 By Joshua Israelsohn, Technical Editor

HOT-SWAP DESIGNS WERE ONCE RESERVED FOR MEDICAL-SUPPORT SYSTEMS, TELCOMM, OR OTHER "HIGH-RELIABILITY" INSTALLATIONS. NOW EVERYBODY WANTS TO SWAP WITH THE LIGHTS ON.

Hot-swapping SIGNALS

It wasn't long ago that hot-swap capabilities were limited to systems demanding the greatest reliability and uptime. Now users expect equipment to tolerate connecting or replacing modules or peripherals without power cycling the host system. Hot-swappable designs, however, are not a homogeneous category. They involve power-management products as well as signal-handling ICs and have architectural implications at the component, bus, and system levels.

WHY NOT SLAM AND SCRAM?

You needn't be an IC designer to design hot-swap tolerant circuits, but familiarity with the I/O structure of candidate logic families goes a long way toward ensuring that your product's hot-swap robustness is a matter of design, not happenstance. A generalized depiction of a CMOS-logic IC includes clamp diodes between I/O pins and the supply rails and is representative of devices from many logic families fabricated on common IC processes (**Figure 1** and **Reference 1**). Diodes D_1 and D_2 provide overshoot and undershoot protection, respectively. They also enhance the chip's ESD robustness at the inputs, protecting the relatively fragile thin-oxide regions that form the MOS devices' gates. IC designers may include explicit output-clamp devices in the chip layout, but the outputs may include parasitic diodes as well.

If, for example, this type of I/O structure is on the module side of the module/host interface, then

large currents can flow through the I/O pins when you plug the device into a live bus (**Figure 2**). If either the logic input or output is connected to a bus trace that is more than a diode drop above ground, the bus signal will clamp to the logic block's V_{DD} rail, which starts at ground potential. The bus signals charge C_1 and all the other bypass capacitors on the plane through D_1 or D_3 until the module's local power-management circuits have time to turn on and take over the local supply rail. The combined dynamic output impedance of the host's bus driver, its power supply, and trace impedances limit the initial current through the bus trace. If sufficient current flows, hot-plugging a module can damage the I/O device, the host system, or both. Practices you normally associate with conservative design—a large bypass capacitance on the I/O device's V_{DD} pin, planar power and ground distribution, and high-speed bus layout—exacerbate the problem.

Even if you limit the I/O currents to nondestructive levels, these currents can still disrupt the host system by interrupting data transactions on shared buses and by inducing transients on the power supply. Module-side logic outputs may load data and control buses until their power supply reaches their minimum operating voltage, increasing the importance of designing swap-safe circuits.

SEEKING A LEVEL

Hot-swap components control power-supply and signal-interface behaviors when users insert or remove subsystem components. At the first, most basic level of hot-swap protection logic, IC designers sever the explicit current paths between I/O pins and the supply by removing the upper clamp diodes. Power-off disable circuits cause unpowered logic outputs to exhibit a high output impedance and block the paths formed by the upper parasitic diodes. Bipolar and BiMOS logic readily dispense with the clamp diodes, and modifications to CMOS gates result in structures meeting these criteria. Several commercially available logic families meet level-1 isolation requirements for hot swap, including ALS, AVC, FAST, and LVC available from Fairchild Semiconductor, Philips, and Texas Instruments, to name a few (**references 1**, **2**, and **3**).

AT A GLANCE

- Hot-swappable interfaces are a simple way of ensuring high reliability in modular designs.
- Hot-swap logic is not homogenous. You need to choose a specific level of swap capability and compatible component technologies.
- Many hot-swap strategies depend on controlled connection sequences, but these may be difficult to arrange with pin-and-socket interfaces.
- Bus switches are simple devices that can safely connect otherwise nonswappable logic. Their low impedances make them attractive for high-speed applications.

Although these logic families protect themselves and the system from destructive currents and eliminate bus loading while unpowered, bus contentions can arise while the power supply ramps. Although you can safely plug and unplug modules of this type while the system power is on, bus activity should be suspended during those events.

A second level of hot-swap isolation maintains a high-impedance state on logic outputs while the supplies ramp either up or down, further reducing the chance of modules loading the bus immediately following insertion. Level-2 isolated logic families include ABT, ALVT, LCX, LVT, GTL, and GTLP. An ABT logic cell serves as a good example of the output-control circuits for level-2 isolation (**Figure 3**). The collector of Q_1 follows the module supply voltage during ramp-up—a logic high from the NOR gate's perspective. The NOR gate asserts the output buffer's tristate control signal irrespective of the signal on the cell's $\overline{OE}$ (output enable) pin. The divider formed by R_1, R_2, D_1, and D_2 determine a threshold voltage. When the supply rises through the threshold voltage, Q_1 turns on, pulling the NOR input low. At this point, the $\overline{OE}$ signal determines the NOR gate's output. During a power-down sequence, the process is reversed, allowing the $\overline{OE}$ pin to control the cell until the supply falls below the threshold.

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, go to our information-request page at www.ednmag.com/info. When you contact any of the following manufacturers directly, please let them know you read about their products in *EDN*.

