an integrated photovoltaic cell for charging the energy storage device. The energy storage device, in turn, can power the tint states and tint state transitions of the ECD within the IGU.

[0198] In some implementations, the communications circuit 612 can additionally or alternatively include a power line communications module 699. The power line communications module 699 can be used in implementations or instances in which data is communicated via the power supply voltage signal V_{Sup1} (and in some cases, also V_{Sup2}) rather than, or in addition to, over communications lines 622 and 624 or wirelessly. As shown, the power line communications module 699 also can communicate with the processing unit 604 via the I^2C bus 697.

Auto-/Semiauto-Commissioning/Self-discovery

[0199] In some implementations, after installation and after the WCs have been turned on, the WCs can request or poll for the 1-Wire IDs within the IGUs 602. These 1-Wire IDs are then sent from the WC to the NC, and ultimately to the MC so that the MC can associate the CANbus ID of the WC to the 1-Wire IDs of the IGUs it controls. In some other implementations, the IGUs also can include wireless transceivers. For example, a Bluetooth transceiver within each

IGU can broadcast a beacon containing the ID of the IGU, which the WC can then pick up. Once the IDs of the IGUs connected with the WC are known, a person can then proceed through the building with a mobile device (phone, IPad, or proprietary device) to associate each of the IGUs with a physical location.

Sleep Modes

[0200] In some implementations, the WC 600 is configured to enter and exit one or more sleep modes in addition to the normal (or "active") operating mode. For example, after a target tint state has been reached and a holding voltage has been applied for a duration of time, the processing unit 604 can stop asserting (or "deassert") the enable signal *EN*, and thus disable the HP downconverter 670. Because the HP down converter 670 supplies power to most of the components within the WC 600, when the enable signal *EN* is deasserted, the WC 600 enters a first sleep mode. Alternatively, instead of turning off or disabling the HP down converter 670, the processing unit can disable each of the components within the WC 600 individually or selectively in groups by deasserting other enable signals (not shown) to such individual components or groups. In some implementations, prior to disabling the HP down converter 670 or otherwise disabling the desired components within the WC 600, the processing unit 604 asserts a control signal *Cntrl* that causes the voltage regulator 680 to enter a high impedance mode, for example, so that when the other components are turned off, charge stored within the EC stacks of the connected IGUs 602 doesn't flow backwards from the IGUs into the WC 600. In some implementations, the LP down converter 668 remains on during the first sleep mode to provide full power to the processing unit 604. In some implementations, the processing unit 604 can enable the differential amplifier 688 and the ADC 692 periodically to determine whether V_{OC} has fallen (or risen) below a threshold level, for example, to determine whether the tint state of the IGU has changed beyond an acceptable level. When V_{OC} has fallen below (or risen above) the threshold, the processing unit 604 can "awaken" the WC 600 (for example, exit the sleep mode and return to the normal active operating mode) by turning on the HP down converter 670 or otherwise turning on the components necessary to drive the EC stack of the IGU to an acceptable level. In some implementations, upon exiting the sleep mode, the processing unit 604 can cause a voltage ramp to be applied to the EC stack followed by a holding voltage.

[0201] In some implementations, the processing unit 604 can be configured to cause the WC 600 to enter a second (or "deep") sleep mode different than the first (or "light") sleep mode. For example, after the WC 600 has been in the first sleep mode for a duration of time, the processing unit 604 can disable some of its functionality to further save power. In effect, the processing unit 604 itself enters a sleep mode. The processing unit 604 still gets the 3.3V from the LP down converter, but it configured in a reduced-functionality, low-power mode in which it consumes

significantly less power than in the normal fully functional mode. While in such a second sleep mode, the processing unit 604 can be awakened in one or more of a number of ways. For example, the processing unit 604 can awaken itself periodically (such as every minute, every few minutes, every 10 minutes). As described above, the processing unit 604 can then enable the differential amplifier 688 and the ADC 692 to determine whether V_{OC} has fallen below (or risen above) a threshold level, for example, to determine whether the tint state of the IGU has changed beyond an acceptable level. When V_{OC} has fallen below (or risen above) the threshold, the processing unit 604 can awaken the WC 600 by turning on the HP down converter 670 or otherwise turning on the components necessary to drive the EC stack of the IGU to an acceptable level. In some implementations, upon exiting the sleep mode, the processing unit 604 can cause a voltage ramp to be applied to the EC stack followed by a holding voltage.

[0202] Additionally or alternatively, the processing unit 604 can be awakened from such a deep sleep mode based on an interrupt such as a command from NC 500 or based on a signal from an occupancy sensor communicatively coupled with the processing unit 604. When such an occupancy sensor detects an occupant, the occupancy sensor can provide a signal to the processing unit 604 that causes the processing unit to awaken and return the WC 600 to the active mode (in some other implementations, the occupancy sensor can be coupled with the NC 500 which then sends an awaken command to the WC 500 based on a signal from the occupancy sensor). In some implementations, for example in scenarios in which users carry devices that include Bluetooth or other suitable types of transceivers that periodically poll or send beacons for pairing, the processing unit 604 can periodically awaken to enable the radio transceiver 698 to determine whether any such devices are in proximity.

[0203] Additionally, to further save power during such sleep modes, the processing unit 604 can enable the voltage regulator 680 via the control signal *Cntrl* to draw the power needed to power the processing unit 604 and the radio transceiver 698 from the charge stored within the EC stack of the IGU 602. More examples of the use of power conservation and intelligent and efficient power distribution are described in PCT Patent Application No. PCT/US16/41176 (Attorney Docket No. VIEWP080WO) filed July 6, 2016 and titled POWER MANAGEMENT FOR ELECTROCHROMIC WINDOW NETWORKS, which is hereby incorporated by reference in its entirety and for all purposes. Additionally, subject matter related to obtaining V_{OC} is further described in U.S. Patent Application No. 13/931,459 (Attorney Docket No. VIEWP052) filed June 28, 2013 and titled CONTROLLING TRANSITIONS IN OPTICALLY SWITCHABLE DEVICES, which is hereby incorporated by reference in its entirety and for all purposes.

Smart Network Controller

[0204] In some implementations, the NC 500 described with reference to Figure 5 can take over some of the functions, processes or operations that are described above as being responsibilities of the MC 400 of Figure 4. Additionally or alternatively, the NC 500 can include additional functionalities or capabilities not described with reference to the MC 400. Figure 8 shows a block diagram of example modules of a network controller in accordance with some implementations. For example, the modules of Figure 8 can be implemented in the NC 500 in any suitable combination of hardware, firmware and software. In some implementations in which the NC 500 is implemented as a network controller application executing within a computer, each of the modules of Figure 9 also can be implemented as an application, task or subtask executing within the network controller application.

[0205] In some implementations, the NC 500 periodically requests status information from the WCs 600 it controls. For example, the NC 500 can communicate a status request to each of the WCs 600 it controls every few seconds, every few tens of seconds, every minute, every few minutes or after any desirable period of time. In some implementations, each status request is directed to a respective one of the WCs 600 using the CAN ID or other identifier of the respective WC 600. In some implementations, the NC 500 proceeds sequentially through all of the WCs 600 it controls during each round of status acquisition. In other words, the NC 500 loops through all of the WCs 600 it controls such that a status request is sent to each of the WCs 600 sequentially in each round of status acquisition. After a status request has been sent to a given WC 600, the NC 500 then waits to receive the status information from the respective WC 600 before sending a status request to the next one of the WCs in the round of status acquisition.

[0206] In some implementations, after status information has been received from all of the WCs 600 that the NC 500 controls, the NC 500 then performs a round of tint command distribution. For example, in some implementations, each round of status acquisition is followed by a round of tint command distribution, which is then followed by a next round of status acquisition and a next round of tint command distribution, and so on. In some implementations, during each round of tint command distribution, the NC 500 proceeds to send a tint command to each of the WCs 600 that the NC 500 controls. In some such implementations, the NC 500 also proceeds sequentially through all of the WCs 600 it controls during the round of tint command distribution. In other words, the NC 500 loops through all of the WCs 600 it controls such that a tint command is sent to each of the WCs 600 sequentially in each round of tint command distribution.

[0207] In some implementations, each status request includes instructions indicating what status information is being requested from the respective WC 600. In some implementations,

responsive to the receipt of such a request, the respective WC 600 responds by transmitting the requested status information to the NC 500 (for example, via the communication lines in the upstream set of cables 616). In some other implementations, each status request by default causes the WC 600 to transmit a predefined set of information for the set of IGUs 602 it controls. Either way, the status information that the WC 600 communicates to the NC 500 responsive to each status request can include a tint status value (S) for the IGUs 602, for example, indicating whether the IGUs 602 is undergoing a tinting transition or has finished a tinting transition. Additionally or alternatively, the tint status value S or another value can indicate a particular stage in a tinting transition (for example, a particular stage of a voltage control profile). In some implementations, the status value S or another value also can indicate whether the WC 600 is in a sleep mode. The status information communicated in response to the status request also can include the tint value (C) for the IGUs 602, for example, as set by the MC 400 or the NC 500. The response also can include a set point voltage set by the WC 600 based on the tint value (for example, the value of the effective applied V_{Eff}). In some implementations, the response also can include a near real-time actual voltage level V_{Act} measured, detected or otherwise determined across the ECDs within the IGUs 602 (for example, via the amplifier 688 and the feedback circuit 610). In some implementations, the response also can include a near real-time actual current level I_{Act} measured, detected or otherwise determined through the ECDs within the IGUs 602 (for example, via the amplifier 690 and the feedback circuit 610). The response also can include various near real-time sensor data, for example, collected from photosensors or temperature sensors integrated on or within the IGUs 602.

[0208] Some protocols such as CANOpen limit the size of each frame of data sent from the WC 600 to the NC 500 and vice versa. In some instances, the sending of each status request and the receiving of status information responsive to such a request actually includes multiple two-way communications, and thus, multiple frames. For example, each status request described above can include a separate sub-request for each of the status values described above. As a more specific example, each status request from the NC 500 to a particular WC 600 can include a first sub-request requesting the status value S. In response to the first sub-request, the WC 600 can transmit to the NC 500 an acknowledgement and a frame including the status value S. The NC 500 can then transmit a second sub-request to the WC 600 requesting the tint value C. In response to the second sub-request, the WC 600 can transmit to the NC 500 an acknowledgement and a frame including the tint value C. The values of V_{Eff}, V_{Act} and I_{Act} as well as sensor data can similarly be obtained with separate respective sub-requests and responses.

[0209] In some other implementations, rather than polling or sending a status request to each of the WCs 600 on a sequential basis, the NC 500 can asynchronously send status requests to

particular WCs 600. For example, it may not be useful to receive status information (including C, S, V_{Eff}, V_{Act} and I_{Act}) from all of the WCs 600 periodically. For example, it may be desirable to asynchronously request such information from only particular ones of the WCs 600 that have recently received or implemented a tint command, that are currently undergoing a tinting transition, that have recently finished a tinting transition, or from which status information has not been collected for a relatively long duration of time.

[0210] In some other implementations, rather than polling or sending status requests to each of the WCs 600 individually, whether on a sequential basis or asynchronously, each of the WCs 600 can periodically broadcast its status information (including C, S, V_{Eff}, V_{Act} and I_{Act}). In some such implementations, each of the WCs 600 can broadcast the status information wirelessly. For example, each WC 600 can broadcast the status information every few seconds, tens of seconds, minutes or tens of minutes. In some implementations, the WCs 600 can be synchronized to broadcast their respective status information at certain times to avoid occupying a large amount of collective bandwidth. Additionally, the broadcast period can be different for different sets (such as the zones described above) of WCs 600 and at different times, for example, based on the positions of the respective IGUs in the building and relative to the sun, or based on whether the rooms adjoining the IGUs are occupied.

[0211] In some other implementations, each of the WCs 600 can broadcast its status information in response to certain conditions, for example, when starting a tinting transition, when finishing a tinting transition, when V_{Act} changes by a threshold, when I_{Act} changes by a threshold, when sensor data (for example, light intensity or temperature) changes by a threshold, when an occupancy sensor indicates the adjoining room is occupied, or when entering or exiting a sleep mode. The NC 500 can listen for such broadcasted status information, and when it hears it, record the status information. Advantageously, in broadcasting implementations, the time required to receive status information from a set of WCs 600 is approximately cut in half because there is no need to request the status information from the WCs 600, and thus, no roundtrip delay associated with each WC 600. Instead, there is only a one-way latency associated with the time required to transmit the status information from each WC 600 to the NC 500.

[0212] In some other implementations, at power on or thereafter, each of the WCs 600 can be configured to read device parameters, drive parameters and lite IDs or other ECD IDs for connected IGUs. The WCs then broadcast their CAN IDs as well as the lite IDs and the associated device and drive parameters. That is, in some implementations, such broadcasting is initiated by one or more processors in a WC without or irrespective of any requests for such data by the NCs or other controllers. When the IDs and parameters are broadcast, the NC 500 can

receive and process the IDs and parameters. In some implementations, lite IDs and parameters from messages broadcasted by the WC are then communicated from the NC to the MC, which stores them, for example, in a table including a list of known CAN IDs. For example, each row of the table can include a CAN ID, a WC location ID associated with the CAN ID, the connected lite IDs, the locations of the respective windows associated with the lite IDs, and the device and drive parameters for the respective ECDs. In some implementations, the MC can store the table in a cloud-based database system so that even if the MC fails, another MC can be instantiated and access the table in the cloud.

[0213] In some instances, during commissioning, a field service technician may intervene and attempt to perform ad hoc lite-to-lite matching based on perceived differences in the tints of two or more neighboring windows. In such cases, the technician may determine that the drive parameters for one or more ECDs should be modified, and these modifications are then implemented. In some implementations, the WC is configured to broadcast the modified parameters to the corresponding NC, from which the parameters can be communicated to the MC. In situations where the WC then fails or experiences an error, the NC or MC can determine that the WC has failed, for instance, because the WC is no longer broadcasting in situations where the WC has been configured to periodically broadcast data such as the WC's CAN ID and/or WC location ID. When the failed WC is replaced with a new WC, which is then powered-on, the new WC will read the corresponding lite IDs and, as described above, broadcast the new WC's CAN ID and the connected lite IDs. When the NC or MC receives this information, the NC or MC can be configured to retrieve the modified drive parameters for the failed WC from a database table by performing a table look-up using the lite IDs. In such instances, the NC or MC is also configured to automatically update the table by assigning the new CAN ID to the WC location ID and associated lite IDs. The NC or MC will then automatically communicate the modified drive parameters to the new WC. In this way, the ECD which had its drive parameters modified during commissioning can still be driven by the modified drive parameters even when the respective WC has been replaced. Other techniques for automatically modifying, updating, and applying drive parameters can be performed in some implementations, as further described in U.S. Provisional Patent Application No. 62/305,892, titled METHOD OF COMMISSIONING ELECTROCHROMIC WINDOWS, by Shrivastava et al., filed March 9, 2016 (Attorney Docket No. VIEWP008X2P), which is hereby incorporated by reference in its entirety and for all purposes,

[0214] In some such implementations, rather than sending a tint command to each of the WCs 600 on a sequential basis, the NC 500 can asynchronously send a tint command to a

particular WC 600 whether through a wired or wireless connection. For example, it may not be useful to send tint commands to all of the WCs 600 periodically. For example, it may be desirable to asynchronously sent tint commands to only particular ones of the WCs 600 that are to be transitioned to a different tint state, for which status information has just been (or has recently been) received, or to which a tint command has not been sent for a relatively long duration of time.

Data Logger

[0215] In some implementations, the NC 500 also includes a data logging module (or "data logger") 802 for recording data associated with the IGUs controlled by the NC 500. In some implementations, the data logger 802 records the status information included in each of some or all of the responses to the status requests. As described above, the status information that the WC 600 communicates to the NC 500 responsive to each status request can include a tint status value (S) for the IGUs 602, a value indicating a particular stage in a tinting transition (for example, a particular stage of a voltage control profile), a value indicating whether the WC 600 is in a sleep mode, a tint value (C), a set point voltage set by the WC 600 based on the tint value (for example, the value of the effective applied V_{Eff}), an actual voltage level V_{Act} measured, detected or otherwise determined across the ECDs within the IGUs 602, an actual current level I_{Act} measured, detected or otherwise determined through the ECDs within the IGUs 602, and various sensor data, for example, collected from photosensors or temperature sensors integrated on or within the IGUs 602. In some other implementations, the NC 500 can collect and queue status information in a messaging queue like RabbitMC, ActiveMQ or Kafka and stream the status information to the MC for subsequent processing such as data reduction/compression, event detection, etc., as further described herein.

[0216] In some implementations, the data logger 802 within the NC 500 collects and stores the various information received from the WCs 600 in the form of a log file such as a comma-separated values (CSV) file or via another table-structured file format. For example, each row of the CSV file can be associated with a respective status request, and can include the values of C, S, V_{Eff}, V_{Act} and I_{Act} as well as sensor data (or other data) received in response to the status request. In some implementations, each row is identified by a timestamp corresponding to the respective status request (for example, when the status request was sent by the NC 500, when the data was collected by the WC 600, when the response including the data was transmitted by the WC 600, or when the response was received by the NC 500). In some implementations, each row also includes the CAN ID or other ID associated with the respective WC 600.

[0217] In some other implementations, each row of the CSV file can include the requested data for all of the WCs 600 controlled by the NC 500. As described above, the NC 500 can

sequentially loop through all of the WCs 600 it controls during each round of status requests. In some such implementations, each row of the CSV file is still identified by a timestamp (for example, in a first column), but the timestamp can be associated with a start of each round of status requests, rather than each individual request. In one specific example, columns 2–6 can respectively include the values C, S, V_{Eff}, V_{Act} and I_{Act} for a first one of the WCs 600 controlled by the NC 500, columns 7–11 can respectively include the values C, S, V_{Eff}, V_{Act} and I_{Act} for a second one of the WCs 600, columns 12–16 can respectively include the values C, S, V_{Eff}, V_{Act} and I_{Act} for a third one of the WCs 600, and so on and so forth through all of the WCs 600 controlled by the NC 500. The subsequent row in the CSV file can include the respective values for the next round of status requests. In some implementations, each row also can include sensor data obtained from photosensors, temperature sensors or other sensors integrated with the respective IGUs controlled by each WC 600. For example, such sensor data values can be entered into respective columns between the values of C, S, V_{Eff}, V_{Act} and I_{Act} for a first one of the WCs 600 but before the values of C, S, V_{Eff}, V_{Act} and I_{Act} for the next one of the WCs 600 in the row. Additionally or alternatively, each row can include sensor data values from one or more external sensors, for example, positioned on one or more facades or on a rooftop of the building. In some such implementations, the NC 500 can send a status request to the external sensors at the end of each round of status requests.

Compact Status

[0218] As described above, some protocols such as CANopen limit the size of each frame sent from the WC 600 to the NC 500 and vice versa. In some instances, the sending of each status request and the receiving of status information responsive to such a request actually includes multiple two-way communications and frames. For example, each status request described above can include a separate sub-request for each of the status values described above. In some implementations, each of two or more of the requested values C, S, V_{Eff}, V_{Act} and I_{Act} can be transmitted together within a single response—a compact status response. For example, in some implementations, the values of two or more of C, S, V_{Eff}, V_{Act} and I_{Act} are formatted so as to fit in one frame. For example, the CANopen protocol limits the size of the data payload that can be sent in each frame to 8 bytes (where each byte includes 8 bits). And in implementations in which the Service Data Object (SDO) sub-protocol of CAN open is used, the maximum size of the data payload portion of the CANopen frame is 4 bytes (32 bits). In some implementations, the size of each of the values V_{Eff}, V_{Act} and I_{Act} is 10 bits. Thus, each of the values of V_{Eff}, V_{Act} and I_{Act} can be packaged within a single SDO frame. This leaves 2 bits left over. In some implementations, each of the values of C and S can be specified with one respective bit. In such

case, all of the values of C, S, V_{Eff}, V_{Act} and I_{Act} can be specified using only 32 bits, and thus, be packaged within one SDO CANopen frame.

[0219] In some implementations, additional time savings can be achieved using a broadcast status request. For example, rather than sending a status request to each of the WCs 600 on an individual (or "unicast" basis), the NC 500 can broadcast a single status request to all of the WCs 600 it controls. As described above, responsive to receiving the status request, each WC 600 can be programmed to respond by communicating status information such as the values C, S, V_{Eff}, V_{Act} and I_{Act} in one or more compact status responses.

Protocol Conversion Module

[0220] As described above, one function of the NC 500 can be in translating between various upstream and downstream protocols, for example, to enable the distribution of information between WCs 600 and the MC 400 or between the WCs and the outward-facing network 310. In some implementations, a protocol conversion module 804 is responsible for such translation or conversion services. In various implementations, the protocol conversion module 904 can be programmed to perform translation between any of a number of upstream protocols and any of a number of downstream protocols. As described above, such upstream protocols can include UDP protocols such as BACnet, TCP protocols such as oBix, other protocols built over these protocols as well as various wireless protocols. Downstream protocols can include, for example, CANopen, other CAN-compatible protocol, and various wireless protocols including, for example, protocols based on the IEEE 802.11 standard (for example, WiFi), protocols based on the IEEE 802.15.4 standard (for example, ZigBee, 6LoWPAN, ISA100.11a, WirelessHART or MiWi), protocols based on the Bluetooth standard (including the Classic Bluetooth, Bluetooth high speed and Bluetooth low energy protocols and including the Bluetooth v4.0, v4.1 and v4.2 versions), or protocols based on the EnOcean standard (ISO/IEC 14543-3-10).

Integrated Analytics

[0221] In some implementations, the NC 500 uploads the information logged by the data logger 802 (for example, as a CSV file) to the MC 400 on a periodic basis, for example, every 24 hours. For example, the NC 500 can transmit a CSV file to the MC 400 via the File Transfer Protocol (FTP) or another suitable protocol over an Ethernet data link 316. In some such implementations, the status information can then be stored in the database 320 or made accessible to applications over the outward-facing network 310.

[0222] In some implementations, the NC 500 also can include functionality to analyze the information logged by the data logger 802. For example, an analytics module 906 can receive and analyze the raw information logged by the data logger 802 in real time. In various implementations, the analytics module 806 can be programmed to make decisions based on the

raw information from the data logger 802. In some other implementations, the analytics module 806 can communicate with the database 320 to analyze the status information logged by the data logger 802 after it is stored in the database 320. For example, the analytics module 806 can compare raw values of electrical characteristics such as V_{Eff}, V_{Act} and I_{Act} with expected values or expected ranges of values and flag special conditions based on the comparison. For example, such flagged conditions can include power spikes indicating a failure such as a short, an error, or damage to an ECD. In some implementations, the analytics module 806 communicates such data to the tint determination module 810 or to the power management module 812.

[0223] In some implementations, the analytics module 806 also can filter the raw data received from the data logger 802 to more intelligently or efficiently store information in the database 320. For example, the analytics module 806 can be programmed to pass only "interesting" information to a database manager 808 for storage in the database 320. For example, interesting information can include anomalous values, values that otherwise deviate from expected values (such as based on empirical or historical values), or for specific periods when transitions are happening. More detailed examples of how raw data can be filtered, parsed, temporarily stored, and efficiently stored long term in a database are described in PCT Patent Application No. PCT/2015/029675 (Attorney Docket No. VIEWP049X1WO) filed May 7, 2015 and titled CONTROL METHOD FOR TINTABLE WINDOWS, which is hereby incorporated by reference in its entirety and for all purposes.

Database Manager

[0224] In some implementations, the NC 500 includes a database manager module (or "database manager") 808 configured to store information logged by the data logger 804 to a database on a periodic basis, for example, every hour, every few hours or every 24 hours. In some implementations, the database can be an external database such as the database 320 described above. In some other implementations, the database can be internal to the NC 500. For example, the database can be implemented as a time-series database such as a Graphite database within the secondary memory 506 of the NC 500 or within another long term memory within the NC 500. In some example implementations, the database manager 808 can be implemented as a Graphite Daemon executing as a background process, task, sub-task or application within a multi-tasking operating system of the NC 500. A time-series database can be advantageous over a relational database such as SQL because a time-series database is more efficient for data analyzed over time

[0225] In some implementations, the database 320 can collectively refer to two or more databases, each of which can store some or all of the information obtained by some or all of the NCs 500 in the network system 300. For example, it can be desirable to store copies of the

information in multiple databases for redundancy purposes. In some implementations, the database 320 can collectively refer to a multitude of databases, each of which is internal to a respective NC 500 (such as a Graphite or other times-series database). It also can be desirable to store copies of the information in multiple databases such that requests for information from applications including third party applications can be distributed among the databases and handled more efficiently. In some such implementations, the databases can be periodically or otherwise synchronized to maintain consistency.

[0226] In some implementations, the database manager 808 also can filter data received from the analytics module 806 to more intelligently or efficiently store information in an internal or external database. For example, the database manager 808 can additionally or alternatively be programmed to store only "interesting" information to a database. Again, interesting information can include anomalous values, values that otherwise deviate from expected values (such as based on empirical or historical values), or for specific periods when transitions are happening. More detailed examples of how raw data can be filtered, parsed, temporarily stored, and efficiently stored long term in a database are described in PCT Patent Application No. PCT/2015/029675 (Attorney Docket No. VIEWP049X1WO) filed May 7, 2015 and titled CONTROL METHOD FOR TINTABLE WINDOWS, which is hereby incorporated by reference in its entirety and for all purposes.

Tint Determination

[0227] In some implementations, the NC 500 or the MC 400 includes intelligence for calculating, determining, selecting or otherwise generating tint values for the IGUs 602. For example, as similarly described above with reference to the MC 400 of Figure 4, a tint determination module 810 can execute various algorithms, tasks or subtasks to generate tint values based on a combination of parameters. The combination of parameters can include, for example, the status information collected and stored by the data logger 802. The combination of parameters also can include time or calendar information such as the time of day, day of year or time of season. Additionally or alternatively, the combination of parameters can include solar calendar information such as, for example, the direction of the sun relative to the IGUs 602. The combination of parameters also can include the outside temperature (external to the building), the inside temperature (within a room adjoining the target IGUs 602), or the temperature within the interior volume of the IGUs 602. The combination of parameters also can include information about the weather (for example, whether it is clear, sunny, overcast, cloudy, raining or snowing). Parameters such as the time of day, day of year, or direction of the sun can be programmed into and tracked by the NC 500. Parameters such as the outside temperature, inside temperature or IGU temperature can be obtained from sensors in, on or around the building or

sensors integrated on or within the IGUs 602. In some implementations, various parameters can be provided by, or determined based on information provided by, various applications including third party applications that can communicate with the NC 500 via an API. For example, the network controller application, or the operating system in which it runs, can be programmed to provide the API.

[0228] In some implementations, the tint determination module 810 also can determine tint values based on user overrides received via various mobile device applications, wall devices or other devices. In some implementations, the tint determination module 810 also can determine tint values based on commands or instructions received various applications, including third party applications and cloud-based applications. For example, such third party applications can include various monitoring services including thermostat services, alert services (for example, fire detection), security services or other appliance automation services. Additional examples of monitoring services and systems can be found in PCT/US2015/019031 (Attorney Docket No. VIEWP061WO) filed 5 March 2015 and titled MONITORING SITES CONTAINING SWITCHABLE OPTICAL DEVICES AND CONTROLLERS. Such applications can communicate with the tint determination module 810 and other modules within the NC 500 via one or more APIs. Some examples of APIs that the NC 500 can enable are described in U.S. Provisional Patent Application Serial No. 62/088,943 (Attorney Docket No. VIEWP073P) filed 8 December 2014 and titled MULTIPLE INTERFACING SYSTEMS AT A SITE.

Power Management

[0229] As described above, the analytics module 806 can compare values of V_{Eff}, V_{Act} and I_{Act} as well as sensor data either obtained in real time or previously stored within the database 320 with expected values or expected ranges of values and flag special conditions based on the comparison. The analytics module 806 can pass such flagged data, flagged conditions or related information to the power management 812. For example, such flagged conditions can include power spikes indicating a short, an error, or damage to an ECD. The power management module 812 can then modify operations based on the flagged data or conditions. For example, the power management module 812 can delay tint commands until power demand has dropped, stop commands to troubled WCs (and put them in idle state), start staggering commands to WCs, manage peak power, or signal for help.

Controlling Multiple Optically Switchable Devices

[0230] Applying the same tint profile to multiple IGUs under control of a single window controller may cause the IGUs to have different tinting characteristics, for example, due to manufacturing differences or other variations between the IGUs. For example, the IGUs may tint at different rates, end a tinting transition at different tint states, or the like. Different tinting

characteristics may cause windows that are proximate and/or adjacent to each other to visually appear different. Applying different tint profiles to each IGU separately may waste resources. For example, separately driving each IGU under control of the window controller may waste the power available from a power supply associated with the window controller. Accordingly, disclosed herein are techniques, methods, systems, and apparatuses to control multiple IGUs using a single window controller such that the multiple IGUs are driven together for a portion of a tint transition, and then fine-tuned individually to achieve a tint uniformity. For example, fine-tuning may occur after a bulk of the tint transition has occurred.

[0231] In some implementations, a single window controller may control multiple optically switchable devices (e.g., IGUs). For example, in some implementations, a window controller may control two IGUs, four IGUs, six IGUs, etc. In some embodiments, each IGUs may be similar in one or more characteristics, such as size (e.g., surface area), dimensions, bus bar dimensions, or the like. In some embodiments, the multiple IGUs may be disposed in and/or located in the same zone of a facility (e.g., in the same region of a building, facing the same geographical direction, located on the same floor of a building, or the like). In some embodiments, a window controller that controls multiple IGUs may be located proximate to the multiple IGUs, such as in a ceiling area near the multiple IGUs. By controlling multiple IGUs with a single window controller, the multiple IGUs may be able to be transitioned in tint in a similar manner (e.g., such that the IGUs transition relatively concurrently). Additionally, use of a single window controller may provide cost savings, because the cost of the window controller may be spread over the multiple IGUs that the window controller controls. For example, multiple IGUs may share the cost of a single power supply. Moreover, use of a single window controller to control multiple IGUs may reduce communications network and/or power distribution network infrastructure, because a relatively smaller number of window controllers may be operatively coupled to a trunk line, while controlling a larger number of IGUs than window controllers coupled to the trunk line.

[0232] In some implementations, a window controller may apply a drive voltage to multiple (e.g., two or more) IGUs. The drive voltage may be applied in connection with a tint transition that is to be applied to each of the IGUs (e.g., a common tint transition from a first tint state to a second tint state). The drive voltage may be provided via a power supply (e.g., a common DC power supply). In some implementations, the window controller may monitor, for each of the IGUs, parameters indicative of a state of the tint transition. In some implementations, in response to determining that at least one of the multiple IGUs has completed the tint transition, the window controller may transition the at least one IGU to a hold voltage while concurrently continuing to apply the drive voltage to the remaining IGUs. The window controller may then

cycle through the remaining IGUs, and may transition each of the remaining IGUs to the hold voltage upon determining that the IGU has completed the tint transition. By initially driving all of the IGUs with a common drive voltage (e.g., until at least one IGU has completed the tint transition), the bulk of ionic charge transfer may be completed synchronously for all of the IGUs. IGUs which are slower to complete the tint transition may then be topped up until the tint transition is completed through continued provision of the drive voltage. This may allow the IGUs to be transitioned to the same tint state in a more synchronous and uniform manner.

[0233] In some implementations, parameters indicative of a state of a tint transition include an open circuit voltage (V*oc*) of an IGU and/or an amount of charge transferred to an IGU. For example, a determination that an IGU has completed a tint transition may be made in response to determining that the V*oc* of the IGU meets or exceeds a target V*oc* and/or in response to determining that an amount of charge delivered to the IGU meets or exceeds a target charge to be delivered to the IGU. In some embodiments, each IGU is operatively coupled to a common power supply via a switch. In some embodiments, the switch may be an H-bridge circuit. In some embodiments, a V*oc* measurement associated with an IGU may be made by causing the corresponding switch that operatively couples the IGU to the common power supply to be in a high impedance mode.

[0234] In some implementations, an amount of charge transferred to or delivered to a IGU is used to determine whether an IGU has completed a tint transition. In some embodiments, the charge delivered is considered per unit of area, which may be referred to as the delivered charge or charge density, or total delivered charge or charge density. In some embodiments, delivered charge or charge density may be determined by determining an amount of current delivered to an IGU over a period of time (e.g., by integrating the current delivered over a period of time).

[0235] In some embodiments, to determine if a tint transition has been completed, the total delivered charge or charge density may be compared to a threshold charge or threshold charge density (also referred to as a target charge or charge density). The threshold charge or threshold charge density may be chosen based on the minimum charge or charge density required to fully complete or nearly complete the optical transition under the likely operating conditions. In various cases, the threshold charge or threshold charge density may be chosen/estimated based on the charge or charge density required to fully complete or nearly complete the optical transition at a defined temperature (e.g., at about -40°C, at about -30°C, at about -20°C, at about -10°C, at about 0°C, at about 10°C, at about 20°C, at about 25°C, at about 30°C, at about 40°C, at about 60°C, etc.). A suitable threshold charge or threshold charge density may also be affected by the leakage current of the electrochromic device. Devices that have higher leakage currents should have higher threshold charge densities. In some embodiments, an appropriate

threshold charge or threshold charge density may be determined empirically for an individual window or window design. In other cases, an appropriate threshold may be calculated/selected based on the characteristics of the window such as the size, bus bar separation distance, leakage current, starting and ending optical states, etc. Example threshold charge densities range between about $1x10^{-5}$ C/cm^2 and about 5 C/cm^2, for example between about $1x10^{-4}$ and about 0.5 C/cm^2, or between about 0.005-0.05 C/cm^2, or between about 0.01-0.04 C/cm^2, or between about 0.01-0.02 in many cases. Smaller threshold charge densities may be used for partial transitions (e.g., fully clear to 25% tinted) and larger threshold charge densities may be used for full transitions. A first threshold charge or charge density may be used for bleaching/clearing transitions, and a second threshold charge or charge density may be used for coloring/tinting transitions. In certain embodiments, the threshold charge or charge density is higher for tinting transitions than for clearing transitions. In a particular example, the threshold charge density for tinting is between about 0.013-0.017 C/cm^2, and the threshold charge density for clearing is between about 0.016-0.020 C/cm^2. Additional threshold charge densities may be appropriate where the window is capable of transitioning between more than two states. For instance, if the device switches between four different optical states: A, B, C, and D, a different threshold charge or charge density may be used for each transition (e.g., A to B, A to C, A to D, B to A, etc.). In some embodiments, the threshold charge or threshold charge density is determined empirically. For instance, the amount of charge required to accomplish a particular transition between desired end states may be characterized for devices of different sizes. A curve may be fit for each transition to correlate the bus bar separation distance with the required charge or charge density. Such information may be used to determine the minimum threshold charge or threshold charge density required for a particular transition on a given window. In some cases, the information gathered in such empirical determinations is used to calculate an amount of charge or charge density that corresponds to a certain level of change (increase or decrease) in optical density.

[0236] Figure 9 shows an example of a process 900 for controlling multiple IGUs in accordance with some embodiments. Process 900 may be executed by a window controller. In some embodiments, blocks of process 900 may be performed in an order other than what is shown in Figure 9. In some implementations, two or more blocks of process 900 may be performed substantially in parallel. In some implementations, one or more blocks of process 900 may be omitted. At 910, process 900 applies a drive voltage to two or more IGUs in connection with a tint transition via a window controller. For example, the tint transition may be a common tint transition such that the two or more IGUs are all to transition from a first tint state to a second tint state. In some implementations, the window controller and each of the two or more

IGUs may each be operatively coupled to a common power supply. At 920, process 900 monitors, for each of the two or more IGUs, parameters indicative of a state of the tint transition. The parameters indicative of the state of the tint transition may include a V*oc* associated with each IGU and/or an amount of charge delivered to each IGU. At 930, process 900 determines whether any of the IGUs have completed the tint transition. For example, process 900 may determine that a particular IGU has completed the tint transition in response to determining that the V*oc* associated with the IGU meets or exceeds a target V*oc*. As another example, process 900 may determine that a particular IGU has completed the tint transition in response to determining that more than a target amount of charge has been delivered to the IGU. If, at 930, process 900 determines that none of the IGUs have completed the tint transition ("no" at 930), process 900 may loop back to block 920. Process 900 may loop through blocks 920 and 930 until process 900 determines that at least one IGU has completed the tint transition. Conversely, if, at 930, process 900 determines that at least one of the IGUs have completed the tint transition ("yes" at 930), process 900 may proceed to 940 and may transition the IGUs that have completed the tint transition (e.g., one or more IGUs that have completed the tint transition) to a hold voltage while continuing to apply the drive voltage to the remaining IGUs that have not completed the tint transition. It should be noted that, in instances in which two or more IGUs are transitioned to a hold voltage, the two or more IGUs may be transitioned to the same hold voltage and/or to different hold voltages. At 950, process 900 can determine whether all of the IGUs have completed the tint transition. For example, in some embodiments, process 900 can determine whether all of the IGUs have been transitioned to the hold voltage. If, at 950, process 900 determines that not all of the IGUs have completed the tint transition (e.g., there are remaining IGUs that have not completed the transition), or "no" at 950, process 900 can loop back to block 920 and can continue monitoring parameters indicative of the state of the tint transition for any remaining IGUs that have not yet completed the tint transition (and which have therefore not yet been transitioned to the hold voltage). Conversely, if, at 950, process 900 determines that all of the IGUs have completed the tint transition ("yes" at 950), process 900 can end.

[0237] In some embodiments, after an IGU has been transitioned to a hold voltage (e.g., in response to determining that the IGU has completed a tint transition), the IGU may be monitored, for example, to determine whether the IGU has deviated from the final tint state associated with the tint transition. In some implementations, in response to determining that the IGU has deviated from the final tint state, an additional voltage pulse may be applied to the IGU. This may have the effect of boosting or topping up the IGU, e.g., to counteract leakage

associated with the IGU that causes the IGU to deviate from the final state. In some embodiments, the voltage pulse may be a voltage pulse associated with the drive voltage.

[0238] In some embodiments, whether an IGU has deviated from a final tint state associated with a tint transition may be determined based on a measurement of a V*oc* associated with the IGU. For example, a determination that the IGU has deviated from the final state may be made in response to determining that the V*oc* associated with the IGU has dropped below a target V*oc*. In some embodiments, in instances in which a window controller controls multiple IGUs, the window controller can cause V*oc* associated with each IGU to be sequentially or simultaneously measured. In some embodiments, the window controller may then apply voltage pulses to one or more IGUs that have deviated from the final tint state.

[0239] Figure 10 shows an example process 1000 for monitoring IGUs after a tint transition in accordance with some embodiments. Blocks of process 1000 may be executed by a window controller. In some embodiments, two or more blocks of process 1000 may be performed substantially simultaneously. In some embodiments, one or more blocks of process 1000 may be omitted. Process 1000 begins at 1010 by monitoring, for two or more IGUs currently in a hold state, a deviation from a final tint state. The final tint state may be a tint state associated with a completed tint transition. The final tint state may be the same for each of the two or more IGUs. In some implementations, monitoring the deviation from the final tint state may include measuring a V*oc* associated with each IGU. At 1020, process 1000 determines whether any of the IGUs have deviated from the final tint state. For example, process 1000 may determine that an IGU has deviated from the final tint state in response to determining that a V*oc* of the IGU has dropped below a target V*oc*. If, at 1020, process 1000 determines that none of the IGUs have deviated from the final tint state ("no" at 1020), process 1000 loops back to block 1010. Conversely, if, at 1020, process 1000 determines that at least one IGU has deviated from the final tint state ("yes" at 1020), process 1000 proceeds to 1030 and applies a voltage pulse to any IGU (e.g., the at least one IGU) that has deviated from the final tint state. In some implementations, the voltage pulse may have a magnitude corresponding to the drive voltage, e.g., the drive voltage used to transition the IGU to the final tint state.

[0240] In some embodiments, a window controller (which may include one or more processors) may control multiple (e.g., two or more) IGUs. In some embodiments, the window controller and the multiple IGUs may each be operatively coupled to a power supply (e.g., a DC power supply). For example, the window controller may transmit instructions to the power supply that cause the power supply to provide a drive voltage to the multiple IGUs. In some implementations, each of the multiple IGUs is operatively coupled to the power supply via a switch. The switch may be toggled between an enabled state (e.g., a low impedance state) and a

disabled state (e.g., a high impedance state) such that, when a switch associated with an IGU is in an enabled state, the IGU receives the voltage (e.g., a drive voltage) provided by the power supply and such that when the switch is in a disabled state, the IGU is in a hold state and does not receive the voltage (e.g., drive voltage) provided by the power supply. In some implementations, the window controller (e.g., the one or more processors of the window controller) provide instructions to each switch that cause the switch to change state (e.g., from an enabled state to a disabled state, or vice versa). In some embodiments, a switch that operatively couples an IGU to the power supply may be an H-Bridge.

[0241] In some embodiments, each IGU is operatively coupled to a voltage measuring circuit that is configured to measure an open circuit voltage (e.g., V*oc*) across the IGU. In some embodiments, the voltage measuring circuit includes a differential amplifier (e.g., as shown in and described in connection with Figure 6). In some embodiments, a voltage measuring circuit is operatively coupled to a corresponding switch (e.g., H-Bridge). For example, a voltage measuring circuit may be configured to measure the V*oc* across an IGU when the corresponding switch is in a disabled state.

[0242] In some embodiments, a switch is operatively coupled to an IGU via a trunk line segment. In some embodiments, a voltage drop (e.g., V*drop*) is measured across each trunk line segment via a voltage measuring circuit. In one example, the voltage measuring circuit includes a differential amplifier (e.g., as shown in and described in connection with Figure 6). In some embodiments, a voltage drop measurement is provided to the window controller. For example, the window controller (e.g., the one or more processors of the window controller) may cause voltage provided by the power supply to be altered based on the voltage drop measurement (e.g., by transmitting instructions to the power supply to increase a voltage provided, where the increase is determined based at least in part on the voltage drop measurement). In some embodiments, voltage drop measurements may be provided to the window controller via a serial communications interface and/or using a serial communications protocol. An example of a serial communications interface includes the 1-WIRE interface.

[0243] In some embodiments, current through the multiple IGUs is measured using a current measuring circuit. In some embodiments, the current measuring circuit may include an operational amplifier (e.g., as shown in and described in connection with Figure 6). In some embodiments, current through the IGUs may be used to determine a charge delivered to the IGUs, for example, by integrating current over a period of time.

[0244] Figure 11 shows an example system 1100 for controlling multiple IGUs using a window controller in accordance with some embodiments. As illustrated, system 1100 includes a window controller 1102. Window controller 1102 may include one or more processors (e.g.,

hardware processors). Window controller 1102 controls IGUs 1104, 1106, 1108, and 1110. Each of IGUs 1104, 1106, 1108, and 1110 receive power from a power supply 1112. For example, in some embodiments, window controller 1102 transmits instructions to power supply 1112 that cause power supply 1112 to provide a drive voltage to any of IGUs 1104, 1106, 1108, and 1110. Each IGU is operatively coupled to power supply 1112 via an H-Bridge. For example, IGU 1104 is operatively coupled to power supply 1112 via an H-Bridge 1114, IGU 1106 is operatively coupled to power supply 1112 via an H-Bridge 1116, IGU 1108 is operatively coupled to power supply 1112 via an H-Bridge 1118, and IGU 1110 is operatively coupled to power supply 1112 via an H-Bridge 1120. In some embodiments, each IGU is operatively coupled to a V*oc* circuit (e.g., a circuit configured to measure an open circuit voltage). In some embodiments, a V*oc* circuit may comprise a differential amplifier. A V*oc* circuit may be operatively coupled to a corresponding H-Bridge. For example, V*oc* circuit 1122 is operatively coupled to H-Bridge 1114 and IGU 1104, V*oc* circuit 1124 is operatively coupled to H-Bridge 1116 and IGU 1106, V*oc* circuit 1126 is operatively coupled to H-Bridge 1118 and IGU 1108, and V*oc* circuit 1128 is operatively coupled to H-Bridge 1120 and IGU 1110. A V*oc* circuit may be configured to measure the open circuit voltage across the corresponding IGU when the corresponding H-Bridge is in a disabled (e.g., high impedance) state. Each IGU is operatively coupled to a V*drop* measuring circuit. A V*drop* measuring circuit measures a voltage drop across power distribution line segments to the IGU. For example, IGU 1104 is operatively coupled to V*drop* measuring circuit 1130, IGU 1106 is operatively coupled to V*drop* measuring circuit 1132, IGU 1108 is operatively coupled to Vdrop measuring circuit 1134, and IGU 1110 is operatively coupled to V*drop* measuring circuit 1136. Voltage drop measurements are transmitted to window controller 1102 via a serial communications interface 1140 (e.g., a 1-WIRE Master interface). Measurements are transmitted via a 1-Wire Bus. Current provided to IGUs 1104, 1106, 1108, and 1110 is measured using a current measuring circuit 1138. Current measuring circuit 1138 may comprise an operational amplifier operatively coupled to power supply 1112 and to each IGU. In some embodiments, data provided to window controller 1102 is digitized using Analog-to-Digital Converter (ADC) 1142.

[0245] In some implementations, a window controller system for controlling multiple IGUs may be implemented using a printed circuit board (PCB). It should be noted that, in instances in which multiple IGUs are controlled by a single window controller system, each IGU may differ in terms of size/dimensions, cable length that couples the IGU to the window controller system, device parameters, etc. In some cases, a window controller system may be capable of driving up to 10 Watts per IGU, or up to 15 Watts per IGU, thereby allowing each IGU to tint relatively fast.

[0246] In some embodiments, a window controller system PCB may be configured such that the PCB may be utilized in connection with various types of power and/or communications protocols/systems, such as a G.hn protocol that utilizes a coaxial cable to provide 48V DC power via a trunk line, G.hn protocol that utilizes a powerline to provide 110V AC power (e.g., from a building power supply), Power over Ethernet (POE), or the like. In other words, a common PCB may be populated with components according to the cable type and/or the protocol type to be used to provide power and/or data signals to the PCB, and therefore, to be utilized by the window controller system. Utilizing a single PCB layout and populating the PCB at manufacturing time may result in substantial cost-savings in window controller system manufacturing. Moreover, some embodiments find particularly advantageous use in buildings such as commercial office buildings or residential buildings. Some embodiments can be particularly suited and adapted for use in the construction of new buildings. For example, some embodiments are designed to work in conjunction with modern or novel heating, ventilation, and air conditioning (HVAC) systems, interior lighting systems, security systems, and power systems, as a single holistic efficient energy control system for an entire building or a campus of buildings . Some embodiments are particularly well-suited for integration with a building management system (BMS). Furthermore, embodiments having Ethernet, Wi-Fi, BLUETOOTH, G.hn or other communication capabilities may be well-suited for use in residential homes and other smaller-scale non-commercial applications. Other examples of embodiments that may be well suited for residential or other smaller-scale applications are described in U.S. Patent Application No. 16/949,855, entitled "CONTROLLER FOR OPTICALLY-SWITCHABLE WINDOWS," (Attorney Docket VIEWP041C3US), which is hereby incorporated by reference herein in its entirety.

[0247] In some implementations, the PCB may include a first set of components that are commonly utilized regardless of the power/communications system utilized. In other words, the first set of components may be affixed (e.g., soldered or otherwise connected to) the PCB regardless of the power/communications system used. The first set of components may include a processing unit, a memory device (e.g., a flash memory device), one or more IGU controllers each corresponding to a different IGU to be controlled using the processing unit, and/or connectors for each IGU to be controlled by the processing unit. The processing unit may perform many of the functions described herein as being performed by a window controller. For example, the processing unit may determine or control tint transitions for each of the multiple IGUs controlled by the processing unit, receive voltage and/or current data that may be utilized to determine error conditions, or the like. The processing unit may be operatively coupled to each IGU controller such that the processing unit controls each IGU controller, and each IGU

controller in turn controls each IGU. The processing unit may be communicatively coupled to another controller (e.g., a higher hierarchy controller, such as a network controller or a master controller). The communicative coupling may be a wired communication channel, or a wireless communication channel. For example, as shown in and described below in connection with Figure 16, the processing unit may utilize Wi-Fi and/or BLUETOOTH to communicate with a higher hierarchy controller or control panel. The processing unit may be configured to cause various information, such as tint scheduling information, window zone information, user preferences, or the like, in the memory device of the PCB such that the processing unit can utilize the stored information regardless of whether the processing unit is communicatively coupled to the control panel or higher hierarchy controller at a given time. For example, in some implementations, the processing unit may utilize the stored information to control each connected IGU in an "offline" mode. In some embodiments, the memory device may store information indicating the power/communications system to be used to provide power and/or data signals to the PCB.

[0248] The PCB may additionally include a second set of components, where the components included in the second set of components may be determined based on the power/communications system utilized. For example, the second set of components may include components such as analog front ends, splitters, adaptors, etc. that are particular to the power/communications system utilized. As a more particular example, in an instance in which a G.hn coaxial input is provided (e.g., as power/communications to the window controller system), the second set of components may include a splitter to operatively couple to the trunk line and/or a coaxial coupling circuit configured to provide power from the trunk line to an analog front end component. As another more particular example, in an instance in which a G.hn powerline input is provided, the second set of components may include a powerline coupling circuit configured to operatively couple to an AC power outlet (e.g., that provides 110V AC or 220V AC) and an analog front end component. The second set of components may further include an A-to-D converter configured to convert AC signals to DC signals. As yet another more particular example, in an instance in which an Ethernet input is provided, the second set of components may include various connectors and/or adaptors configured to receive an Ethernet cable, process power and/or data signals using PoE, or the like.

[0249] In some implementations, a cabling network (e.g., that provides power and/or communications) to one or more downstream devices (e.g., window controllers, sensor units or ensembles, etc.) may include at least one distribution junction, which may be referred to herein as a "splitter" or a "junction." In some embodiments, the distribution junction may be incorporated into a window controller system PCB (e.g., as shown in and described below in

connection with Figure 12). For example, in some implementations, the window controller PCB may be populated with one or more components that perform functions of the distribution junction. Distribution junctions (e.g., circuitry and functions performed by distribution junctions) are described in more detail in PCT/US/2021/017946, entitled "DATA AND POWER NETWORK OF A FACILITY," which is hereby incorporated by reference herein in its entirety and for all purposes. In some implementations, the distribution junction may include at least one connector. The distribution junction may distribute one or more time-varying signals and/or DC power within a network infrastructure. The distribution junction may couple together two or more circuits. As an example, the distribution junction may couple together at least two of an upstream circuit, a downstream circuit, and a branch circuit. The upstream and downstream circuits may be part of a network bus (also referred to herein as a trunk line). In some embodiments, a bus is a subsystem that is used to connect components transfer data (e.g., signal) and/or DC power between those components. The distribution junction can be passive, or active. The distribution junction may comprise active and passive components. The distribution junction may include ground one or more paths in the upstream, downstream, and branch circuits that are electrically coupled together. The distribution junction can include, or be operatively coupled to, a microprocessor. The cabling network may include a passive distribution junction and/or an active distribution junction. An active distribution junction has at least one active component. A passive distribution junction has passive component(s) and no active components. In some embodiments, the distribution junction provides impedance matching. In some embodiments, the distribution junction may comprise a transformer. For example, implementations of a distribution junction that utilizes a transformer can provide impedance matching. The impedance matching may serve to reduce (e.g., eliminate) unwanted signal reflections off of distribution junctions within the network infrastructure. In some embodiments, the distribution junction includes at least one filter. The distribution junction may include one or more low-pass filters, high-pass filters, and/or band-pass filters. The filters may serve to block certain frequencies from a branch circuit (e.g., when such frequencies are not utilized by that branch circuit) and/or from a downstream circuit (e.g., when no downstream circuits utilize such frequencies). By blocking such frequencies (e.g., signal portions), the filters may reduce noise in the network, e.g., as the signal propagates through the network (e.g., through the bus). In some embodiments, the distribution junction includes frequency shifting capabilities. In particular, the control panel and distribution junctions may frequency-shift one or more of the time-varying signals to reduce interference as the signals travel through the network. Signals may be shifted into regions of the spectrum available on the medium (e.g., coaxial cable) that are not being used. The distribution junction may include passive or active components that remove this frequency shift when

conveying signals from a network bus to a branch circuit and that insert this frequency shift when conveying signals from the branch circuit to the network bus. The control panel may include a G.hn head-end (or other component) that adds and removes frequency shifts to the time-varying signals as they are transmitted by and received at the control panel.