Fairchild Semiconductor
www.fairchildsemi.com
Circle No. 374

Integrated Device Technology
www.idt.com
Circle No. 375

National Semiconductor
www.national.com
Circle No. 376

Pericom
www.pericom.com
Circle No. 377

Philips Semiconductor
www.philips.com
Circle No. 378

Texas Instruments
www.ti.com
Circle No. 379

SUPER INFO NUMBER
For more information on the products available from all of the vendors listed in this box, Enter No. 380 at www.ednmag.com/info.

ACRONYMS

ABT: advanced BiCMOS technology
ALS: advanced low-power Schottky logic
ALVT: advanced low-voltage technology
AVC: advanced very-low-voltage CMOS
BTL: backplane transceiver logic
ESD: electrostatic discharge
ETL: enhanced transceiver logic
FAST: Fairchild advanced-Schottky TTL
FCT: fast CMOS, TTL-compatible
GTL: Gunning transceiver logic
GTLP: Gunning transceiver logic plus
LCX: low-voltage CMOS-logic product family with translation
LVC: low-voltage CMOS
LVDS: low voltage differential signal
LVT: low-voltage BiCMOS technology
PCI: Peripheral Component Interconnect
TTL: transistor-to-transistor logic

At this point, Q_1 shuts off, and the cell's output enters a tristate condition.

With the upper clamps gone and level-2 isolation managing the cell's output impedance during supply transitions, you'd think you could declare victory and move on. Unfortunately, even the stray capacitances of inactive circuits in a high-impedance state combined with the module's trace and connector strays can transiently load the active bus. A simple model of the event includes only the source impedances and stray load capacitances (**Figure 4a**). This model is suitable for physically small systems operating at moderate speeds. As system size or speed increases, trace inductances become non-negligible elements in the model (**Figure 4b**).

Figure 1

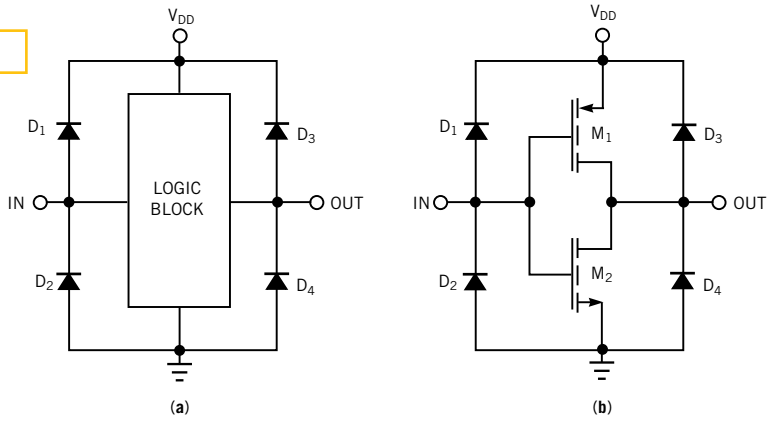

Common CMOS logic families incorporate I/O clamp diodes (a). D_3 and D_4 are often parasitic devices (b).

Referring to the simple model of **Figure 4a**, the source voltage, V_I, corresponds to the voltage on the bus immediately prior to a module insertion. The bus impedance, R_B, is the combination of the bus driver's dynamic source impedance, that of its power supply, the effect of other static bus loads, and wiring resistances. V_B corresponds to the voltage that receivers monitoring the bus observe, and C_M represents the net additional stray capacitance presented to the bus as a result of inserting a module. The disturbance observed by devices monitoring the bus is a step and exponential recovery given by:

$$V_B(t) = V_I - (V_I - V_C(0)) \bullet e^{\frac{-t}{R_B C_B}},$$

where $V_C(0)$ is the voltage on the module's strays immediately before making contact with the bus. The glitch width within a receiver's threshold is:

$$t = -LN\left(\frac{V_I - V_{TH}}{V_I - V_C(0)}\right) R_B C_M,$$

where V_{TH} is the receiver's threshold voltage (**Reference 2**). Note that these equations do not take into account nonidealities during module insertion, such as contact noise.