[0250] In some embodiments, a window controller system PCB may distribute power and/or data signals to and from a processing unit of the PCB. The data signals may include time-varying signals, which may be referred to as digital communications, digital data, or the like. The data signals may include power line communications signals, such as G.hn, HomePlug®, or HD-PLC compatible signals. The data signals may include signals compatible with the MoCA (multimedia over coax alliance) protocol. The data signals may include signals compatible with other protocols including Ethernet protocols such as 802.3bw, 802.3bp, 802.3ch, and/or 802.3cq. The first frequency window may extend from approximately 2 to approximately 200 MHz (e.g., such as used in the G.hn protocol). As an example, the first frequency windows may extend from approximately 500 MHz to approximately 600 MHz, from approximately 875 MHz to approximately 1 Ghz, and/or from approximately 1.15 to approximately 1.5 GHz.

[0251] Figure 12 is a schematic diagram of an example window controller system PCB according to certain implementations. As illustrated the PCB may include a first set of components which may be used regardless of the power/communication system utilized. Accordingly, the PCB may be populated with the first set of components for multiple different types of power/communication systems/protocols. The first set of components may include a DC/DC converter 1202, a processing unit 1204, a memory device 1206, an authentication device 1208, a set of IGU controllers (1210a, 1210b, 1210c, and 1210d) corresponding to a set of IGUs controlled by the window controller system, and a corresponding set of IGU connectors (1212a, 1212b, 1212c, 1212d).

[0252] DC/DC converter 1202 may be configured to convert a first DC voltage to a second DC voltage (e.g., from 48V to 5V, from 23 V to 5 V, from 5V to 48V etc.). DC/DC converter 1202 may provide DC power utilized by other components of the window controller system PCB, such as processing unit 1206, Ethernet controller 1218, memory device 1206, IGU control circuitry 1210a-1210d, etc. Note that DC/DC converter 1202 may be configured to boost or reduce (as appropriate) DC voltage.

[0253] Processing unit 1204 may be a processor that implements many of the functions described herein as performed by a window controller. For example, processing unit 1204 may determine tint levels for a particular IGU, determine drive voltages to be used, determine scheduling information that indicates that a particular IGU is to transition to a particular tint state at a given date/time, communicate with a higher level controller (e.g., a network controller or a

master controller) or a control panel, receive electrical characteristics (e.g., voltage information, current information such as current leakage) to identify error states or malfunctioning IGUs, or the like.

[0254] In some implementations, processing unit 1204 may control the set of IGUs in conjunction such that the IGUs achieve a given tint state concurrently or nearly concurrently (e.g., as described above in connection with Figures 9 and/or 10). In some implementations, processing unit 1204 may control initial ramp times across the set of IGUs such that the start times of a ramp-to-drive step are staggered across the IGUs. This may reduce power consumption by the window controller system to allow the IGUs to tint in conjunction with each other. In some implementations, processing unit 1204 may determine a ramp type and/or a ramp rate for a given tint transition to be performed by a given IGU. Note that ramp types and ramp rates are described above in more detail in connection with Figure 2. In some implementations, processing unit 1204 may determine a different ramp type and/or a different ramp rate for two different IGUs controlled by the window controller system. Accordingly, processing unit 1204 may then transmit instructions to each IGU controller (e.g., 1210a-1210d) that indicate the determined ramp type and/or the determined ramp rate such that two different IGUs controlled by the same window controller may be tinted using different ramp types and/or different ramp rates.

[0255] In some implementations, processing unit 1204 may receive override information. For example, the override may override a tint state determined based on scheduling information, current photosensor information, or the like. The override may be received via a local network of a facility. The override may be received from a user device (e.g., a mobile phone, a tablet computer, a laptop computer, etc.), e.g., via an application executing on the user device. The override may be received from a higher level controller or other computing device (which may include a cloud computing device, a control panel computing device installed in the facility, or any other suitable computing device). In some implementations, the override may be based on photosensor information (e.g., a photosensor that is part of a sky sensor device configured to determine sky temperature, a photosensor that is disposed proximate to the window being tinted, or the like).

[0256] Note that, in some implementations, processing unit 1204 may be a G.hn baseband processing unit, and there may be an additional processor or microcontroller 1205 operatively coupled between processing unit 1204 and each IGU controller 1210a-1210d. In such implementations, the functionality described herein as performed by processing unit 1204 may be performed by processing unit 1204, additional processor or microcontroller 1205, or split

between processing unit 1204 and additional processor or microcontroller 1205 in any suitable manner.

[0257] In some implementations, processing unit 1204 may be configured to enter a sleep mode. When in the sleep mode, processing unit 1204, as well as other components of the window controller system, may draw less power. A window controller sleep mode is described above in connection with Figure 6, as well as in PCT/US2016/041176, entitled "POWER MANAGEMENT FOR ELECTROCHROMIC WINDOW NETWORKS," which is hereby incorporated by reference herein in its entirety for all purposes. In some embodiments, processing unit 1204 may enter into a sleep mode responsive to a command from a higher level controller, processor, or other computing device to which the window controller system (and therefore, processing unit 1204) is communicatively coupled. The higher level controller, processor, or other computing device may be within the facility (e.g., connected via one or more trunk line segments to the window controller system, connected via a local area network, communicatively coupled via a wireless network, or the like), or remote from the facility (e.g., as a cloud controller or other remote computing device). Examples of higher level controllers, processors, or computing devices include a network controller, a master controller, an edge server, a control panel, or the like. In some embodiments, the window controller system (e.g., including processing unit 1204) may awaken from the sleep mode. In some implementations, awakening from the sleep mode may be responsive to a subsequent awakening command from the higher level controller, processor, or other computing device. Additionally or alternatively, in some embodiments, awakening from the sleep mode may be responsive to a predetermined duration of time elapsing (e.g., five minutes, thirty minutes, one hour, or the like).

[0258] Memory device 1206 may store device parameters (e.g., associated with each IGU controlled by the window controller system), tint scheduling information, or the like. In some implementations, processing unit 1204 may retrieve device parameters, scheduling information, or the like from memory device 1206. In some implementations, updated information (e.g., pushed from a network controller, a master controller, a cloud computing device, etc.) may be stored in memory device 1206. The updates may be stored on a periodic basis (e.g., every 24 hours, every week, every month, etc.) or on an aperiodic basis.

[0259] Authentication device 1208 may store any suitable encrypted information and may allow processing unit 1204 to securely communicate with and/or control the set of IGUs. In particular, authentication device 1208 may be a cryptography co-processor configured to accelerate the time required to perform various cryptography algorithms (e.g., those used in secure public key authentication). An example of a component that may be utilized is ATECC508A.

[0260] The window controller system may include a set of IGU controllers, where each IGU controller is operatively coupled to a given IGU. Processing unit 1204 may control each IGU controller, which in turn controls the corresponding IGU. Examples of IGU controllers are 1210a-1210d. Each IGU may be operatively coupled to the corresponding IGU controller via an IGU connected (e.g., 1212a-1212d). It should be noted that, in some implementations, processing unit 1204 may be capable of performing more computationally intensive computations or the like relative to the IGU controllers.

[0261] As described above, the PCB may additionally include a second set of components, where the second set of components which are affixed to the PCB based on the power/communications systems being utilized. In other words, the PCB may include, in the layout, space and traces for all possible seconds sets of components, where the PCB is populated based on the power/communication system to be utilized.

[0262] For example, in an instance in which Ethernet is being utilized (e.g., for PoE), the PCB may be populated with Ethernet connecter 1214, Ethernet controller 1216, and PoE circuitry 1218. Ethernet connector 1214 may be an Ethernet connector with Power over Ethernet (PoE) capabilities such that Ethernet connector 1214 is configured to provide power and data transmission capabilities. Ethernet connector 1214 may provide 100Base Ethernet and/or 1000Base Ethernet connectivity. Ethernet connector 1214 may be an RJ45 connector. Ethernet controller 1216 may be bidirectionally coupled with processing unit 1204 and may be configured to provide data in an appropriate physical layer format for subsequent transmission as Ethernet transmission. For example, Ethernet controller 1216 may be configured to decode Ethernet data from an end device (e.g., an IGU controller) and provide the unencoded data to processing unit 1204 for subsequent upstream transmission (e.g., to a control panel, to a higher level controller/processor/computing device, or the like). Ethernet controller 1214 may be configured to receive downstream data (e.g., from a control panel, from a higher level controller/processor/computing device, etc.) and provide the downstream data to processing unit 1204. PoE circuitry 1218 may be configured for sending and receiving data. PoE circuitry 1218 may be configured to make electrical power available for tramssion on physical lines that can also carry Ethernet formatted data.

[0263] As another example, in an instance in which G.hn is utilized in connection with a coaxial input, the PCB may be populated with a splitter 1220, a coaxial coupling circuit 1222, and a filter 1224. Splitter 1220 may be configured to split or tap power and/or data from a trunk line, which may be received via a coaxial cable. Note that the data may be formatted and/or encoded according to a G.hn protocol. Note that splitter 1220 may be considered part of a distribution junction, as described above. Coaxial coupling circuit 1222 may be configured to

couple signals from splitter 1220 to analog front end 1226, which is in turn coupled to processing unit 1204.

[0264] As yet another example, in an instance in which G.hn is utilized in connection with a powerline input, the PCB may be populated with a powerline coupling circuit 1228, an analog front end 1230, an AC to DC voltage converter 1234, and a filter 1232. Powerline coupling circuit 1228 may be configured to couple signals from the powerline input to analog front end 1230, which is in turn coupled to processing unit 1204. AC to DC voltage converter 1234 may be configured to convert the 120V (or 220V) AC voltage to a DC voltage, such as 48V. It should be noted that, in some implementations, analog front end 1226 may be utilized in connection with G.hn inputs for both a coaxial input and a powerline input.

[0265] It should be understood, that, in some implementations various other components may be provided in a window controller system. In some embodiments, such components may be mounted on a daughterboard that is operatively coupled to a primary window controller system PCB.

[0266] In some implementations, a window controller system may be housed in an enclosure suitable for plugging into a wall outlet. For example, in some implementations, a window controller system may be housed in a standard three-gang or four-gang outlet box. Such an enclosure may allow for relatively straightforward installation in residences (e.g., single family homes, apartment buildings, etc.) or small commercial spaces. The enclosure may include the outlet box, in which a window controller system PCB is disposed. An example of such a window controller system PCB is shown in and described above in connection with Figure 12. The window controller system PCB may include connectors for connecting to multiple (e.g., two, three, four, eight, etc.) IGUs. In some embodiments (e.g., in instances in which a G.hn powerline power/communications system is utilized) The enclosure may additionally include a powerline PCB for, e.g., receiving power and/or data signals via a G.hn powerline. In some implementations, a window controller system PCB and a powerline PCB may be separated by a fixed spacing. For example, in some implementations, the powerline PCB may be disposed above the window controller system PCB with a fixed gap between the two PCBs. The fixed gap may be achieved via a spacing element disposed on the back portion of the upper PCB (e.g., the PCB closer to a face plate that covers the enclosure) that rests on the upper portion of the lower PCB.

[0267] The window controller system may additionally include a face plate that covers the internal components (e.g., those disposed within the outlet box). The face plate may be fixed to the outlet box via screws or any other suitable mechanism. In some embodiments, at least one outer edge of the face plate may be angled, e.g., downward toward the outlet box. The face plate

may include one or more vents to allow for airflow. In some embodiments, the one or more vents may be along an outer edge, which may be an angled outer edge. Alternatively, in some embodiments, the one or more vents may be in a central portion of the face plate. It should be noted that, in some implementations, a heat sink may additionally be disposed within the outlet box. In such implementations, the heat sink and the one or more vents may in conjunction provide cooling for the window controller system. In some implementations, the face plate may additionally include one or more LEDs that provide indications of a current status of any components. For example, in some embodiments, there may be one LED corresponding to each IGU the window controller system may be connected to and/or control, where each LED indicates a current status of the corresponding IGU. By way of example, a red color may indicate an error status associated with the IGU (e.g., that the IGU is not currently properly connected to the window controller system, that there is an error or fault associated with the IGU, etc.), and a green color may indicate proper functioning of the corresponding IGU. Note that there may be one or more light pipes configured to guide light from the PCB within the enclosure to the face plate.

[0268] Figures 13A and 13B illustrate different views of a window controller system disposed within an enclosure. As illustrated, Figure 13A depicts a four-gang outlet box 1302. However, note that, in some implementations, a three-gang outlet box may be used instead. Disposed within the outlet box are window controller system PCB 1304 and powerline PCB 1306. Face plate 1310 may form a cover (e.g., to cover the internal components, such as PCB 1304 and PCB 1306) and may be affixed to outlet box 1302 via one or more screws 1316. Note the angled outer edges of face plate 1310. Window controller system PCB 1304 may include a set of IGU connectors, such as IGU connector 1308. Each IGU connector may be configured to connect to a single IGU controller by the window controller system. Although four IGU connectors are illustrated in Figure 13A, in some implementations, any suitable number of IGU connectors (e.g., two, three, eight, etc.) may be included corresponding to a number of IGUs that may be controller by the window controller system. Note that face plate 1310 includes a set of LEDs, such as LED 1312. Each LED may indicate a current status of a given component associated with the window controller system, such as a POE system, an IGU connected to the window controller system, or the like. Each LED may have a corresponding light pipe, such as light pipe 1314 for transmitting light up to face plate 1310. Additionally, face plate 1310 includes vents, such as vent 1318, which may allow for airflow and to provide cooling of components of the window controller system.

[0269] Figure 13B is a cross-sectional side view of an example window controller system and enclosure. As illustrated, Figure 13B depicts a portion of outlet box 1352. Face plate 1354 is

affixed (e.g., via one or more screws, not shown) to outlet box 1352. Note the angled outer edges (e.g., top and bottom outer edges) of face plate 1354. Within the enclosure formed by outlet box 1352 and face plate 1354 are a powerline PCB 1356 and a window controller system PCB 1358. Powerline PCB 1356 may be configured with an AC power connector 1360 configured to receive an AC power cable. Window controller system PCB 1358 may be configured with IGU connectors 1362 to couple multiple IGUs to the window controller system (e.g., to window controller system PCB 1358). Figure 13B additionally shows a heat sink 1364, which may be omitted in some implementations.

[0270] Figure 13C shows views of a window controller system PCB 1370 and a powerline PCB 1372 in accordance with some implementations. Note that window controller system PCB 1370 and powerline PCB 1372 may be disposed within an enclosure formed by an outlet box and a face plate, as shown in and described above in connection with Figures 13A and 13B. As illustrated, window controller system PCB 1370 may include multiple IGU connectors, such as IGU connector 1382. Each IGU connector may be configured to receive a cable which may connect window controller system PCB 1370 to a corresponding IGU controlled by window controller system PCB 1370. The IGU connectors may be M8K connectors in one example. Window controller system PCB 1370 may additionally include an Ethernet connector 1384. Ethernet connector 1384 may be an IX connector in one example. Panel 1376 illustrates a side view of window controller system 1370. IGU connectors 1382 can be seen in the side view of window controller system 1370.

[0271] As illustrated, powerline PCB 1372 may include an AC power connector 1386, which is also seen in the side view of powerline PCB 1372 in panel 1374. In some implementations, powerline PCB 1372 may be stacked above window controller system PCB 1370 within the enclosure. In such implementations, there may be a notch 1388 in window controller system PCB 1370 to allow for the cable that connects to AC power connector 1386.

[0272] As described above, in some implementations, a window controller system may be operatively coupled to a higher level controller (e.g., a network controller or a master controller) and/or a control panel, which may include a higher level processor that controls the window controller. In some embodiments, such a control panel may include an enclosure, which may be sized to fit comfortably in e.g., a residential electrical enclosure panel. By way of example, in some embodiments, a height and/or a width of the control panel may be less than about 24 inches, or less than about 18 inches. In some implementations, a depth of a control panel may be less than about 6 inches, or less than about 4 inches, allowing the control panel to easily be disposed flush to a fall when installed.

[0273] A control panel may include one or more processors or other computing devices. The processor may be configured for building or industrial automation. The control panel may additionally include a DC power supply (e.g., a 48V power supply). The control panel may include one or more AC line filters. The control panel may include one or more powerline bridges configured for providing networking capabilities to the control panel and any connected devices. The control panel may include an Ethernet port switch, which may be configured for any suitable number of ports (e.g., 8, 10, 12, or the like).

[0274] Figure 14 illustrates a diagram of an example of the Powerline communications control panel 1400. As illustrated, control panel 1400 includes an enclosure 1402. Control panel 1400 includes a set of DIN rail terminal blocks, such as DIN rail terminal block 1404. Each DIN rail terminal block may serve a protective function to electrical equipment that is part of control panel 1400, for example, by preventing short circuits. Control panel 1400 may include one or more AC line filters, such as AC line filter 1406. Line filters may prevent the downstream communication signals from propagating upstream to the AC circuits outside of the desired network broadcast. This may be important, e.g., when creating separate control units when used in, for example, an apartment building. Control panel 1400 may include one or more powerline bridges, such as powerline bridge 1408. Each powerline bridge 1408 may be configured to provide networking (e.g., data) capabilities to connected devices via the connection to the AC power circuit that feeds the downstream devices. The powerline bridge may incorporate the communication signals on top of the AC power carried signal. Control panel 1400 may include an Ethernet port switch 1410. Control panel 1400 may include a DC power supply 1412, which may be a 24V power supply, a 48V power supply, or the like.

[0275] In some implementations, a window controller system may be communicatively coupled with a higher level controller or processor, such as a network controller or a master controller. For example, the window controller system may receive instructions and/or data from the higher level controller or processor, where the instructions may include tint schedules (e.g., tint transition times and/or tint levels for various IGUs under control of the window controller system), instructions for implementing tint override commands, updates to firmware utilized by the window controller system, or the like. The window controller system may transmit data to the higher level controller or processor, such as electrical characteristics associated with the IGUs controlled by the window controller system. In some implementations, the higher level controller or processor may utilize the electrical characteristics (e.g., voltages across the IGU, leakage current, etc.) to identify faults or failures of the IGUs, to identify that a particular component of an IGU may be likely to fail in the near future, or the like. More detailed techniques for identifying error conditions are described above. As described above, in

some implementations, instructions received at the window controller system from the higher level controller or processor may be stored in local memory of the window controller system such that the window controller system is configured to use the stored instructions even when a communications channel between the window controller system and the higher level controller or processor is non-operational, e.g., in an "offline" mode. The window controller system may then update stored instructions when updated instructions are transmitted by the higher level controller or processor, e.g., in an "online" mode.

[0276] In some implementations, a higher level controller or processor may be implemented on an edge server that is within the facility and remote from the window controller system. For example, in an instance in which a window controller system is implemented in a given room of a residence (e.g., the living room, a bedroom, etc.) or a other facility (e.g., a conference room of an office building, etc.), the edge server may be within the facility but remote from the window controller system, such as in a building common area or other location where building operations equipment may be stored. Additionally or alternatively, the higher level controller or processor may be a remote cloud device that is remote from the facility.

[0277] In some implementations, the higher level controller or processor may receive data from one or more sensors disposed throughout the facility. The one or more sensors may include light detectors (e.g., photosensors), sound sensors, motion detectors, volatile organic compound (VOC) sensors, humidity sensors, temperature sensors, carbon dioxide sensors, or the like. The higher level controller or processor may then determine tint levels for particular windows or zones of windows based on the sensor data using, e.g., one or more machine learning models. The tint levels may then be transmitted to a given window controller system that controls the particular windows or zones of windows. In some implementations, the higher level controller or processor may in turn be communicatively coupled with another controller or processor, which may be local or remote to the facility and which may be configured to transmit updates to machine learning models or the like to the intermediate controller (e.g., the edge computer). By way of example, as described above, a network controller may be coupled to a master controller.

[0278] Figures 15 and 16 illustrate two example manners in window controller systems may be communicatively coupled to higher level controllers or processors. In the example shown in Figure 15, a wired communication system is utilized. In contrast, in the example shown in Figure 16, a wireless communication system is utilized. It should be understood that the window controller systems shown in and described below in connection with Figures 15 and 16 may be configured using the window controller system PCB shown in and described above in connection with Figure 12, and may optionally be disposed in an enclosure that includes an outlet box and a face plate, as shown in and described above in connection with Figures 13A-

13C. For example, to utilize the wired communication system shown in Figure 15, the window controller system PCB may be configured with ethernet or G.hn powerline power/communication system components. To utilize the wireless communication system shown in Figure 16, the window controller system PCB may be configured with components that provide for Wi-Fi and/or BLUETOOTH communications. It should be noted that, in some implementations, components for performing wireless communications (e.g., Wi-Fi, BLUETOOTH, etc.) may be provided in a window controller system PCB via a daughterboard that is operatively coupled to a primary window controller system PCB, such as the PCB shown in and described above in connection with Figure 12.

[0279] Turning to Figure 15, window controllers 1502 and 1504 are shown. Each window controller may control multiple IGUs. For example, window controller 1502 controls IGUs 1506a-1506d, and window controller 1504 controls IGUs 1508a-1508d. Note that although four IGUs are depicted as controlled by each window controller, this is merely exemplary, and, in some implementations, each window controller may control any suitable number of IGUs (e.g., one, two, five, ten, fifteen, etc.). Communication and/or data may be provided to each window controller via networking components 1510, which may include a residential gateway 1512, an Ethernet/G.hn adapter 1514, and a thread border router 1516. In some implementations, thread border router 1516 may wirelessly communicate (e.g., using a Wi-Fi protocol) with one or more IoT devices communicatively coupled with thread border router 1516. In some embodiments, the components of networking components 1510 may be packaged together in a single housing.

[0280] As illustrated, residential gateway 1512 may be operatively coupled to an ethernet switch 1524, which may in turn be operatively coupled to an edge server 1522. Edge server 1522 may implement a higher level controller or processor relative to window controllers 1502 and/or 1504. For example, edge server 1522 may determine building operations control parameters, which may include tint transitions for IGUs coupled to window controllers 1502 and/or 1504. Edge server 1522 may determine the building operations control parameters (including tint transitions) based on sensor device data from one or more sensors disposed in the facility. Edge server 1522 may be operatively coupled to the internet 1526 via a router 1520. In some implementations, edge server 1522 may transmit and/or receive data, programs, etc. from a cloud device or remote server. For example, edge server 1522 may receive updates to one or more machine learning algorithms from the cloud device or remote server, where the machine learning algorithms are configured to utilize sensor data to determine building operations parameters. In some implementations, edge server 1522 may push collected building data, which may include sensor data and/or IGU data, to the cloud device or the remote server for data analysis.

[0281] Figure 16 illustrates an example system in which window controllers utilize a wireless communication protocol to communicate with a higher level controller in accordance with some implementations. As illustrated, a window controller 1602 controls IGUs 1606a-1606d, and a window controller 1604 controls IGUs 1608a-1608d. Each window controller may be configured to utilize one or more wireless communication protocols, such as Wi-Fi, BLUETOOTH, etc. For example, each window controller system PCB (as shown in and described above in connection with Figure 12) may be populated with one or more chips or components configured to provide wireless communication capabilities to the window controller system. Note that both the number of window controllers illustrated and the number of IGUs controlled by each window controller is merely exemplary. There may be any suitable number of window controllers within a facility (e.g., ten, twenty, one hundred, etc.) each controlling any suitable number of IGUs (one, two, five, ten, etc.). Each window controller may be communicatively coupled to a Wi-Fi access point 1610. Wi-Fi access point 1610 may be disposed within the facility. Each window controller may communicate with a remote cloud controller 1612 using a wireless communication protocol, such as Wi-Fi. Note that, in this example, the facility need not have an edge computer, control panel, or other higher level controller local to the facility, which may be advantageous in some residential contexts, such as a single family home or the like. Similar to what is described above, cloud controller 1612 may be communicatively coupled to window controllers 1602 and/or 1604 to send and receive data, instructions, programs, etc. For example, cloud controller 1612 may be configured to push instructions, tint schedules, or the like to window controllers 1602 and/or 1604. As another example, window controllers 1602 and/or 1604 may be configured to transmit electrical characteristics and/or other IGU data to cloud controller 1612, e.g., to diagnose faults or error conditions. Similar to what is described above in connection with Figures 12 and 15, window controllers 1602 and/or 1604 may be configured to operate in an "offline" mode where stored instructions may be utilized when the window controller is not communicatively coupled to cloud controller 1612.

[0282] In some implementations, multiple window controllers within a facility or a region of a facility (e.g., window controllers 1602 and/or 1604) configured for Wi-Fi and/or BLUETOOTH communications, may be configured to form a mesh network. In some such implementations, window controllers may be capable of wirelessly communicating (e.g., using Wi-Fi and/or BLUETOOTH) with one another without utilizing an intermediate node. Moreover, in some implementations, where multiple window controllers are connected in a network (e.g., a mesh network), wireless power may be delivered from one window in the network to another. In some implementations, each window controller configured with wireless

capability (e.g., Wi-Fi and/or BLUETOOTH capability) may be considered a radio node or a client. The mesh network may additionally include mesh routers and gateways, where the routers forward traffic to and from the gateways. In some implementations, a gateway may be connected with the network. In some embodiments, the radio nodes or clients (e.g., the window controllers of the mesh network) may form a "mesh cloud," to store and/or process data collected by individual window controllers. In some implementations, data from the mesh cloud may be transferred to, e.g., a cloud controller or other cloud computing or storage device for, e.g., further processing/analysis. Wireless mesh networks are further described in U.S. Patent No. 10,303,035, entitled "SELF-CONTAINED EC IGU," which is hereby incorporated by reference herein in its entirety for all purposes.

[0283] Example Embodiments:

Clause 1: A window controller system for controlling multiple optically switchable devices, the window controller system comprising: a printed circuit board comprising a first plurality of footprints to which a first plurality of components is mounted and a second plurality of footprints, wherein a subset of the second plurality of footprints is populated by a second plurality of components, and wherein the first plurality of components comprises: a plurality of insulated glass unit (IGU) controllers, each configured to control an IGU of a corresponding plurality of IGUs operatively coupled to the window controller system, and a processing unit configured to control each of the plurality of IGU controllers, and wherein the second plurality of components is selected based on a cable type and/or a protocol type used to provide power and data signals to the printed circuit board.

Clause 2: The window controller system of clause 1, wherein the protocol type comprises at least one of: a home networking protocol, or a power over ethernet (PoE) protocol.

Clause 3: The window controller system of clause 2, wherein the home networking protocol is a G.hn protocol.

Clause 4: The window controller system of any one of clauses 1-3, wherein the cable type comprises at least one of: a coaxial cable, a powerline cable, or a twisted conductor pair cable.

Clause 5: The window controller system of any one of clauses 1-4, wherein the protocol type comprises G.hn, and wherein the second plurality of components comprise one of: (i) a first subset of components configured to utilize a G.hn coaxial input, or (ii) a second subset of components configured to utilize a G.hn powerline input.

Clause 6: The window controller system of clause 5, wherein the first subset of components configured to utilize the G.hn coaxial input comprise a splitter configured to operatively couple to a trunk line, and a coaxial coupling circuit configured to provide power from the trunk line to an analog front end component.

Clause 7: The window controller system of clause 6, wherein the analog front end component is operatively coupled to the processing unit.

Clause 8: The window controller system of clause 5, wherein the second subset of components configured to utilize the G.hn powerline input comprise a powerline coupling circuit configured to operatively couple to an AC power outlet, and an analog front end component.

Clause 9: The window controller system of clause 8, wherein the analog front end component is operatively coupled to the processing unit.

Clause 10: The window controller system of clause 9, wherein the second set of components further comprise an A-to-D converter (ADC) configured to convert AC signals from the AC power outlet to DC signals.

Clause 11: The window controller system of any one of clauses 1-10, wherein the protocol type comprises G.hn, and wherein the second plurality of components comprise a third subset of components to be utilized with either a G.hn coaxial input or a G.hn powerline input.

Clause 12: The window controller system of any one of clauses 1-11, wherein the first plurality of components further comprises a memory device.

Clause 13: The window controller system of clause 12, wherein the memory device comprises a flash memory device.

Clause 14: The window controller system of any one of clauses 12 or 13, wherein the memory device is configured to store an indication of the cable type and/or the protocol type used to provide the power and the data signals.

Clause 15: The window controller system of any one of clauses 1-14, wherein the processing unit is configured to characterize leakage current from one or more IGUs of the corresponding plurality of IGUs.

Clause 16: The window controller system of any one of clauses 1-15, wherein the processing unit is further configured to communicate with a computing device.

Clause 17: The window controller system of clause 16, wherein the computing device is a cloud device remote from a premises on which the window controller system is installed.

Clause 18: The window controller system of any one of clauses 16 or 17, wherein the processing unit is further configured to: detect an anomalous event based on the characterized leakage current; and responsive to detecting the anomalous event, cause data associated with the characterized leakage current to be transmitted to the computing device.

Clause 19: The window controller system of clause 18, wherein the computing device is a cloud computing device.

Clause 20: The window controller system of clause 18, wherein the computing device is housed in a control panel installed in a facility in which the window controller system is installed.

Clause 21: The window controller system of any one of clauses 1-20, further comprising a memory configured to store window zone information associated with the plurality of IGUs and/or tint schedule information associated with the plurality of IGUs.

Clause 22: The window controller system of clause 21, wherein the processing unit is configured to: receive the window zone information and/or the tint schedule information from a computing device; and cause the window zone information and/or the tint schedule information to be stored in the memory.

Clause 23; The window controller system of any one of clauses 1-22, wherein the processing unit is configured to: receive an override to a tint command for at least one IGU of the plurality of IGUs; and cause the at least one IGU to transition to a tint state based on the received override.

Clause 24: The window controller system of clause 23, wherein the override is received via an application executing on a mobile device.

Clause 25: The window controller system of any one of clauses 23 or 24, wherein the override is received via a local network to which the processing unit is communicatively coupled.

Clause 26: The window controller system of any one of clauses 23-25, wherein the override is based on information obtained by a sky sensor device configured to measure solar radiation external to a facility in which the window controller system is installed.

Clause 27: The window controller system of clause 26, wherein the override is received via a local network to which the sky sensor device and the processing unit are communicatively coupled.

Clause 28: The window controller system of any one of clauses 1-27, wherein the processing unit is configured to cause each IGU controller to apply a drive voltage having a ramp rate to the corresponding IGU, wherein a first ramp rate associated with a first IGU of the plurality of IGUs is different than a second ramp rate associated with a second IGU of the plurality of IGUs.

Clause 29: The window controller system of clause 28, wherein the drive voltage applied to the first IGU is at a time point that is offset from application of the drive voltage to the second IGU.

Clause 30: The window controller system of any one of clauses 1-29, wherein the processing unit is configured to communicate with at least one computing device configured to control building operations of a facility the window controller system is disposed in.

Clause 31: The window controller system of clause 30, wherein the at least one computing device is an edge server.

Clause 32: The window controller system of any one of clauses 1-31, wherein the processing unit is configured to enter a sleep mode responsive to a signal from at least one computing device.

Clause 33: The window controller system of any one of clauses 1-32, wherein the first plurality of components further comprises at least one component configured to provide one or more wireless communications channels.

Clause 34: The window controller system of clause 33, wherein the one or more wireless communications channels utilize at least one of a Wi-Fi communication protocol or a BLUETOOTH communication protocol.

Clause 35: The window controller system of clause 34, wherein the processing unit is configured to communicate using the Wi-Fi communication protocol with a cloud device.

Clause 36: The window controller system of clause 35, wherein the processing unit is configured to receive one or more updates from the cloud device using the Wi-Fi communication protocol.

Clause 37: The window controller system of any one of clauses 34-36, wherein the processing unit is configured to participate in a mesh network with at least one other processing unit associated with at least one other window controller system.

Clause 38: The window controller system of any one of clauses 1-37, further comprising a second processing unit operatively coupled between the processing unit and each of the plurality of IGU controllers, wherein the second processing unit is configured to control each of the plurality of IGU controllers.

Clause 39: A window controller system for controlling multiple optically switchable devices, the window controller system comprising: an enclosure, wherein the enclosure comprises a 3-gang or a 4-gang outlet box and a faceplate; and a printed circuit board comprising a plurality of components, the plurality of components comprising: a plurality of insulated glass unit (IGU) controllers, each configured to control an IGU of a corresponding plurality of IGUs operatively coupled to the window controller system; and a processing unit configured to control each of the plurality of IGU controllers.

Clause 40: The window controller system of clause 39, further comprising a heat sink disposed in the enclosure.

Clause 41: The window controller system of any one of clauses 39 or 40, wherein the face plate is affixed to the 3-gang or 4-gang outlet box via a plurality of screws.

Clause 42: The window controller system of any one of clauses 39-41, wherein at least one edge of the face plate is angled.

Clause 43: The window controller system of any one of clauses 39-42, wherein the face plate comprises a plurality of air vents.

Clause 44: The window controller system of any one of clauses 39-43, further comprising a cooling fan disposed within the enclosure.

Clause 45: The window controller system of any one of clauses 39-44, further comprising a plurality of light-emitting diode (LED) indicators.

Clause 46: The window controller system of any one of clauses 39-45, further comprising a second printed circuit board configured for providing power and/or data via a powerline system.

Clause 47: The window controller system of any one of clauses 39-46, wherein the second printed circuit board comprises an AC power connector.

Clause 48: The window controller system of any one of clauses 46 or 47, wherein the second printed circuit board is disposed above the printed circuit board and closer to the face plate than the printed circuit board.

Clause 49: The window controller system of clause 48, wherein the printed circuit board comprises a notch configured to pass through a powerline cable through the notch to the second printed circuit board.

Clause 50: A method for controlling tinting of multiple optically switchable devices, comprising: (a) applying, via a window controller, a drive voltage to two or more optically switchable devices in connection with a tint transition from a first tint state to a second tint state, wherein applying the drive voltage to the two or more optically switchable device comprises providing a voltage from a power supply common to the two or more optically switchable devices; (b) monitoring, for each of the two or more optically switchable devices, parameters indicative of a state of the tint transition; (c) repeating (a) and (b) until it is determined that one or more optically switchable devices have completed the tint transition to the second tint state; (d) in

response to determining at (c) that the one or more optically switchable devices have completed the tint transition to the second state, applying hold voltages to the one or more optically switchable devices while concurrently continuing to apply the drive voltage to a remaining one or more optically switchable devices of the two or more optically switchable devices that have not completed the tint transition; and (e) repeating (b) – (d) until the two or more optically switchable devices have completed the tint transition.

Clause 51: The method of clause 50, wherein monitoring the parameters indicative of the state of the tint transition comprises measuring, for each of the two or more optically switchable devices separately, an open circuit voltage (Voc) and/or a charge delivered.

Clause 52: The method of clause 51, wherein determining that the one or more optically switchable devices have completed the tint transition to the second tint state comprises (i) determining that the Voc measured for the one or more optically switchable devices meets or exceeds a target Voc and/or (ii) determining that the charge delivered to the one or more optically switchable devices meets or exceeds a target charge to be delivered to the one or more optically switchable devices.

Clause 53: The method of any one of clauses 50-52, further comprising: (f) measuring, for an optically switchable device for which the hold voltage is being applied, a hold state Voc; (g) determining whether the hold state Voc is less than a predetermined Voc threshold; and (h) in response to determining that the hold state Voc is less than the predetermined Voc threshold, applying a voltage pulse to the optically switchable device

Conclusion

[0284] In one or more aspects, one or more of the functions described may be implemented in hardware, digital electronic circuitry, analog electronic circuitry, computer software, firmware, including the structures disclosed in this specification and their structural equivalents thereof, or in any combination thereof. Certain implementations of the subject matter described in this document also can be implemented as one or more controllers, computer programs, or physical structures, for example, one or more modules of computer program instructions, encoded on a computer storage media for execution by, or to control the operation of window controllers, network controllers, and/or antenna controllers. Any disclosed implementations presented as or for electrochromic windows can be more generally implemented as or for switchable optical devices (including windows, mirrors, etc.).

[0285] Various modifications to the embodiments described in this disclosure may be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other implementations without departing from the spirit or scope of this disclosure. Thus, the claims are not intended to be limited to the implementations shown herein, but are to be accorded the widest scope consistent with this disclosure, the principles and the novel features disclosed herein. Additionally, a person having ordinary skill in the art will readily appreciate, the terms "upper" and "lower" are sometimes used for ease of describing the figures, and indicate relative positions corresponding to the orientation of the figure on a properly oriented page, and may not reflect the proper orientation of the devices as implemented.

[0286] Certain features that are described in this specification in the context of separate implementations also can be implemented in combination in a single implementation. Conversely, various features that are described in the context of a single implementation also can be implemented in multiple implementations separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claimed combination may be directed to a subcombination or variation of a subcombination.

[0287] Similarly, while operations are depicted in the drawings in a particular order, this does not necessarily mean that the operations are required to be performed in the particular order shown or in sequential order, or that all illustrated operations be performed, to achieve desirable results. Further, the drawings may schematically depict one more example processes in the form of a flow diagram. However, other operations that are not depicted can be incorporated in the example processes that are schematically illustrated. For example, one or more additional operations can be performed before, after, simultaneously, or between any of the illustrated operations. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system components in the implementations described above should not be understood as requiring such separation in all implementations, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products. Additionally, other implementations are within the scope of the following claims. In some cases, the actions recited in the claims can be performed in a different order and still achieve desirable results.

CLAIMS

What is claimed is:

1. A window controller system for controlling multiple optically switchable devices, the window controller system comprising:

a printed circuit board comprising a first plurality of footprints to which a first plurality of components is mounted and a second plurality of footprints, wherein a subset of the second plurality of footprints is populated by a second plurality of components, and wherein the first plurality of components comprises:

a plurality of insulated glass unit (IGU) controllers, each configured to control an IGU of a corresponding plurality of IGUs operatively coupled to the window controller system, and

a processing unit configured to control each of the plurality of IGU controllers,

and wherein the second plurality of components is selected based on a cable type and/or a protocol type used to provide power and data signals to the printed circuit board.

2. The window controller system of claim 1, wherein the protocol type comprises at least one of: a home networking protocol, or a power over ethernet (PoE) protocol.

3. The window controller system of claim 2, wherein the home networking protocol is a G.hn protocol.

4. The window controller system of claim 1, wherein the cable type comprises at least one of: a coaxial cable, a powerline cable, or a twisted conductor pair cable.

5. The window controller system of claim 1, wherein the protocol type comprises G.hn, and wherein the second plurality of components comprise one of: (i) a first subset of components configured to utilize a G.hn coaxial input, or (ii) a second subset of components configured to utilize a G.hn powerline input.

6. The window controller system of claim 5, wherein the first subset of components configured to utilize the G.hn coaxial input comprise a splitter configured to operatively couple to a trunk line, and a coaxial coupling circuit configured to provide power from the trunk line to an analog front end component.

7. The window controller system of claim 6, wherein the analog front end component is operatively coupled to the processing unit.

8. The window controller system of claim 5, wherein the second subset of components configured to utilize the G.hn powerline input comprise a powerline coupling circuit configured to operatively couple to an AC power outlet, and an analog front end component.

9. The window controller system of claim 8, wherein the analog front end component is operatively coupled to the processing unit.

10. The window controller system of claim 9, wherein the second set of components further comprise an A-to-D converter (ADC) configured to convert AC signals from the AC power outlet to DC signals.

11. The window controller system of claim 1, wherein the protocol type comprises G.hn, and wherein the second plurality of components comprise a third subset of components to be utilized with either a G.hn coaxial input or a G.hn powerline input.

12. The window controller system of claim 1, wherein the first plurality of components further comprises a memory device.

13. The window controller system of claim 12, wherein the memory device comprises a flash memory device.

14. The window controller system of claim 12, wherein the memory device is configured to store an indication of the cable type and/or the protocol type used to provide the power and the data signals.

15. The window controller system of claim 1, wherein the processing unit is configured to characterize leakage current from one or more IGUs of the corresponding plurality of IGUs.

16. The window controller system of claim 1, wherein the processing unit is further configured to communicate with a computing device.

17. The window controller system of claim 16, wherein the computing device is a cloud device remote from a premises on which the window controller system is installed.

18. The window controller system of claim 16, wherein the processing unit is further configured to:

detect an anomalous event based on the characterized leakage current; and

responsive to detecting the anomalous event, cause data associated with the characterized leakage current to be transmitted to the computing device.

19. The window controller system of claim 18, wherein the computing device is a cloud computing device.

20. The window controller system of claim 18, wherein the computing device is housed in a control panel installed in a facility in which the window controller system is installed.

21. The window controller system of claim 1, further comprising a memory configured to store window zone information associated with the plurality of IGUs and/or tint schedule information associated with the plurality of IGUs.

22. The window controller system of claim 21, wherein the processing unit is configured to:

receive the window zone information and/or the tint schedule information from a computing device; and

cause the window zone information and/or the tint schedule information to be stored in the memory.

23. The window controller system of claim 1, wherein the processing unit is configured to:

receive an override to a tint command for at least one IGU of the plurality of IGUs; and

cause the at least one IGU to transition to a tint state based on the received override.

24. The window controller system of claim 23, wherein the override is received via an application executing on a mobile device.

25. The window controller system of claim 23, wherein the override is received via a local network to which the processing unit is communicatively coupled.

26. The window controller system of claim 23, wherein the override is based on information obtained by a sky sensor device configured to measure solar radiation external to a facility in which the window controller system is installed.

27. The window controller system of claim 26, wherein the override is received via a local network to which the sky sensor device and the processing unit are communicatively coupled.

28. The window controller system of claim 1, wherein the processing unit is configured to cause each IGU controller to apply a drive voltage having a ramp rate to the corresponding IGU, wherein a first ramp rate associated with a first IGU of the plurality of IGUs is different than a second ramp rate associated with a second IGU of the plurality of IGUs.

29. The window controller system of claim 28, wherein the drive voltage applied to the first IGU is at a time point that is offset from application of the drive voltage to the second IGU.

30. The window controller system of claim 1, wherein the processing unit is configured to communicate with at least one computing device configured to control building operations of a facility the window controller system is disposed in.

31. The window controller system of claim 30, wherein the at least one computing device is an edge server.

32. The window controller system of claim 26, wherein the processing unit is configured to enter a sleep mode responsive to a signal from the at least one computing device.

33. The window controller system of claim 1, wherein the first plurality of components further comprises at least one component configured to provide one or more wireless communications channels.

34. The window controller system of claim 33, wherein the one or more wireless communications channels utilize at least one of a Wi-Fi communication protocol or a BLUETOOTH communication protocol.

35. The window controller system of claim 34, wherein the processing unit is configured to communicate using the Wi-Fi communication protocol with a cloud device.

36. The window controller system of claim 35, wherein the processing unit is configured to receive one or more updates from the cloud device using the Wi-Fi communication protocol.

37. The window controller system of claim 34, wherein the processing unit is configured to participate in a mesh network with at least one other processing unit associated with at least one other window controller system.

38. The window controller system of claim 1, further comprising a second processing unit operatively coupled between the processing unit and each of the plurality of IGU controllers, wherein the second processing unit is configured to control each of the plurality of IGU controllers.

39. A window controller system for controlling multiple optically switchable devices, the window controller system comprising:

an enclosure, wherein the enclosure comprises a 3-gang or a 4-gang outlet box and a faceplate; and

a printed circuit board comprising a plurality of components, the plurality of components comprising: a plurality of insulated glass unit (IGU) controllers, each configured to control an IGU of a corresponding plurality of IGUs operatively coupled to the window controller system; and a processing unit configured to control each of the plurality of IGU controllers.

40. The window controller system of claim 39, further comprising a heat sink disposed in the enclosure.

41. The window controller system of claim 39, wherein the face plate is affixed to the 3-gang or 4-gang outlet box via a plurality of screws.

42. The window controller system of claim 39, wherein at least one edge of the face plate is angled.

43. The window controller system of claim 39, wherein the face plate comprises a plurality of air vents.

44. The window controller system of claim 39, further comprising a cooling fan disposed within the enclosure.

45. The window controller system of claim 39, further comprising a plurality of light-emitting diode (LED) indicators.

46. The window controller system of claim 39, further comprising a second printed circuit board configured for providing power and/or data via a powerline system.

47. The window controller system of claim 46, wherein the second printed circuit board comprises an AC power connector.

48. The window controller system of claim 46, wherein the second printed circuit board is disposed above the printed circuit board and closer to the face plate than the printed circuit board.

49. The window controller system of claim 48, wherein the printed circuit board comprises a notch configured to pass through a powerline cable through the notch to the second printed circuit board.

50. A method for controlling tinting of multiple optically switchable devices, comprising:

(a) applying, via a window controller, a drive voltage to two or more optically switchable devices in connection with a tint transition from a first tint state to a second tint state, wherein applying the drive voltage to the two or more optically switchable device comprises providing a voltage from a power supply common to the two or more optically switchable devices;

(b) monitoring, for each of the two or more optically switchable devices, parameters indicative of a state of the tint transition;

(c) repeating (a) and (b) until it is determined that one or more optically switchable devices have completed the tint transition to the second tint state;

(d) in response to determining at (c) that the one or more optically switchable devices have completed the tint transition to the second state, applying hold voltages to the one or more optically switchable devices while concurrently continuing to apply the drive voltage to a remaining one or more optically switchable devices of the two or more optically switchable devices that have not completed the tint transition; and

(e) repeating (b) – (d) until the two or more optically switchable devices have completed the tint transition.

51. The method of claim 50, wherein monitoring the parameters indicative of the state of the tint transition comprises measuring, for each of the two or more optically switchable devices separately, an open circuit voltage (Voc) and/or a charge delivered.

52. The method of claim 51, wherein determining that the one or more optically switchable devices have completed the tint transition to the second tint state comprises (i) determining that the Voc measured for the one or more optically switchable devices meets or exceeds a target Voc and/or (ii) determining that the charge delivered to the one or more optically switchable devices meets or exceeds a target charge to be delivered to the one or more optically switchable devices.

53. The method of claim 50, further comprising:

(f) measuring, for an optically switchable device for which the hold voltage is being applied, a hold state Voc;

(g) determining whether the hold state Voc is less than a predetermined Voc threshold; and

(h) in response to determining that the hold state Voc is less than the predetermined Voc threshold, applying a voltage pulse to the optically switchable device.

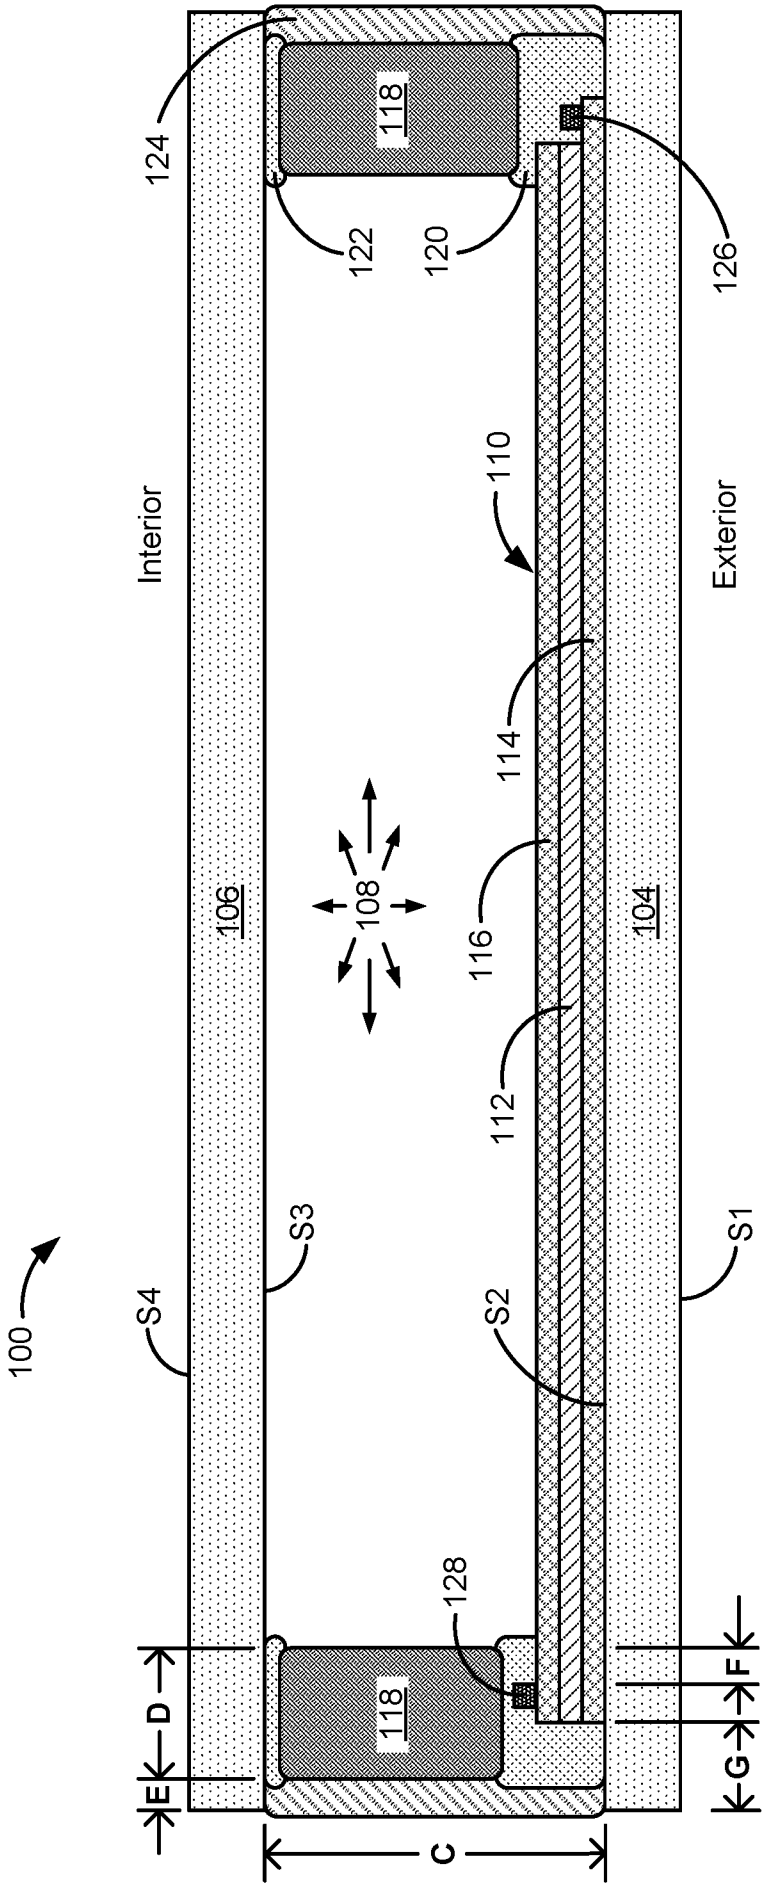

100
124
118
122
120
126
Interior
110
Exterior
114
106
108
104
116
112
S3
S1
S4
S2
128
118
D
E
F
G
C

FIGURE 1

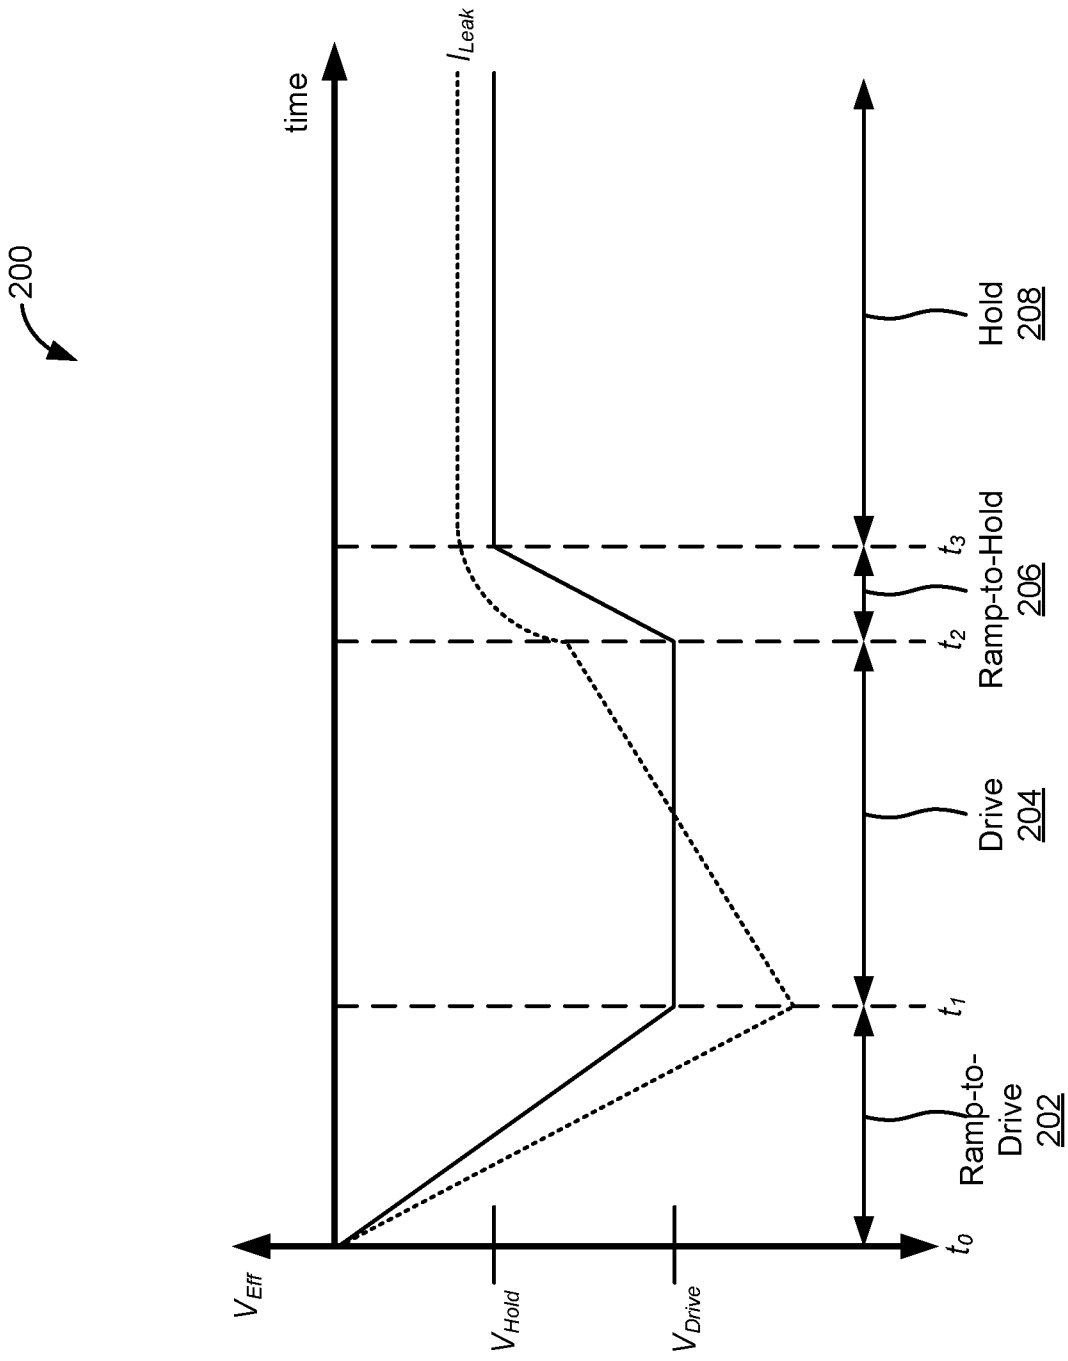

200
time
I_{Leak}
V_{Eff}
V_{Hold}
V_{Drive}
t_0
t_1
t_2
t_3
Ramp-to-Drive 202
Drive 204
Ramp-to-Hold 206
Hold 208

FIGURE 2

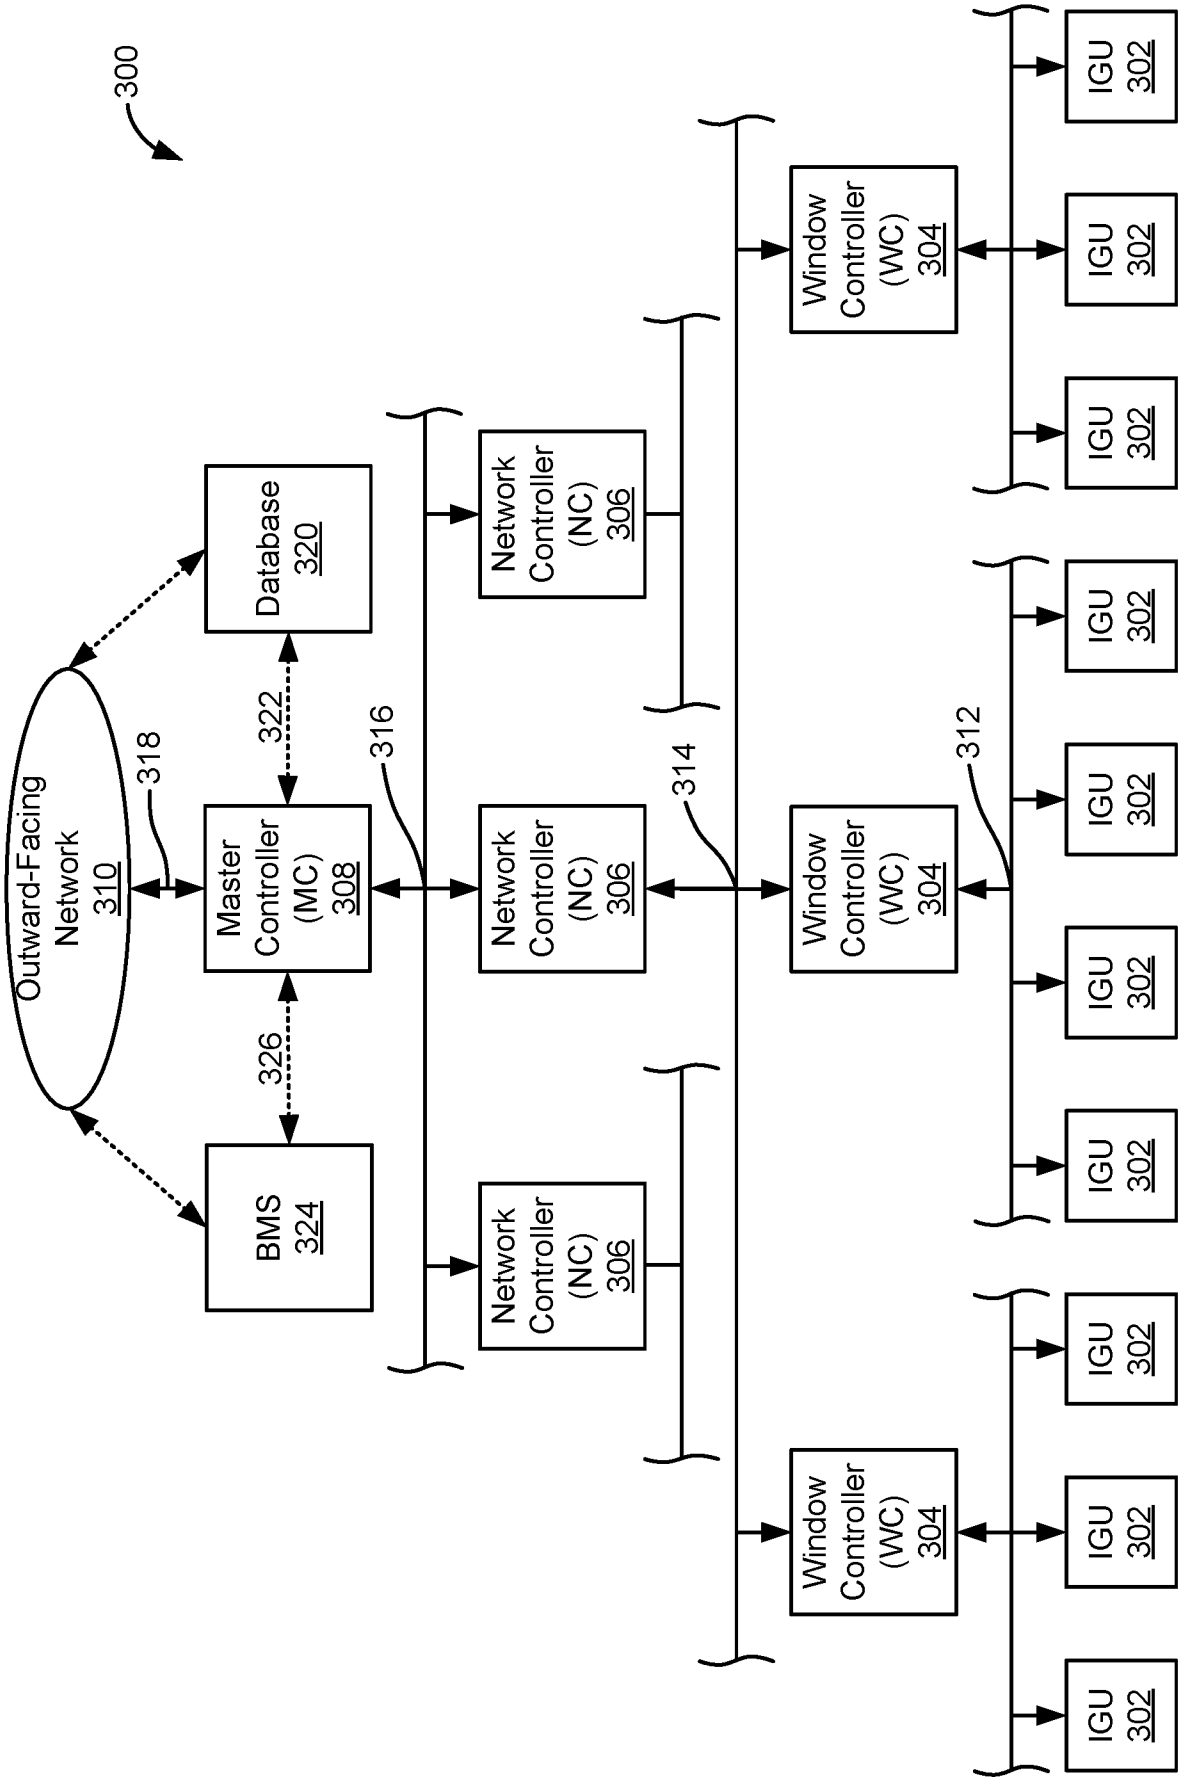

300
Outward-Facing Network 310
318
Master Controller (MC) 308
322
Database 320
326
BMS 324
316
Network Controller (NC) 306
Network Controller (NC) 306
Network Controller (NC) 306
314
Window Controller (WC) 304
Window Controller (WC) 304
Window Controller (WC) 304
312
IGU 302

FIGURE 3

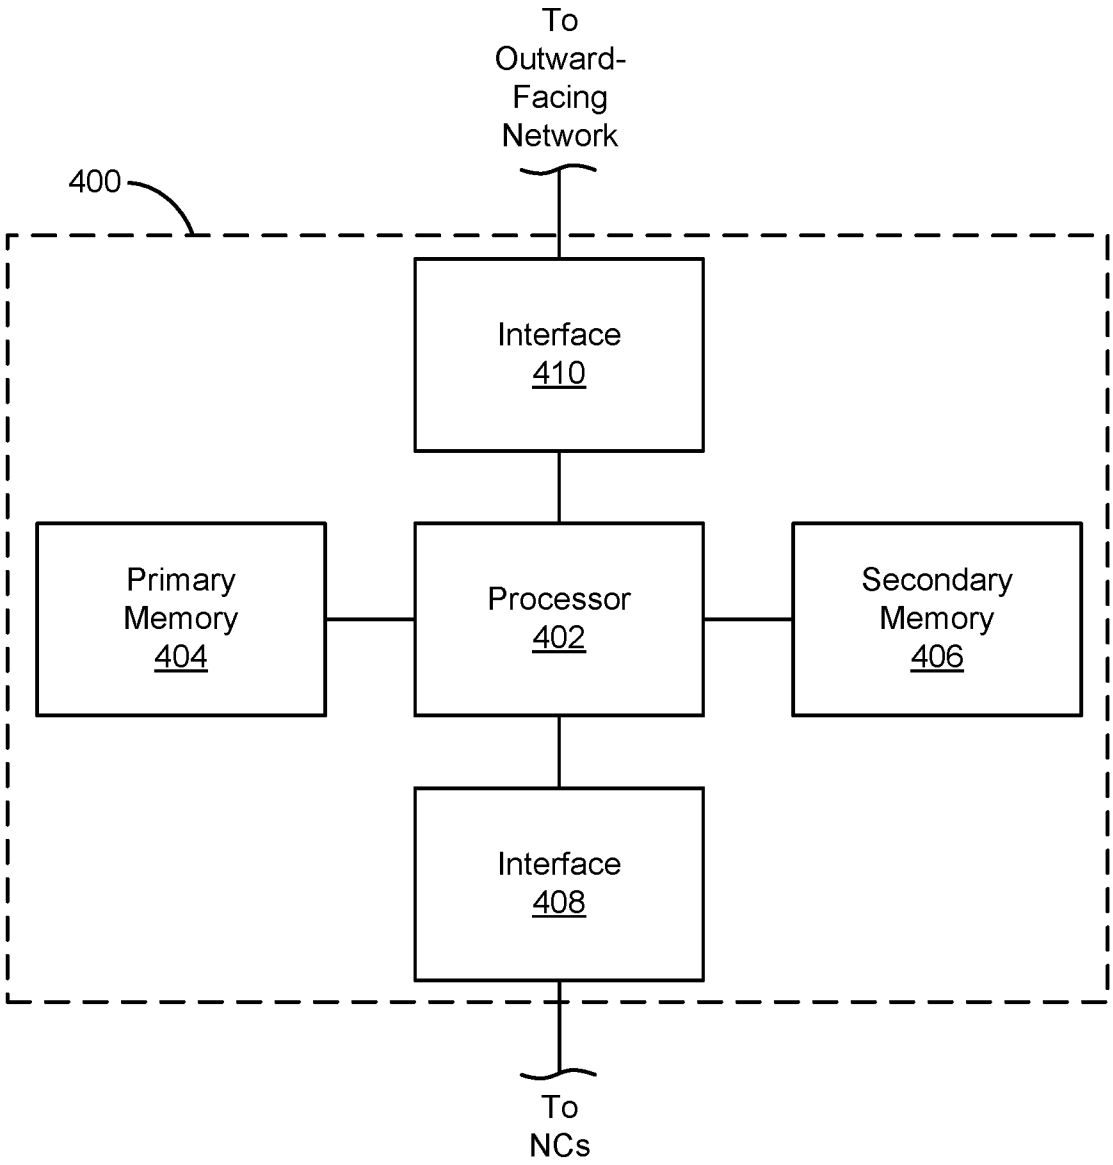

To
Outward-
Facing
Network
400
Interface
410
Primary
Memory
404
Processor
402
Secondary
Memory
406
Interface
408
To
NCs

FIGURE 4

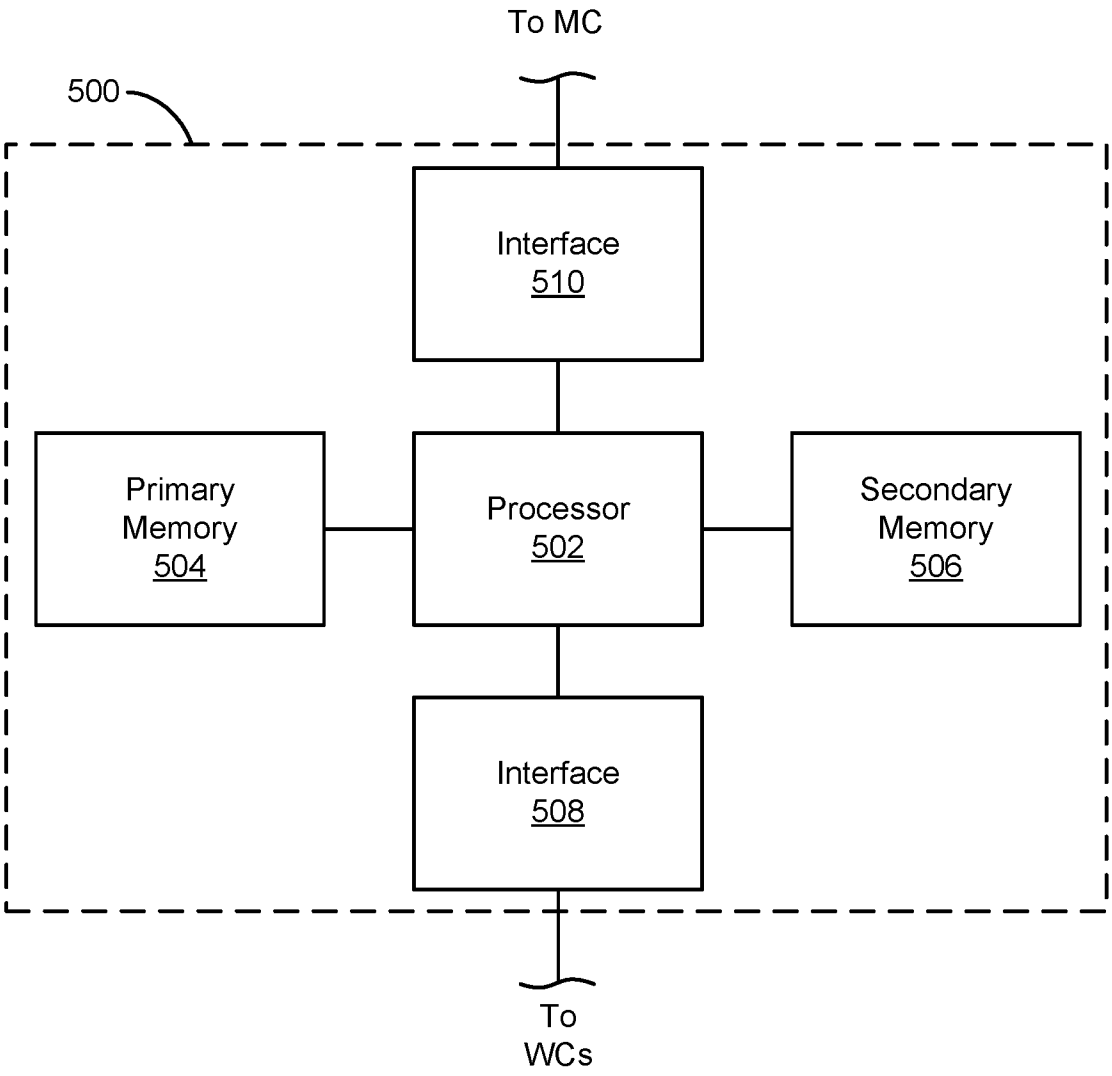

To MC
500
Interface
510
Primary
Memory
504
Processor
502
Secondary
Memory
506
Interface
508
To
WCs

FIGURE 5

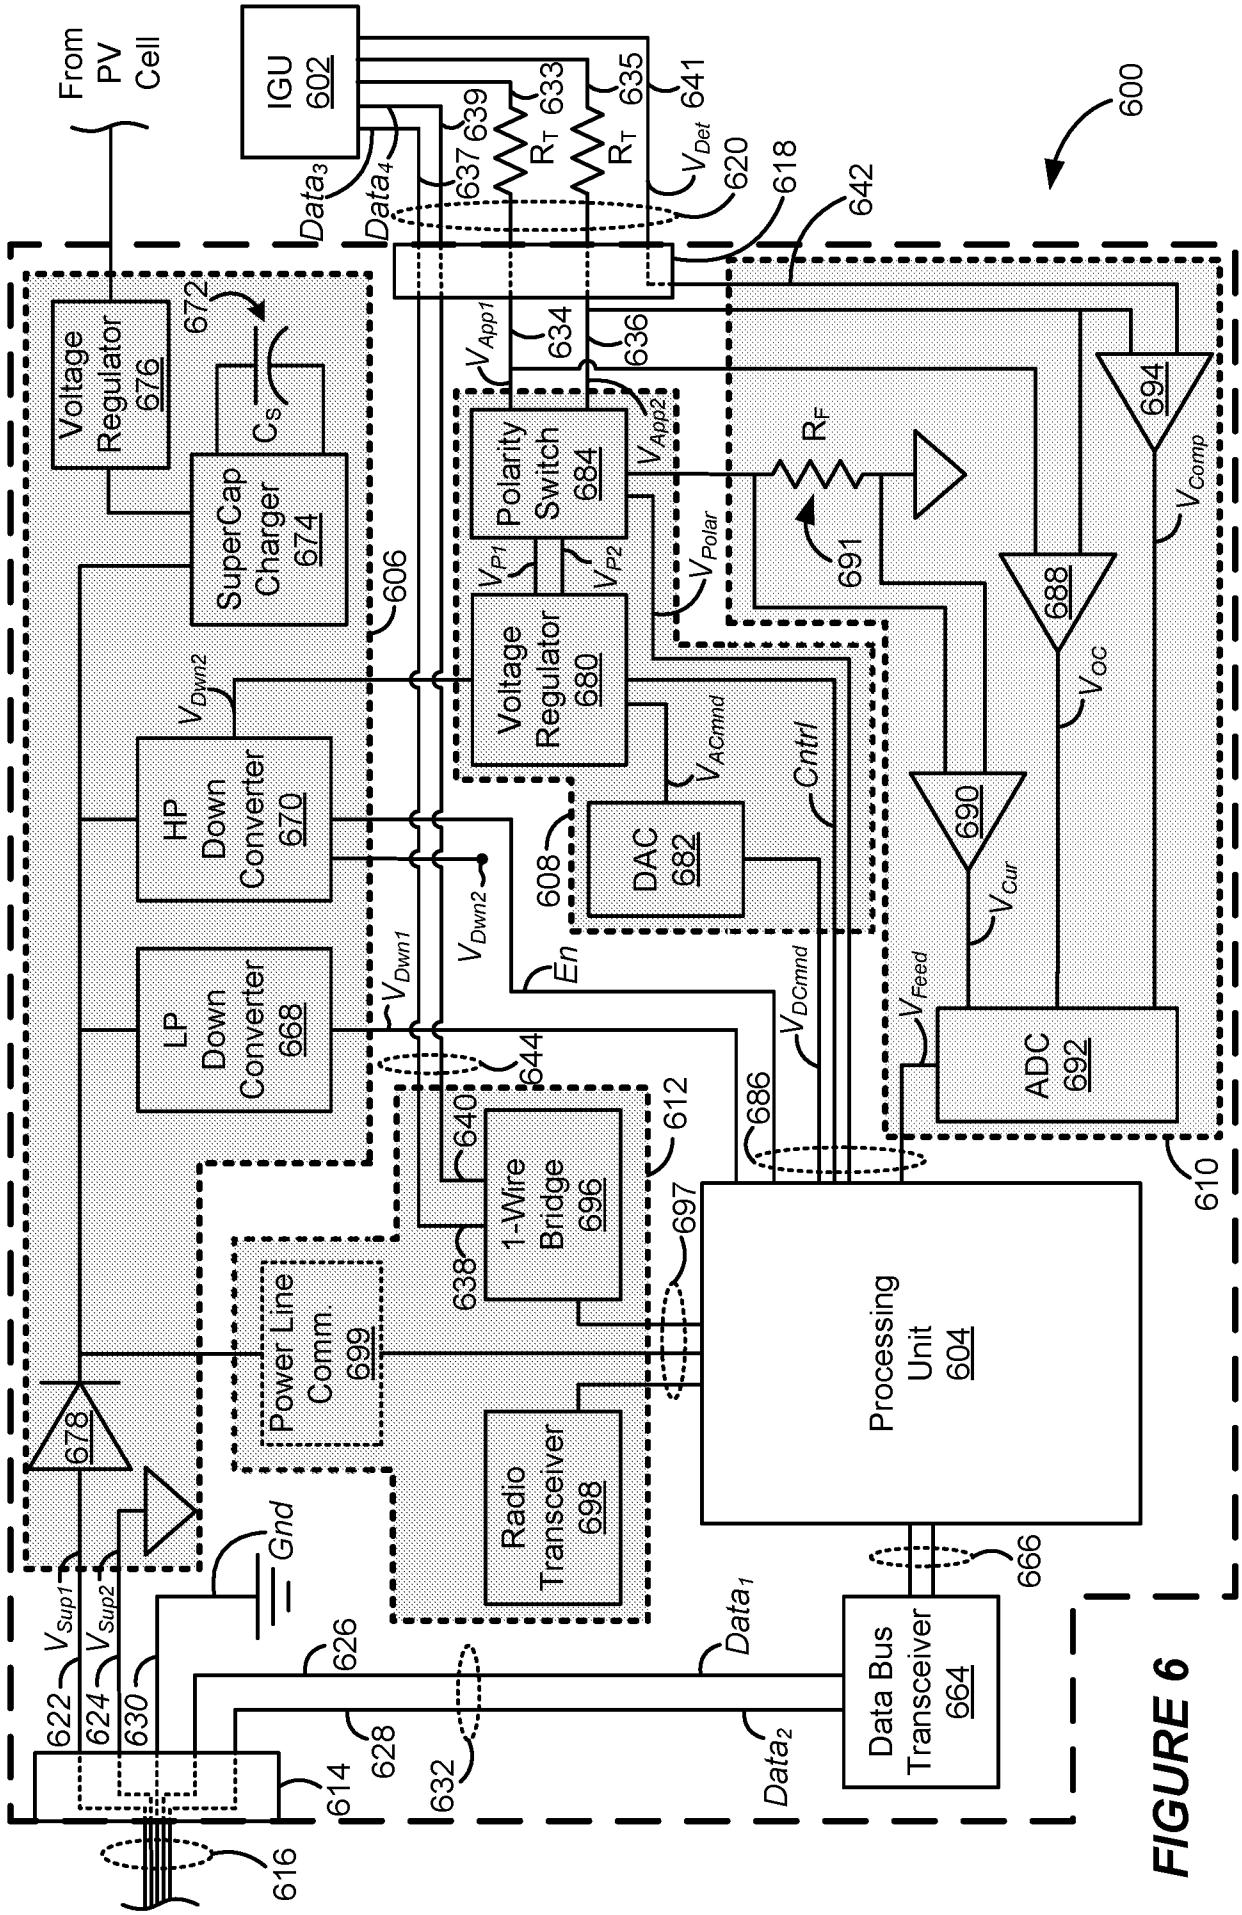

From PV Cell
IGU 602
Data3
Data4
637
639
RT 633
RT 635
641
VDet
620
618
642
600
Voltage Regulator 676
672
CS
SuperCap Charger 674
606
VApp1
634
636
Polarity Switch 684
VApp2
VP1
VP2
VPolar
RF
691
694
VComp
688
VOC
690
VCur
HP Down Converter 670
VDwn2
Voltage Regulator 680
VACmnd
Cntrl
DAC 682
608
VDwn2
En
LP Down Converter 668
VDwn1
644
VDCmnd
VFeed
ADC 692
686
612
640
1-Wire Bridge 696
638
697
610
Power Line Comm. 699
Processing Unit 604
678
Radio Transceiver 698
Gnd
VSup1
VSup2
622
624
630
626
628
632
614
616
Data1
Data2
Data Bus Transceiver 664
666

FIGURE 6

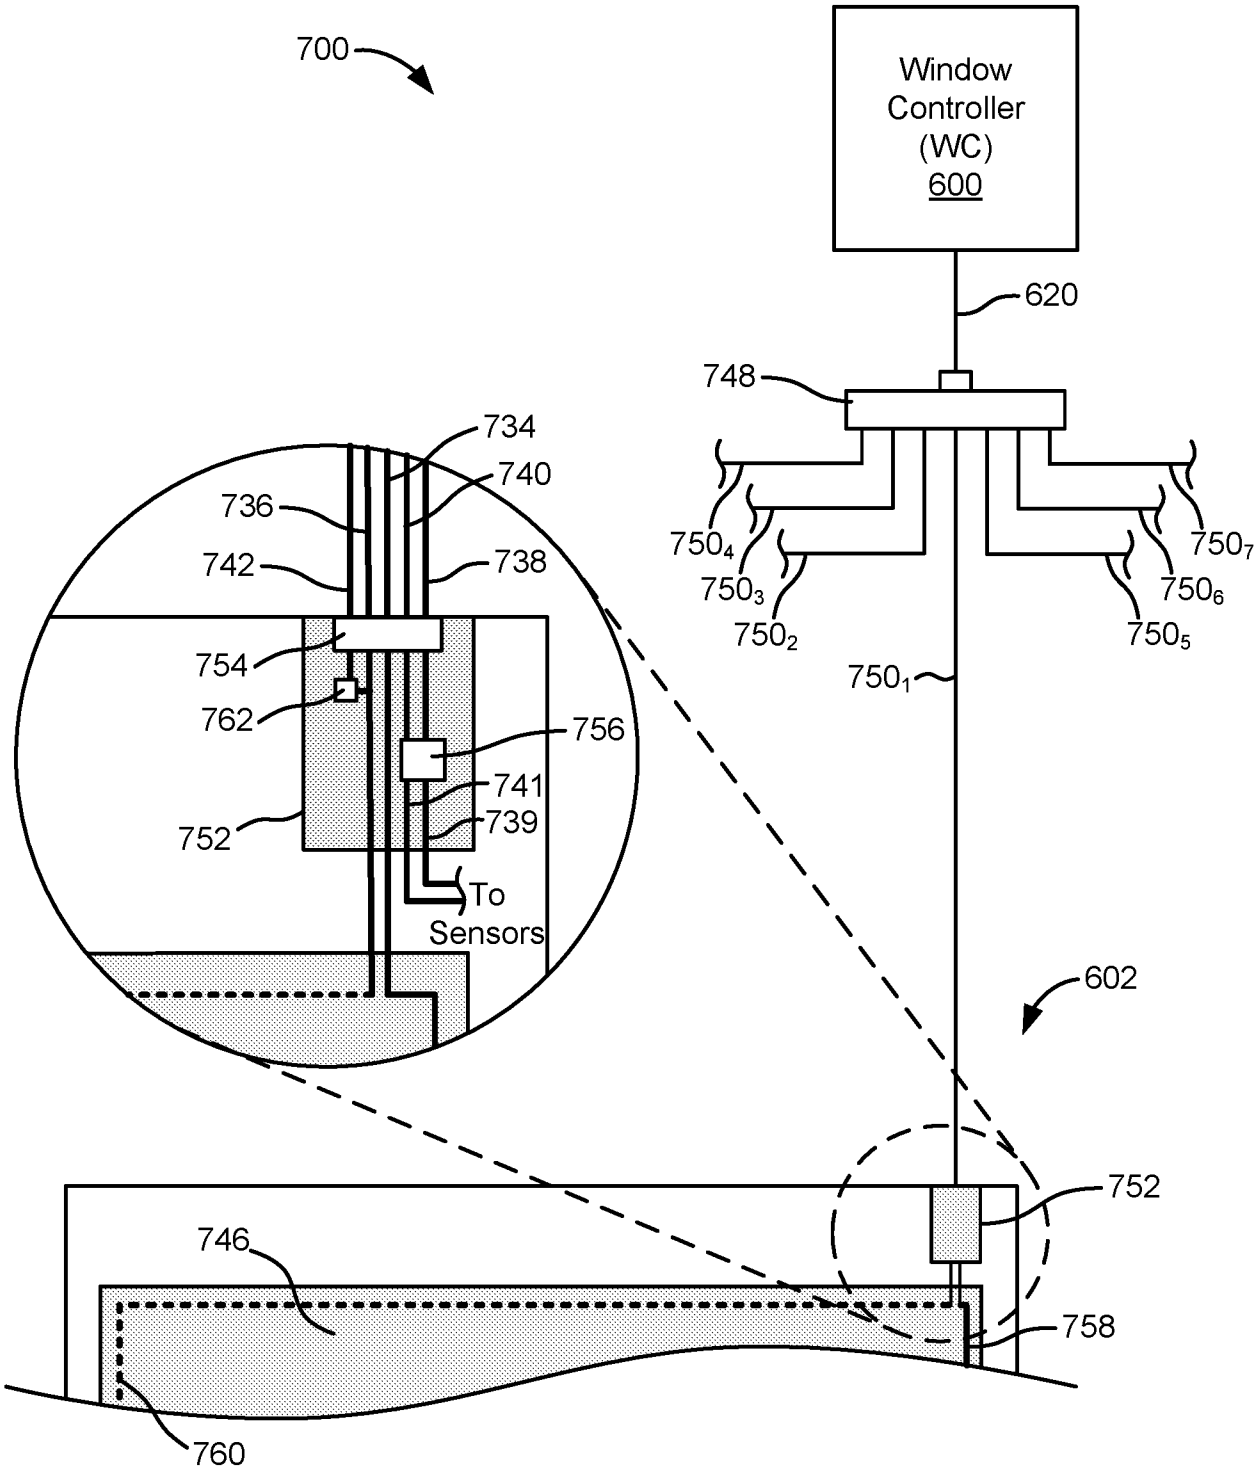

700
Window Controller (WC) 600
620
748
734
740
736
738
742
754
762
756
741
739
752
To Sensors
750_1
750_2
750_3
750_4
750_5
750_6
750_7
602
752
746
758
760

FIGURE 7

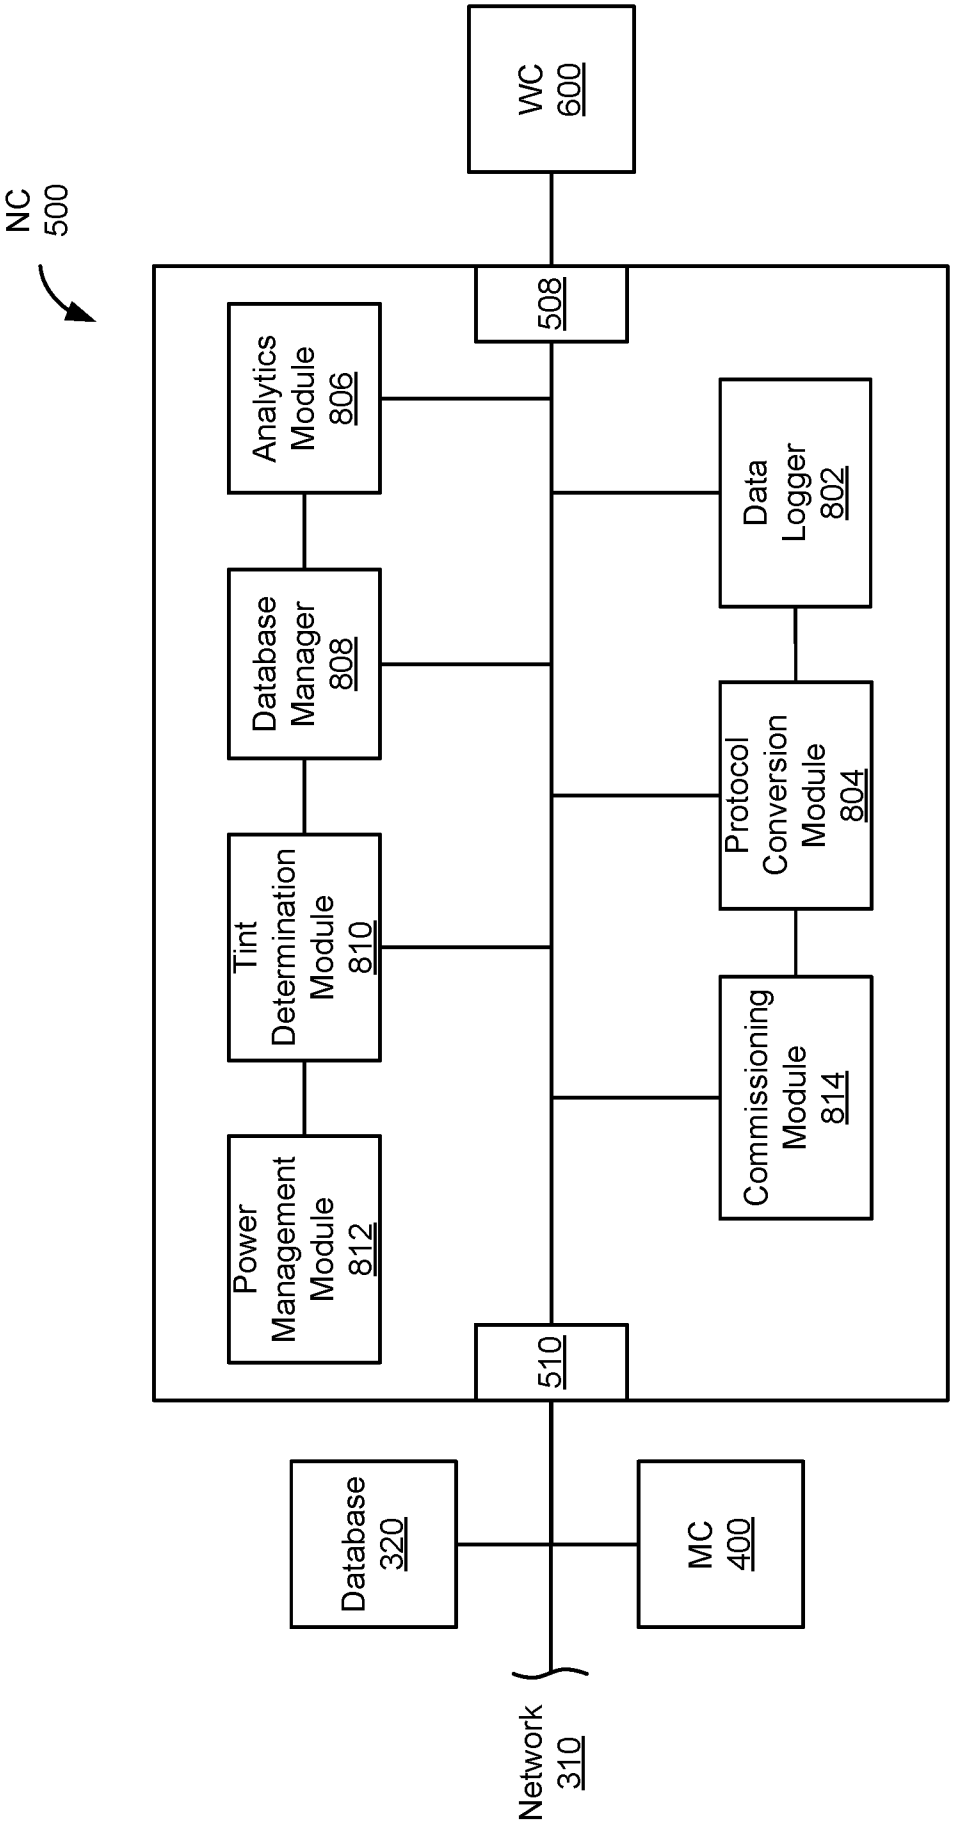

NC
500
WC
600
508
Analytics Module
806
Database Manager
808
Tint Determination Module
810
Power Management Module
812
Data Logger
802
Protocol Conversion Module
804
Commissioning Module
814
510
Database
320
MC
400
Network
310

FIGURE 8

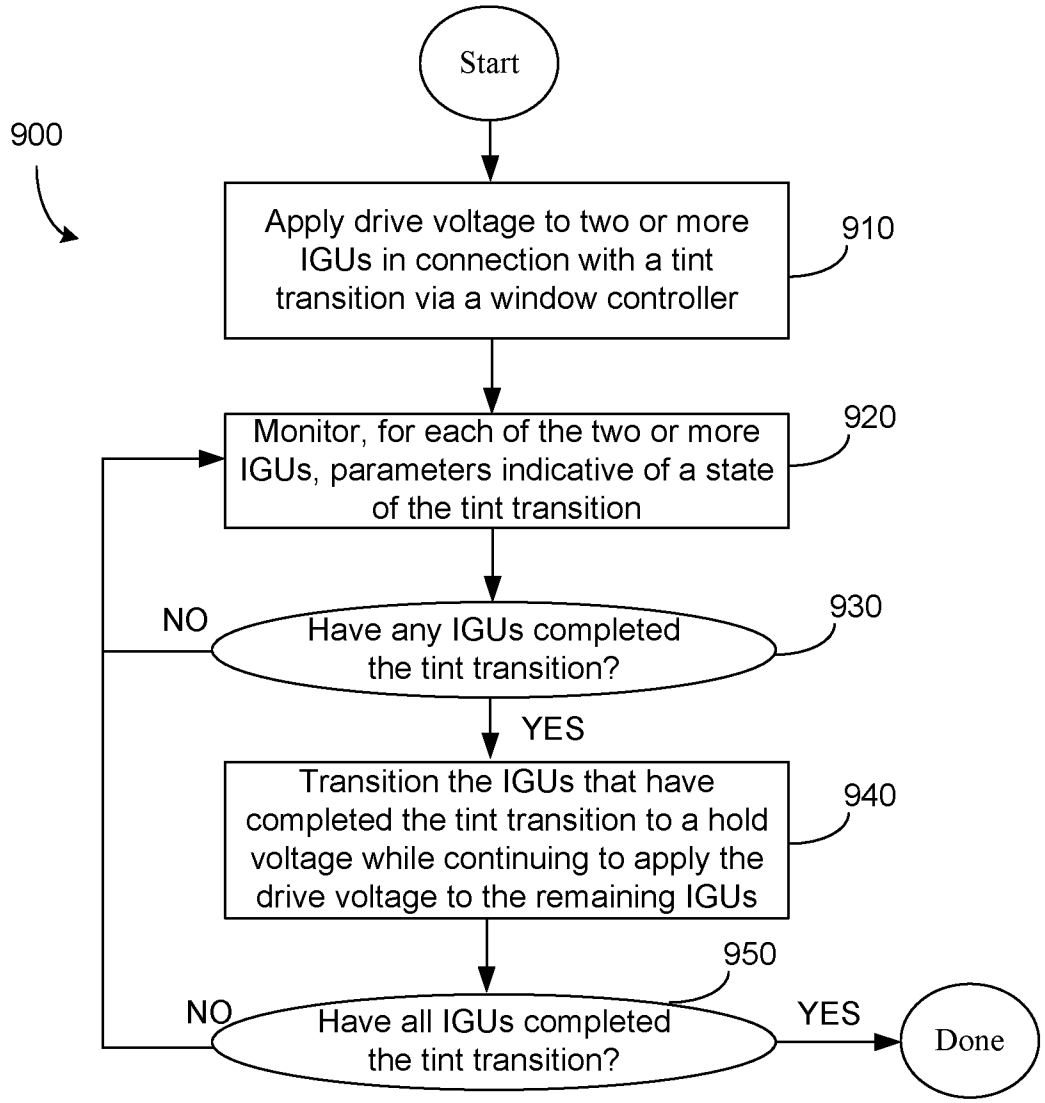

Start
900
Apply drive voltage to two or more IGUs in connection with a tint transition via a window controller
910
Monitor, for each of the two or more IGUs, parameters indicative of a state of the tint transition
920
NO
Have any IGUs completed the tint transition?
930
YES
Transition the IGUs that have completed the tint transition to a hold voltage while continuing to apply the drive voltage to the remaining IGUs
940
950
NO
Have all IGUs completed the tint transition?
YES
Done

FIGURE 9

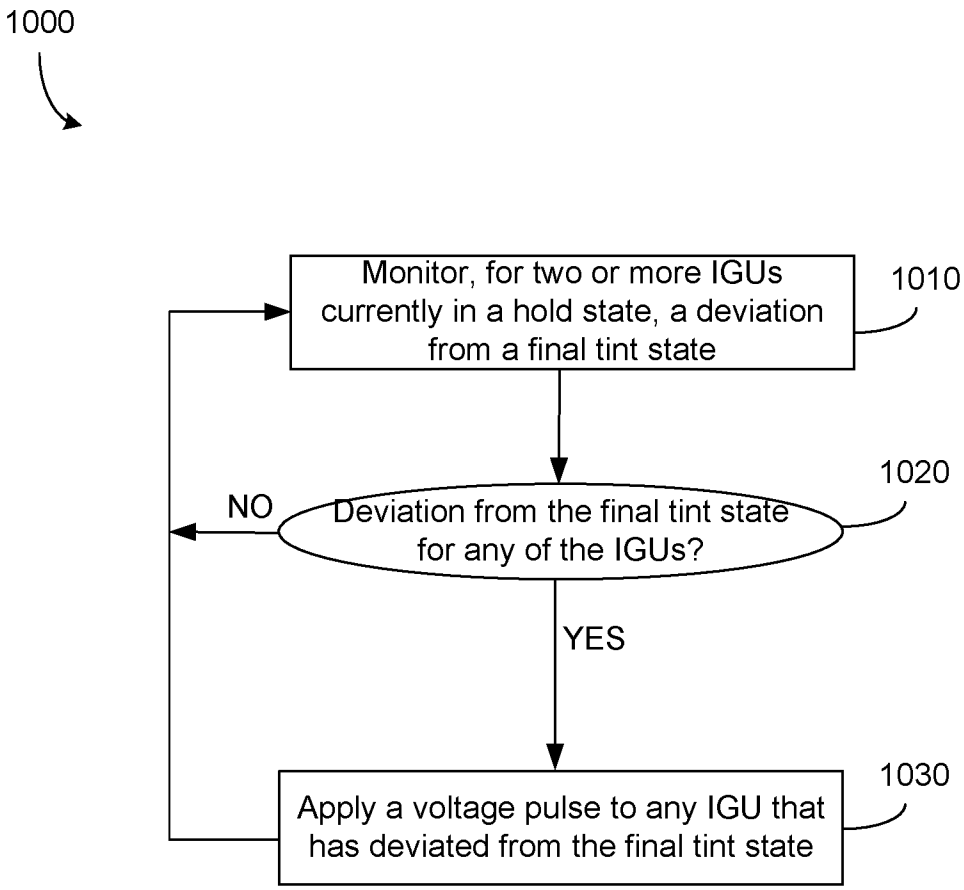

1000
Monitor, for two or more IGUs currently in a hold state, a deviation from a final tint state
1010
Deviation from the final tint state for any of the IGUs?
1020
NO
YES
Apply a voltage pulse to any IGU that has deviated from the final tint state
1030

FIGURE 10

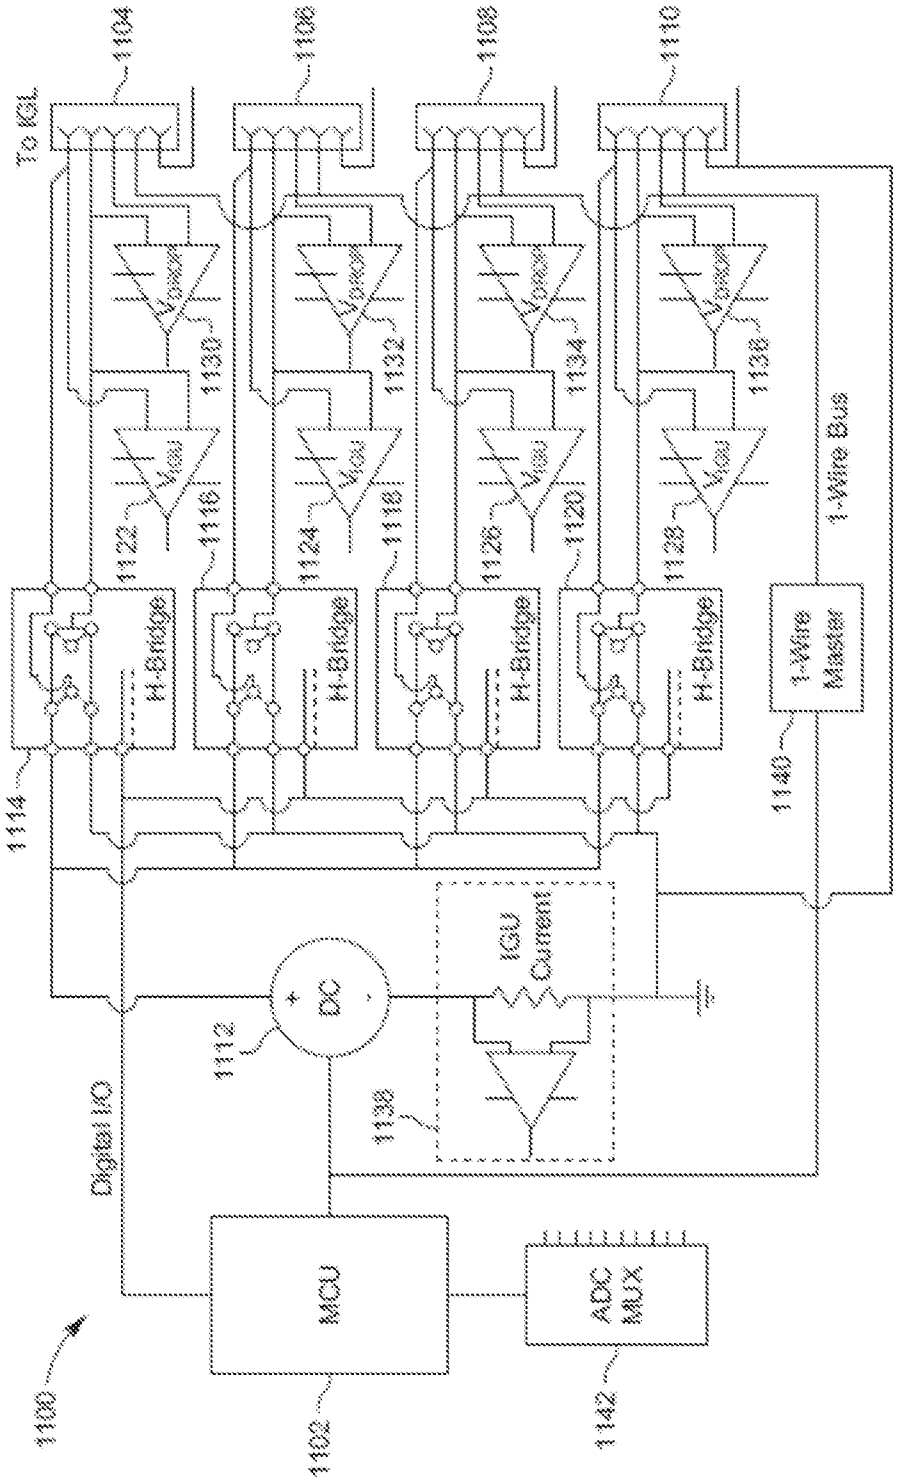

To IGL
1104
1106
1108
1110
1130
1132
1134
1136
1122
1116
1124
1118
1126
1120
1128
IGU
H-Bridge
1-Wire Bus
1-Wire Master
1140
1114
DC
1112
IGU Current
1138
Digital I/O
MCU
ADC MUX
1100
1102
1142

FIGURE 11

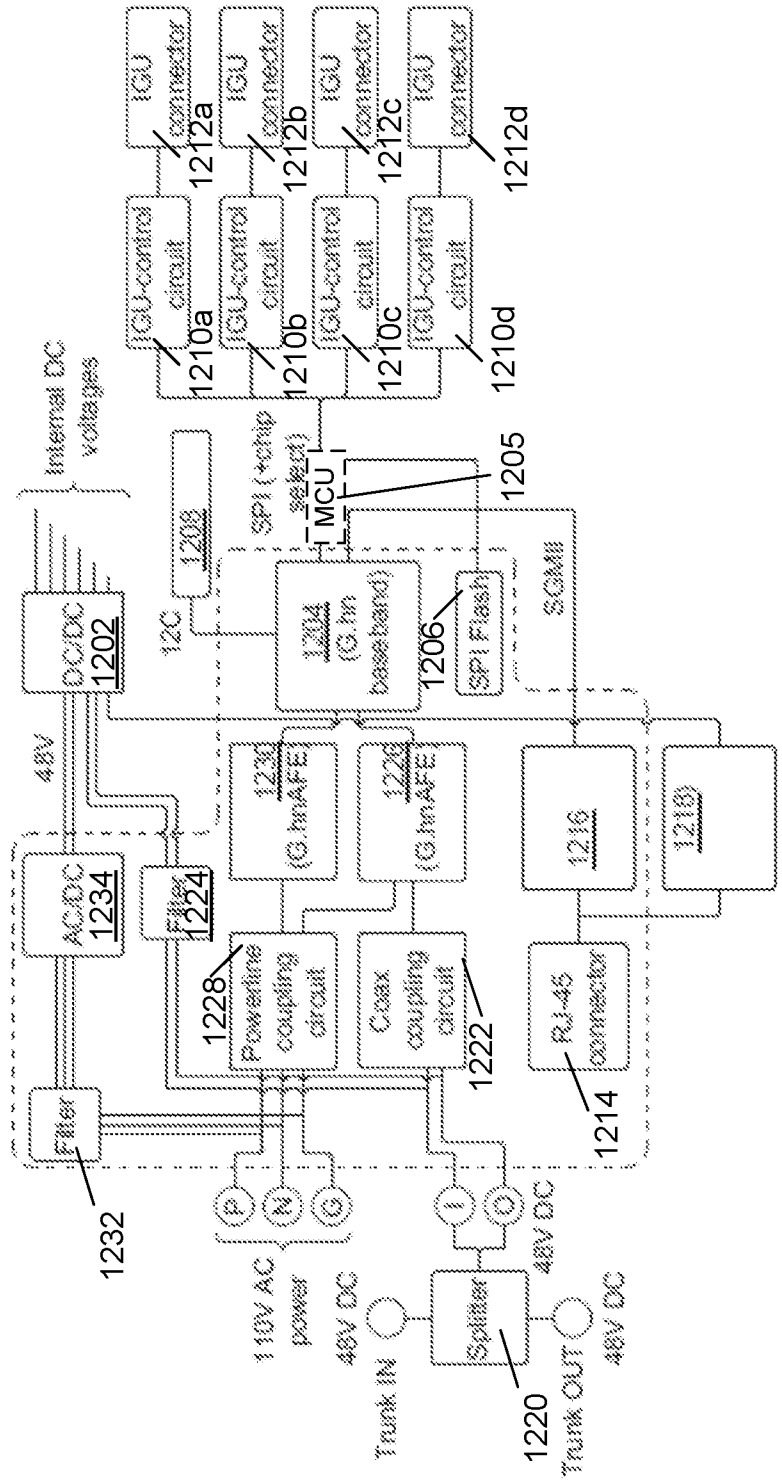

IGU connector
1212a
1212b
1212c
1212d
IGU-control circuit
1210a
1210b
1210c
1210d
Internal DC voltages
DC/DC
1202
1208
12C
SPI (+chip select)
MCU
1205
1204
(G.hn baseband)
1206
SPI Flash
SGMII
48V
AC/DC
1234
Filter
1224
1230
(G.hnAFE)
1220
(G.hnAFE)
1216
1218
1228
Powerline coupling circuit
Coax coupling circuit
1222
RJ-45 connector
1214
Filter
1232
P
N
G
I
O
110V AC power
48V DC
Trunk IN
Splitter
1220
48V DC
Trunk OUT
48V DC

FIGURE 12

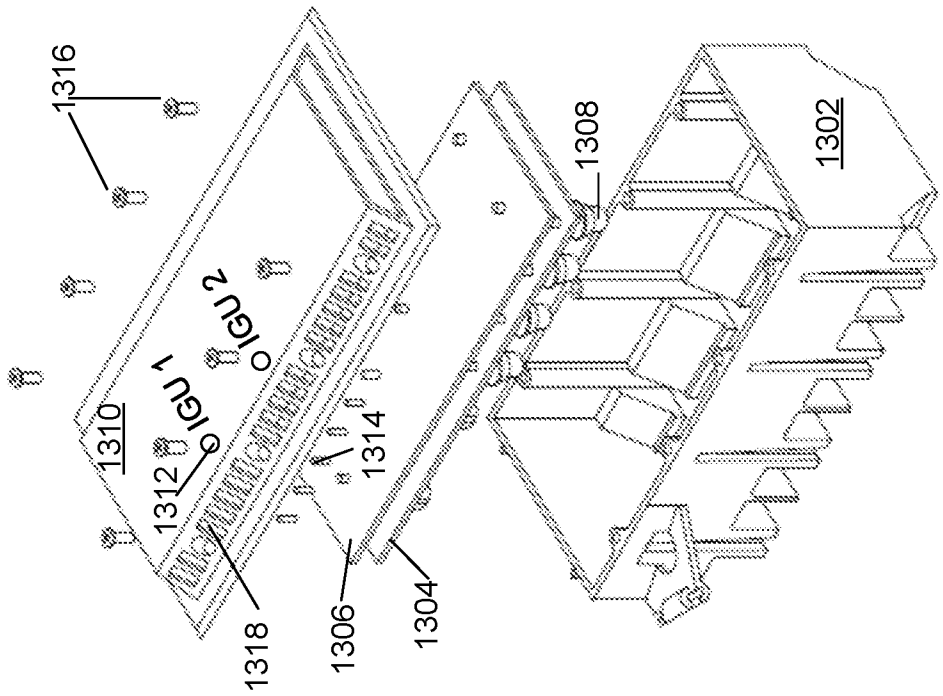


FIGURE 13A

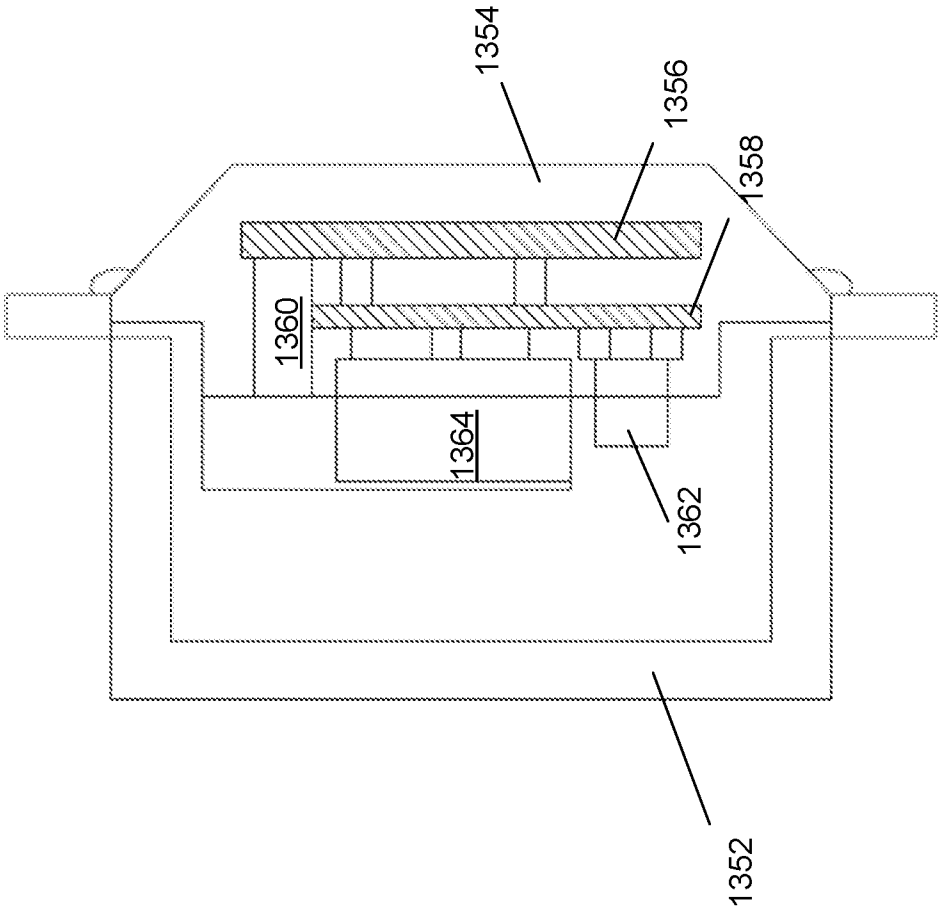


FIGURE 13B

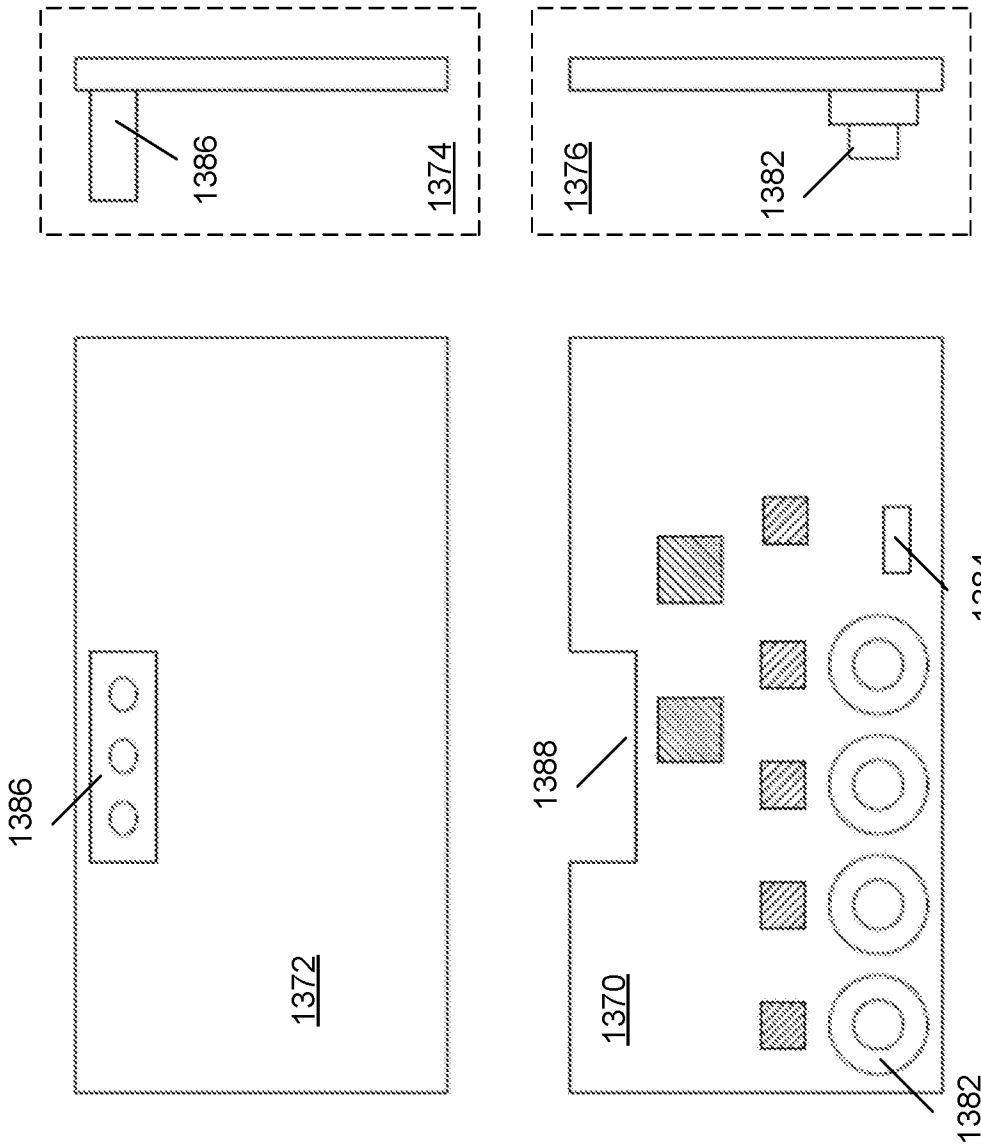


FIGURE 13C

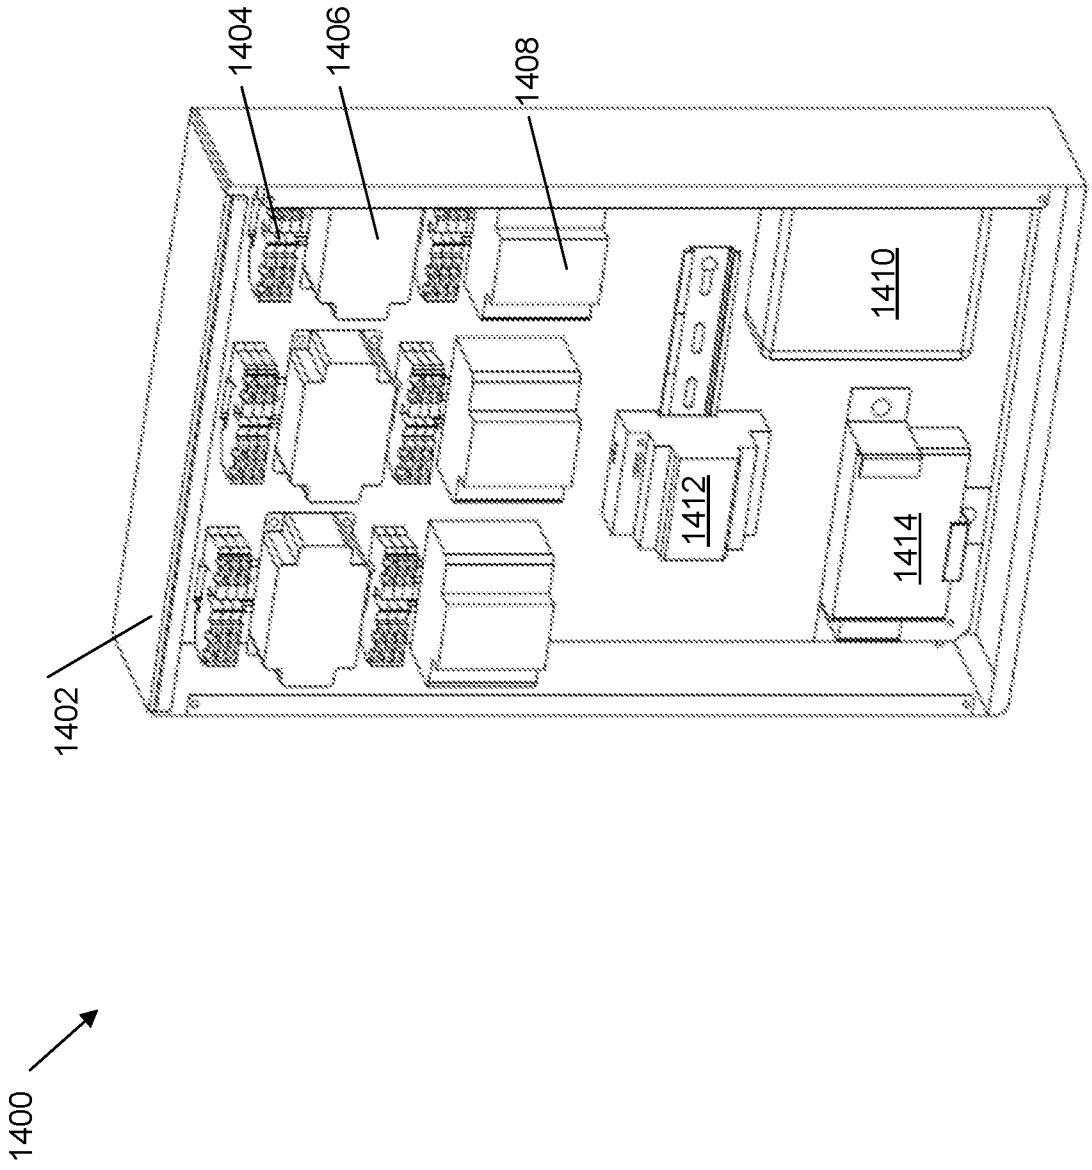


FIGURE 14

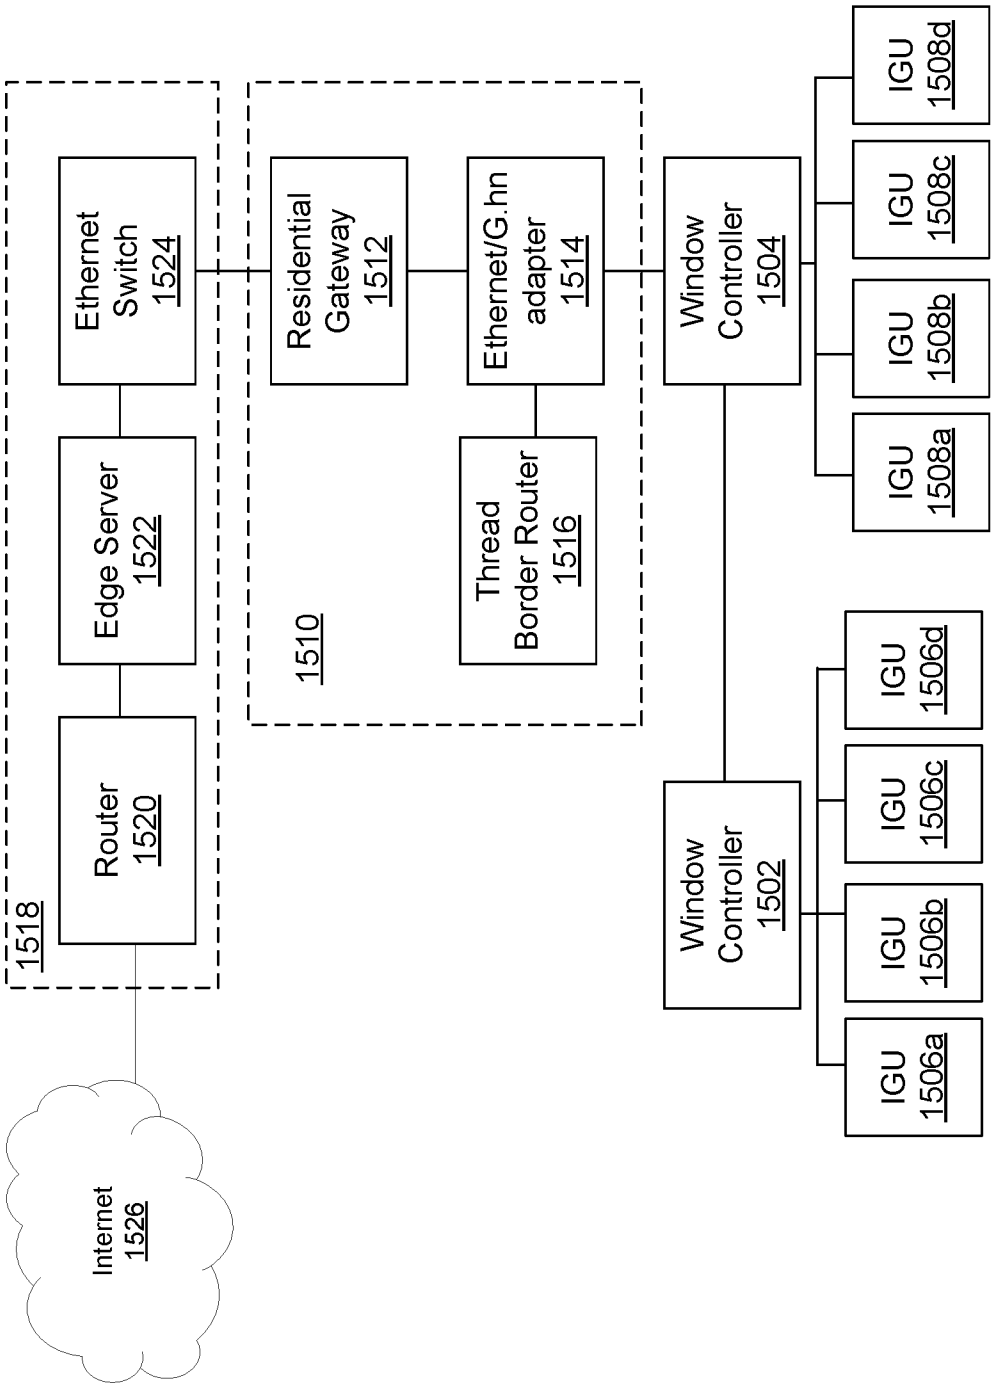

Internet
1526
1518
Router
1520
Edge Server
1522
Ethernet
Switch
1524
1510
Residential
Gateway
1512
Thread
Border Router
1516
Ethernet/G.hn
adapter
1514
Window
Controller
1502
Window
Controller
1504
IGU
1506a
IGU
1506b
IGU
1506c
IGU
1506d
IGU
1508a
IGU
1508b
IGU
1508c
IGU
1508d

FIGURE 15

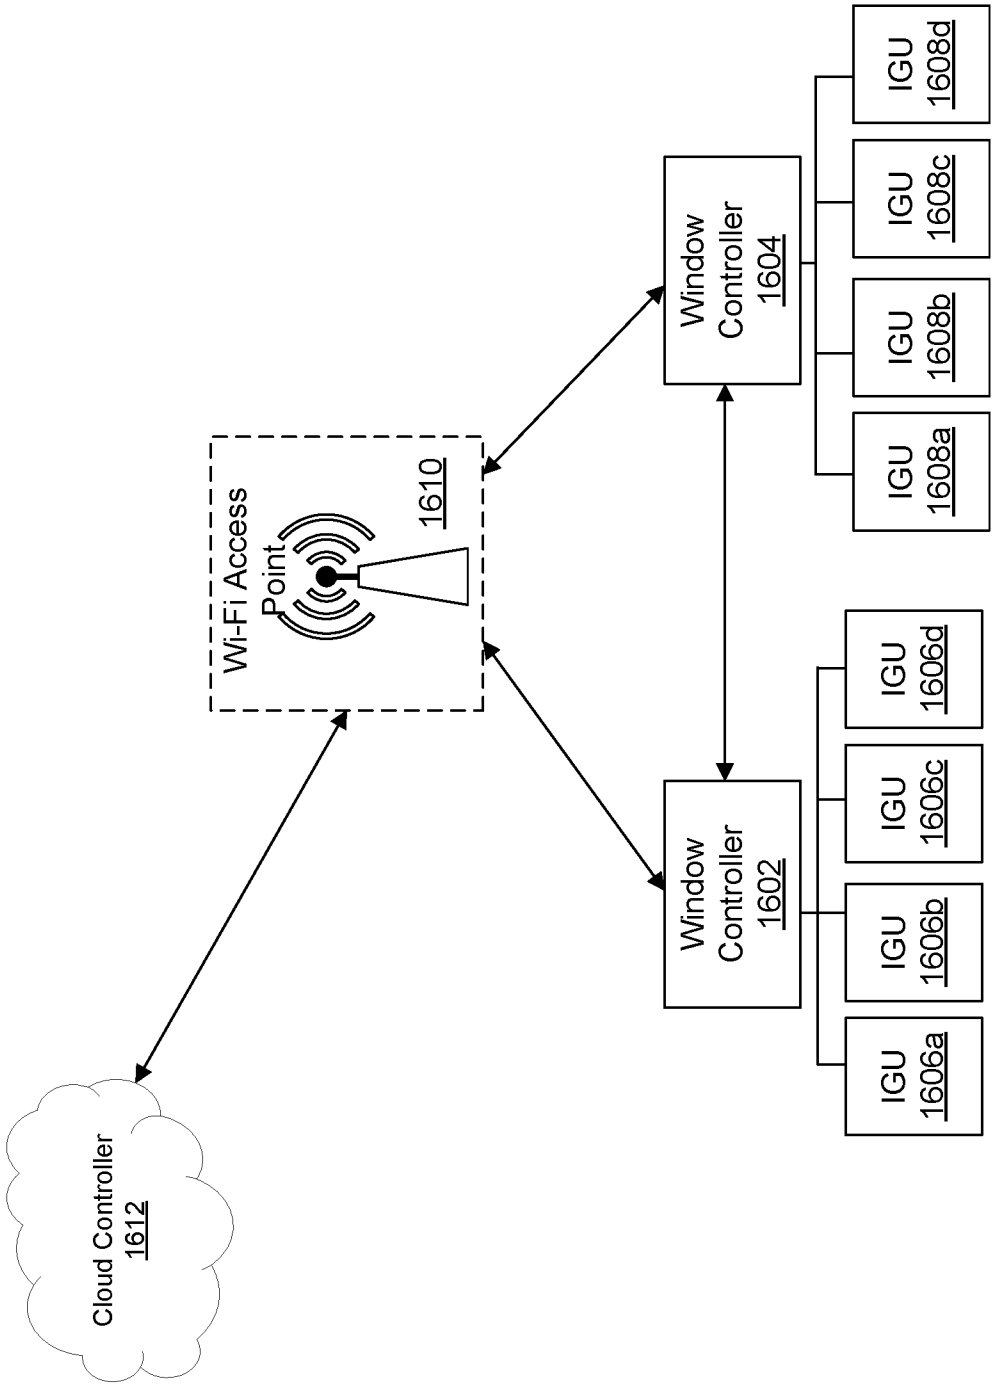

Cloud Controller
1612
Wi-Fi Access Point
1610
Window Controller
1602
Window Controller
1604
IGU
1606a
IGU
1606b
IGU
1606c
IGU
1606d
IGU
1608a
IGU
1608b
IGU
1608c
IGU
1608d

FIGURE 16

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2022/074066

A. CLASSIFICATION OF SUBJECT MATTER

E06B 9/24(2006.01)i; **G05B 15/02**(2006.01)i; **G02F 1/163**(2006.01)i; **H02M 3/158**(2006.01)i; **H04L 67/125**(2022.01)i

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

E06B 9/24(2006.01); E06B 3/67(2006.01); G02F 1/15(2006.01); G09G 3/38(2006.01); H04B 10/25(2013.01); H04W 76/02(2009.01); H05B 37/02(2006.01); H05K 5/02(2006.01)

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

Korean utility models and applications for utility models
Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: window controller system, insulated glass unit, printed circuit board, processing unit, power, data signal, protocol, cable, computing device, enclosure

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
	WO 2021-067505 A1 (VIEW, INC.) 08 April 2021 (2021-04-08)	
X	paragraphs [0096]-[0102], [0128], [0168], [0189]-[0204] and figures 1A-1B, 26, 32	1-4,11-14,16-17,39-49
Y		5-10,21-38
A		15,18-20
	US 2017-0006694 A1 (KORTEK INDUSTRIES PTY. LTD.) 05 January 2017 (2017-01-05)	
Y	paragraphs [0045], [0073] and figure 2	5-10,33-37
	WO 2020-037055 A1 (VIEW, INC.) 20 February 2020 (2020-02-20)	
Y	paragraphs [0128], [0149]-[0150], [0166] and figures 11, 24	21-27
	WO 2018-112241 A1 (VIEW, INC.) 21 June 2018 (2018-06-21)	
X	paragraphs [0016], [0061]-[0069], [0096], [0099] and figures 2, 9A-9D	50-53
Y		28-32,38

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

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"D" document cited by the applicant in the international application
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
18 November 2022	**18 November 2022**

Name and mailing address of the ISA/KR	Authorized officer
Korean Intellectual Property Office **189 Cheongsa-ro, Seo-gu, Daejeon** **35208, Republic of Korea**	**PARK, Tae Wook**
Facsimile No. **+82-42-481-8578**	Telephone No. **+82-42-481-3405**

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2022/074066

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	KR 10-2018-0135741 A (INWOO SYSTEM CO., LTD.) 21 December 2018 (2018-12-21) paragraphs [0060]-[0081] and figures 1-4	1-53
A	CN 207652442 U (ZHUHAI ZHONGTENG TECH. CO., LTD.) 24 July 2018 (2018-07-24) claims 1-4 and figures 1-2	1-53

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2022/074066

Patent document cited in search report	Publication date (day/month/year)	Patent family member(s)	Publication date (day/month/year)
WO 2021-067505 A1	08 April 2021	CN 113940014 A	14 January 2022
		EP 3966963 A2	16 March 2022
		EP 4038449 A1	10 August 2022
		JP 2022-531368 A	06 July 2022
		KR 10-2022-0006601 A	17 January 2022
		KR 10-2022-0065074 A	19 May 2022
		US 11231633 B2	25 January 2022
		US 11346710 B2	31 May 2022
		US 2020-0318426 A1	08 October 2020
		US 2020-0321682 A1	08 October 2020
		US 2020-0387041 A1	10 December 2020
		US 2020-0393733 A1	17 December 2020
		US 2021-0041759 A1	11 February 2021
		US 2021-0063835 A1	04 March 2021
		US 2021-0063836 A1	04 March 2021
		US 2021-0108960 A1	15 April 2021
		US 2021-0119318 A1	22 April 2021
		US 2021-0132458 A1	06 May 2021
		US 2021-0174804 A1	10 June 2021
		US 2021-0191221 A1	24 June 2021
		US 2021-0208467 A1	08 July 2021
		US 2021-0255519 A1	19 August 2021
		US 2021-0294173 A1	23 September 2021
		US 2021-0325754 A1	21 October 2021
		US 2021-0376445 A1	02 December 2021
		US 2022-0021099 A1	20 January 2022
		US 2022-0026267 A1	27 January 2022
		US 2022-0113184 A1	14 April 2022
		WO 2020-227702 A2	12 November 2020
		WO 2020-227702 A3	07 January 2021
		WO 2021-141991 A2	15 July 2021
		WO 2021-141991 A3	02 September 2021
		WO 2021-154915 A1	05 August 2021
		WO 2021-163287 A1	19 August 2021
		WO 2021-163552 A1	19 August 2021
		WO 2021-194944 A1	30 September 2021
		WO 2021-195180 A1	30 September 2021
		WO 2021-211798 A1	21 October 2021
		WO 2021-226182 A1	11 November 2021
		WO 2021-237019 A1	25 November 2021
		WO 2022-026366 A1	03 February 2022
		WO 2022-046541 A1	03 March 2022
		WO 2022-072454 A1	07 April 2022
		WO 2022-072461 A2	07 April 2022
		WO 2022-072461 A3	02 June 2022
		WO 2022-072462 A1	07 April 2022
		WO 2022-093629 A1	05 May 2022
		WO 2022-098630 A1	12 May 2022
US 2017-0006694 A1	05 January 2017	AU 2015-230675 A1	29 September 2016
		CA 2942104 A1	17 September 2015

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2022/074066

Patent document cited in search report			Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
				CN	106465514	A	22 February 2017
				EP	3117691	A1	18 January 2017
				JP	2017-511966	A	27 April 2017
				WO	2015-135033	A1	17 September 2015
WO	2020-037055	A1	20 February 2020	AU	2019-320784	A1	04 March 2021
				CA	3109581	A1	20 February 2020
				CN	112771241	A	07 May 2021
				EP	3837414	A1	23 June 2021
				JP	2021-535971	A	23 December 2021
				KR	10-2021-0058839	A	24 May 2021
				TW	202034051	A	16 September 2020
WO	2018-112241	A1	21 June 2018	AU	2013-249621	A1	30 October 2014
				AU	2013-249621	B2	07 September 2017
				AU	2013-249706	B2	20 October 2016
				AU	2014-219076	A1	10 September 2015
				AU	2014-219076	B2	23 November 2017
				AU	2015-227056	A1	08 September 2016
				AU	2015-227056	B2	01 October 2020
				AU	2015-255913	A1	24 November 2016
				AU	2015-255913	B2	08 October 2020
				AU	2015-324000	A1	06 April 2017
				AU	2015-324000	B2	06 May 2021
				AU	2015-353569	A1	25 May 2017
				AU	2015-353569	B2	17 December 2020
				AU	2015-353606	A1	25 May 2017
				AU	2015-353606	B2	21 May 2020
				AU	2015-360714	A1	29 June 2017
				AU	2016-334038	A1	10 May 2018
				AU	2016-346328	A1	17 May 2018
				AU	2017-200334	A1	02 February 2017
				AU	2017-257789	A1	22 November 2018
				AU	2017-260101	A1	22 November 2018
				AU	2017-260101	B2	22 July 2021
				AU	2017-270472	A1	21 December 2017
				AU	2017-270472	B2	19 December 2019
				AU	2017-363581	A1	20 June 2019
				AU	2017-376447	A1	04 July 2019
				AU	2018-201341	A1	15 March 2018
				AU	2018-201341	B2	05 March 2020
				AU	2018-201341	C1	19 November 2020
				AU	2018-203436	A1	07 June 2018
				AU	2018-260906	A1	29 November 2018
				AU	2019-240134	A1	08 October 2020
				AU	2020-202011	A1	09 April 2020
				AU	2020-202135	A1	16 April 2020
				AU	2020-202135	B2	28 October 2021
				AU	2020-202135	B9	11 November 2021
				AU	2020-220165	A1	10 September 2020
				AU	2020-220165	B2	27 January 2022
				AU	2020-226999	A1	17 September 2020

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2022/074066

Patent document cited in search report	Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
		AU	2020-250299	A1	05 November 2020
		AU	2020-273368	A1	17 December 2020
		AU	2021-200070	A1	18 March 2021
		AU	2021-201145	A1	11 March 2021
		AU	2021-205049	A1	12 August 2021
		AU	2021-215134	A1	02 September 2021
		AU	2021-250838	A1	04 November 2021
		CA	170770	S	15 February 2018
		CA	2870627	A1	24 October 2013
		CA	2870627	C	14 December 2021
		CA	2870673	A1	24 October 2013
		CA	2871047	A1	31 October 2013
		CA	2880920	A1	13 February 2014
		CA	2902106	A1	28 August 2014
		CA	2916862	A1	31 December 2014
		CA	2941526	A1	11 September 2015
		CA	2948668	A1	12 November 2015
		CA	2963096	A1	07 April 2016
		CA	2968665	A1	02 June 2016
		CA	2968840	A1	02 June 2016
		CA	2970300	A1	16 June 2016
		CA	2980477	A1	29 September 2016
		CA	3000852	A1	13 April 2017
		CA	3001233	A1	13 April 2017
		CA	3003639	A1	04 May 2017
		CA	3017138	A1	14 September 2017
		CA	3022396	A1	02 November 2017
		CA	3023072	A1	09 November 2017
		CA	3034630	A1	01 March 2018
		CA	3038974	A1	05 April 2018
		CA	3039342	A1	12 April 2018
		CA	3039606	A1	12 April 2018
		CA	3044823	A1	31 May 2018
		CA	3047093	A1	21 June 2018
		CA	3047110	A1	21 June 2018
		CA	3062815	A1	01 November 2018
		CA	3062817	A1	01 November 2018
		CA	3062818	A1	01 November 2018
		CA	3066285	A1	07 November 2019
		CA	3094479	A1	26 September 2019
		CA	3102449	A1	12 December 2019
		CA	3102958	A1	19 December 2019
		CA	3103480	A1	26 December 2019
		CN	103238107	A	07 August 2013
		CN	103238107	B	07 June 2019
		CN	103261960	A	21 August 2013
		CN	103261960	B	08 August 2017
		CN	103492940	A	01 January 2014
		CN	103492940	B	15 February 2017
		CN	103547965	A	29 January 2014

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2022/074066

Patent document cited in search report	Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
		CN	103547965	B	08 August 2017
		CN	103649826	A	19 March 2014
		CN	103649826	B	01 May 2018
		CN	104114804	A	22 October 2014
		CN	104246594	A	24 December 2014
		CN	104246594	B	24 September 2019
		CN	104321497	A	28 January 2015
		CN	104321497	B	24 August 2016
		CN	104321696	A	28 January 2015
		CN	104321696	B	18 October 2019
		CN	104335595	A	04 February 2015
		CN	104335595	B	18 September 2018
		CN	104364706	A	18 February 2015
		CN	104364706	B	18 August 2017
		CN	104603686	A	06 May 2015
		CN	104603686	B	08 March 2017
		CN	105143586	A	09 December 2015
		CN	105143586	B	16 November 2018
		CN	105431772	A	23 March 2016
		CN	105431772	B	16 April 2019
		CN	106125444	A	16 November 2016
		CN	106125444	B	14 August 2020
		CN	106164973	A	23 November 2016
		CN	106164973	B	01 January 2021
		CN	106462023	A	22 February 2017
		CN	106462023	B	24 December 2019
		CN	106796380	A	31 May 2017
		CN	106796380	B	03 September 2021
		CN	106837094	A	13 June 2017
		CN	106837094	B	25 June 2021
		CN	106896613	A	27 June 2017
		CN	106896613	B	07 August 2020
		CN	107109892	A	29 August 2017
		CN	107109892	B	12 June 2020
		CN	107110704	A	29 August 2017
		CN	107110704	B	24 December 2021
		CN	107111287	A	29 August 2017
		CN	107112620	A	29 August 2017
		CN	107112620	B	31 December 2019
		CN	107272296	A	20 October 2017
		CN	107340664	A	10 November 2017
		CN	107340664	B	22 January 2021
		CN	107390446	A	24 November 2017
		CN	107533267	A	02 January 2018
		CN	108291424	A	17 July 2018
		CN	108291424	B	12 June 2020
		CN	108351251	A	31 July 2018
		CN	108351251	B	09 July 2021
		CN	108352144	A	31 July 2018
		CN	108352144	B	14 December 2021

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2022/074066

Patent document cited in search report	Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
		CN	108388058	A	10 August 2018
		CN	108388058	B	21 May 2021
		CN	108922149	A	30 November 2018
		CN	108922149	B	18 June 2021
		CN	109060126	A	21 December 2018
		CN	109060126	B	03 December 2021
		CN	109313498	A	05 February 2019
		CN	109314307	A	05 February 2019
		CN	109314307	B	15 October 2021
		CN	109653669	A	19 April 2019
		CN	109791338	A	21 May 2019
		CN	109844631	A	04 June 2019
		CN	109863425	A	07 June 2019
		CN	109923482	A	21 June 2019
		CN	110058472	A	26 July 2019
		CN	110100072	A	06 August 2019
		CN	110100072	B	16 November 2021
		CN	110168185	A	23 August 2019
		CN	110187583	A	30 August 2019
		CN	110214293	A	06 September 2019
		CN	110609428	A	24 December 2019
		CN	110709765	A	17 January 2020
		CN	110709784	A	17 January 2020
		CN	110716362	A	21 January 2020
		CN	110720120	A	21 January 2020
		CN	110832828	A	21 February 2020
		CN	111061110	A	24 April 2020
		CN	111106426	A	05 May 2020
		CN	111106426	B	03 December 2021
		CN	111550172	A	18 August 2020
		CN	111550173	A	18 August 2020
		CN	112004984	A	27 November 2020
		CN	112055927	A	08 December 2020
		CN	112236576	A	15 January 2021
		CN	112243522	A	19 January 2021
		CN	112262341	A	22 January 2021
		CN	112313909	A	02 February 2021
		CN	112615135	A	06 April 2021
		CN	112627704	A	09 April 2021
		CN	112731720	A	30 April 2021
		CN	113204147	A	03 August 2021
		CN	113219756	A	06 August 2021
		CN	113219757	A	06 August 2021
		CN	113362585	A	07 September 2021
		CN	113588075	A	02 November 2021
		CN	113608392	A	05 November 2021
		CN	113889744	A	04 January 2022
		CN	113934066	A	14 January 2022
		CN	113960844	A	21 January 2022
		EP	2517332	A2	31 October 2012

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2022/074066

Patent document cited in search report	Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
		EP	2517332	B1	26 September 2018
		EP	2638429	A1	18 September 2013
		EP	2638429	B1	24 February 2021
		EP	2649490	A2	16 October 2013
		EP	2649490	B1	11 July 2018
		EP	2686728	A2	22 January 2014
		EP	2686729	A2	22 January 2014
		EP	2686729	B1	08 May 2019
		EP	2686730	A2	22 January 2014
		EP	2686730	B1	06 May 2020
		EP	2791451	A1	22 October 2014
		EP	2791451	B1	26 September 2018
		EP	2837205	A1	18 February 2015
		EP	2837205	B1	15 February 2017
		EP	2839336	A1	25 February 2015
		EP	2839337	A1	25 February 2015
		EP	2839337	B1	13 December 2017
		EP	2841671	A1	04 March 2015
		EP	2841987	A1	04 March 2015
		EP	2841987	B1	22 January 2020
		EP	2880492	A1	10 June 2015
		EP	2880492	B1	06 October 2021
		EP	2959088	A1	30 December 2015
		EP	2959088	B1	02 January 2019
		EP	3014349	A1	04 May 2016
		EP	3014349	B1	22 August 2018
		EP	3114640	A1	11 January 2017
		EP	3140694	A1	15 March 2017
		EP	3140694	B1	29 July 2020
		EP	3195301	A1	26 July 2017
		EP	3195301	B1	31 March 2021
		EP	3201582	A1	09 August 2017
		EP	3223532	A1	27 September 2017
		EP	3223532	B1	09 October 2019
		EP	3224442	A1	04 October 2017
		EP	3224901	A1	04 October 2017
		EP	3230943	A1	18 October 2017
		EP	3230943	B1	07 July 2021
		EP	3271782	A1	24 January 2018
		EP	3359930	A1	15 August 2018
		EP	3360125	A1	15 August 2018
		EP	3360125	B1	22 September 2021
		EP	3364241	A1	22 August 2018
		EP	3364241	B1	03 June 2020
		EP	3368736	A1	05 September 2018
		EP	3415984	A1	19 December 2018
		EP	3415984	B1	04 August 2021
		EP	3422095	A1	02 January 2019
		EP	3425450	A1	09 January 2019
		EP	3425450	B1	24 November 2021

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2022/074066

Patent document cited in search report	Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
		EP	3426874	A1	16 January 2019
		EP	3434853	A1	30 January 2019
		EP	3444664	A1	20 February 2019
		EP	3449341	A1	06 March 2019
		EP	3453072	A1	13 March 2019
		EP	3470611	A1	17 April 2019
		EP	3470611	B1	01 September 2021
		EP	3492973	A1	05 June 2019
		EP	3500891	A1	26 June 2019
		EP	3519889	A1	07 August 2019
		EP	3520067	A1	07 August 2019
		EP	3522456	A1	07 August 2019
		EP	3522456	B1	05 May 2021
		EP	3523614	A1	14 August 2019
		EP	3545159	A1	02 October 2019
		EP	3555701	A2	23 October 2019
		EP	3568562	A1	20 November 2019
		EP	3587723	A1	01 January 2020
		EP	3594743	A1	15 January 2020
		EP	3611707	A1	19 February 2020
		EP	3616000	A1	04 March 2020
		EP	3616008	A1	04 March 2020
		EP	3616189	A2	04 March 2020
		EP	3751089	A1	16 December 2020
		EP	3766162	A1	20 January 2021
		EP	3768934	A1	27 January 2021
		EP	3768935	A1	27 January 2021
		EP	3788769	A1	10 March 2021
		EP	3803506	A2	14 April 2021
		EP	3803815	A1	14 April 2021
		EP	3811574	A1	28 April 2021
		EP	3848924	A1	14 July 2021
		EP	3859444	A1	04 August 2021
		EP	3869265	A2	25 August 2021
		EP	3869265	A3	10 November 2021
		EP	3876030	A1	08 September 2021
		EP	3926572	A1	22 December 2021
		EP	3929394	A1	29 December 2021
		EP	3929395	A1	29 December 2021
		EP	3933498	A1	05 January 2022
		ES	2625003	T3	18 July 2017
		ES	2718671	T3	03 July 2019
		HK	1202600	A1	02 October 2015
		HK	1205273	A1	11 December 2015
		HK	1206108	A1	31 December 2015
		HK	1206528	A1	08 January 2016
		HK	1206821	A1	15 January 2016
		HK	1207675	A1	05 February 2016
		HK	1214848	A1	05 August 2016
		HK	1220775	A1	12 May 2017

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2022/074066

Patent document cited in search report	Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
		HK	1244592	A1	10 August 2018
		HK	1258581	A1	15 November 2019
		HK	1258923	A1	22 November 2019
		IN	2371KON2014	A	01 May 2015
		JP	2013-515457	A	02 May 2013
		JP	2015-530613	A	15 October 2015
		JP	2016-001994	A	07 January 2016
		JP	2016-516921	A	09 June 2016
		JP	2017-515160	A	08 June 2017
		JP	2018-077498	A	17 May 2018
		JP	2018-507337	A	15 March 2018
		JP	2019-008302	A	17 January 2019
		JP	2019-085871	A	06 June 2019
		JP	2019-502145	A	24 January 2019
		JP	2019-520552	A	18 July 2019
		JP	2019-521411	A	25 July 2019
		JP	2020-098336	A	25 June 2020
		JP	2020-184876	A	12 November 2020
		JP	2020-507697	A	12 March 2020
		JP	2020-514576	A	21 May 2020
		JP	2021-008813	A	28 January 2021
		JP	2021-059969	A	15 April 2021
		JP	2021-177698	A	11 November 2021
		JP	2021-518498	A	02 August 2021
		JP	2021-523400	A	02 September 2021
		JP	2021-528017	A	14 October 2021
		JP	2022-003213	A	11 January 2022
		JP	2022-003392	A	11 January 2022
		JP	5805658	B2	04 November 2015
		JP	6266618	B2	24 January 2018
		JP	6414782	B2	31 October 2018
		JP	6541003	B2	10 July 2019
		JP	6629824	B2	15 January 2020
		JP	6639415	B2	05 February 2020
		JP	6724076	B2	15 July 2020
		JP	6818386	B2	20 January 2021
		JP	6919968	B2	18 August 2021
		JP	6938608	B2	22 September 2021
		JP	6949957	B2	13 October 2021
		KR	10-2015-0003271	A	08 January 2015
		KR	10-2015-0008414	A	22 January 2015
		KR	10-2015-0040985	A	15 April 2015
		KR	10-2015-0133728	A	30 November 2015
		KR	10-2016-0127762	A	04 November 2016
		KR	10-2017-0016336	A	13 February 2017
		KR	10-2017-0065599	A	13 June 2017
		KR	10-2017-0087931	A	31 July 2017
		KR	10-2017-0093890	A	16 August 2017
		KR	10-2017-0095892	A	23 August 2017
		KR	10-2018-0051639	A	16 May 2018

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2022/074066

Patent document cited in search report	Publication date (day/month/year)	Patent family member(s)	Publication date (day/month/year)
		KR 10-2018-0064532 A	14 June 2018
		KR 10-2018-0064548 A	14 June 2018
		KR 10-2019-0004308 A	11 January 2019
		KR 10-2019-0005197 A	15 January 2019
		KR 10-2019-0052140 A	15 May 2019
		KR 10-2019-0077583 A	03 July 2019
		KR 10-2020-0012025 A	04 February 2020
		KR 10-2020-0024338 A	06 March 2020
		KR 10-2020-0060787 A	01 June 2020
		KR 10-2020-0133268 A	26 November 2020
		KR 10-2021-0005681 A	14 January 2021
		KR 10-2021-0014222 A	08 February 2021
		KR 10-2021-0023868 A	04 March 2021
		KR 10-2021-0054592 A	13 May 2021
		KR 10-2021-0112404 A	14 September 2021
		KR 10-2021-0115048 A	24 September 2021
		KR 10-2021-0122310 A	08 October 2021
		KR 10-2071736 B1	30 January 2020
		KR 10-2084811 B1	04 March 2020
		KR 10-2117731 B1	02 June 2020
		KR 10-2212618 B1	05 February 2021
		KR 10-2249718 B1	11 May 2021
		KR 10-2306274 B1	30 September 2021
		KR 10-2336168 B1	07 December 2021
		RU 2014144632 A	10 June 2016
		RU 2014145565 A	10 June 2016
		RU 2014145822 A	10 June 2016
		RU 2014147152 A	20 June 2016
		RU 2015107563 A	27 September 2016
		RU 2015139884 A	27 March 2017
		RU 2016102399 A	03 August 2017
		RU 2016139012 A	25 April 2018
		RU 2016139012 A3	24 September 2018
		RU 2016148196 A	13 June 2018
		RU 2016148196 A3	10 December 2018
		RU 2017114996 A	02 November 2018
		RU 2017114996 A3	01 April 2019
		RU 2017120790 A	15 November 2018
		RU 2017123902 A	10 January 2019
		RU 2017123902 A3	17 July 2019
		RU 2017140180 A	12 February 2019
		RU 2018103078 A	22 February 2019
		RU 2018114264 A	04 March 2019
		RU 2018117871 A	02 November 2018
		RU 2018119036 A	07 November 2018
		RU 2018119490 A	02 December 2019
		RU 2018119490 A3	21 May 2020
		RU 2018121503 A	05 March 2019
		RU 2019109013 A	06 May 2019
		RU 2019112834 A	06 June 2019

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2022/074066

Patent document cited in search report	Publication date (day/month/year)	Patent family member(s)			Publication date (day/month/year)
		RU	2623920	C2	29 June 2017
		RU	2636811	C2	01 December 2017
		RU	2644085	C2	07 February 2018
		RU	2656013	C2	30 May 2018
		RU	2657684	C2	14 June 2018
		RU	2658096	C2	19 June 2018
		RU	2660395	C2	06 July 2018
		RU	2684751	C2	12 April 2019
		RU	2688844	C2	22 May 2019
		RU	2711515	C2	17 January 2020
		SG	10201608572	A	29 December 2016
		SG	10201807277	A	27 September 2018
		SG	11201406676	A	30 March 2015
		SG	11201406722	A	27 November 2014
		SG	11201606966	A	29 September 2016
		TW	201231787	A	01 August 2012
		TW	201235757	A	01 September 2012
		TW	201243470	A	01 November 2012
		TW	201248286	A	01 December 2012
		TW	201307975	A	16 February 2013
		TW	201341927	A	16 October 2013
		TW	201348828	A	01 December 2013
		TW	I528096	B	01 April 2016
		TW	I551933	B	01 October 2016
		TW	I559065	B	21 November 2016
		TW	I563160	B	21 December 2016
		TW	I575294	B	21 March 2017
		TW	I596413	B	21 August 2017
		TW	I596414	B	21 August 2017
		TW	I600956	B	01 October 2017
		TW	I600957	B	01 October 2017
		TW	I607269	B	01 December 2017
		TW	I609224	B	21 December 2017
		TW	I615665	B	21 February 2018
		TW	I616705	B	01 March 2018
		TW	I621766	B	21 April 2018
		TW	I626364	B	11 June 2018
		TW	I631403	B	01 August 2018
		TW	I636311	B	21 September 2018
		TW	I655489	B	01 April 2019
		TW	I656393	B	11 April 2019
		TW	I660227	B	21 May 2019
		TW	I670554	B	01 September 2019
		TW	I681294	B	01 January 2020
		TW	I683954	B	01 February 2020
		TW	I685707	B	21 February 2020
		TW	I686656	B	01 March 2020
		TW	I687738	B	11 March 2020
		TW	I693334	B	11 May 2020
		TW	I702754	B	21 August 2020

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.
PCT/US2022/074066

Patent document cited in search report	Publication date (day/month/year)	Patent family member(s)	Publication date (day/month/year)
		TW I716664 B	21 January 2021
		TW I725495 B	21 April 2021
		TW I725999 B	01 May 2021
		TW I727931 B	21 May 2021
		TW I729293 B	01 June 2021
		TW I731881 B	01 July 2021
		TW I734365 B	21 July 2021
		TW I737796 B	01 September 2021
		TW I738174 B	01 September 2021
		TW I739224 B	11 September 2021
KR 10-2018-0135741 A	21 December 2018	KR 10-1951391 B1	29 April 2019
CN 207652442 U	24 July 2018	None	