Figure 2

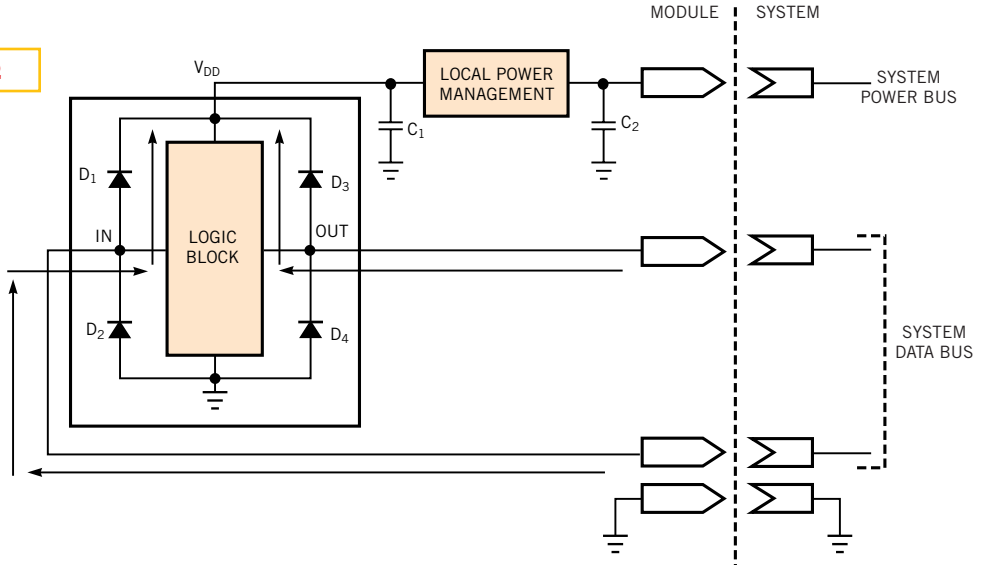

Destructive currents can flow through the clamp diodes if the bus is allowed to drive the module's supply rail during insertion.

The third and highest level of hot-swap-logic isolation includes the second-level features and adds pre-bias circuits to charge the I/O pins' stray capacitances, which reduces contact glitching. Logic families supporting level-3 isolation include BTL, ETL, GTLP with precharge, and LVDS. The pre-bias, however, requires a further complication to the contact sequence during insertion. When mod-

Figure 3

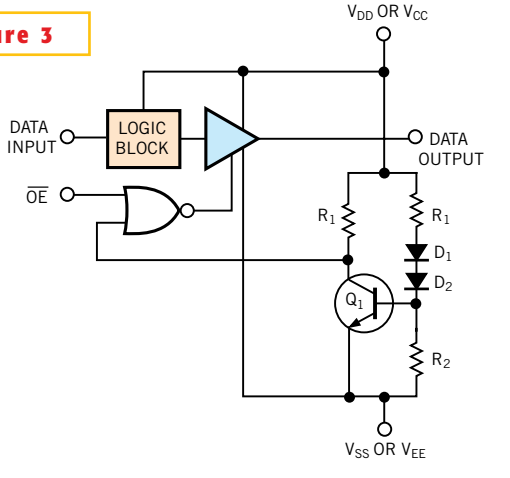

An on-chip supply monitor keeps an ABT output in a tristate condition during power-up and power-down intervals.

ules contact the host through edge connectors, the staggered lengths of the contact fingers control the contact sequence. Ground connections, which the modules must establish first, occupy the longest fingers, followed by output-enable signals derived from the host, supply, and logic-signal pins. Pre-bias adds yet another signal class to the list. To be effective, however, the prebias supply must make contact after the ground and sufficiently before the logic signals to do its job.

Figure 4

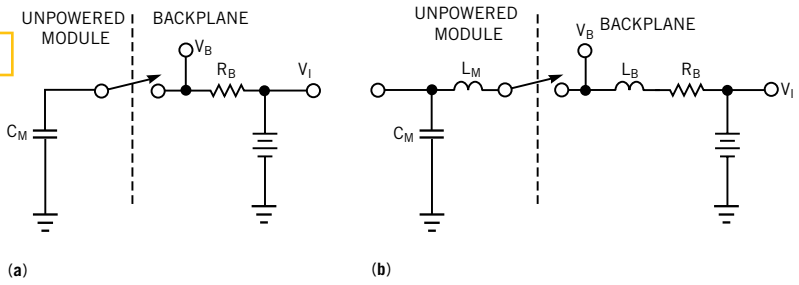

A simple insertion model (a) reveals the magnitude and width of bus disturbances resulting from hot swapping. Include inductive terms when modeling large systems or those operating with high bus speeds (b).

ISSUES OF SAFE SWAPPING

System designers express several concerns about this arrangement: With four of five staggered finger lengths, the differences between contact times become fairly small, and requirements for mechanical precision in the mating connector rise. Large insertion alignment and speed variations from operator to operator and from instance to instance with a given operator reduce the efficacy of staggered finger lengths as determiners of the contact-delay times between signal classes. Spring-loaded insertion mechanisms reduce these variations at some cost of mechanical complexity, but even they are subject to non-negligible variations.

Figure 5

Although it is inexpensive to stagger edge fingers on a pc board, buses that make use of pin-and-socket connectors have no inexpensive equivalents. The most economical option is to choose pin-and-socket connectors that include groundable guide pins that you can use to ensure that the ground systems make the first contact. However, such connectors only provide a two-step sequence of ground and all other signals.

For designers who can use staggered edge fingers, connection-sequence recommendations vary among manufacturers and technologies. For example, Philips recommends, after establishing a ground connection, a sequence of $\overline{OE}$, then supply, with data and other bus signals last (**Figure 5a**). Pericom's recommended sequence for its FCT logic groups all control signals with the $\overline{OE}$, followed by signal and supply, which are connected at the same time (**Figure 5b** and **Reference 4**). Signals mirrored at both ends of the finger columns in the **figure** compensate for board misalignment during insertion, a commonly recommended practice.

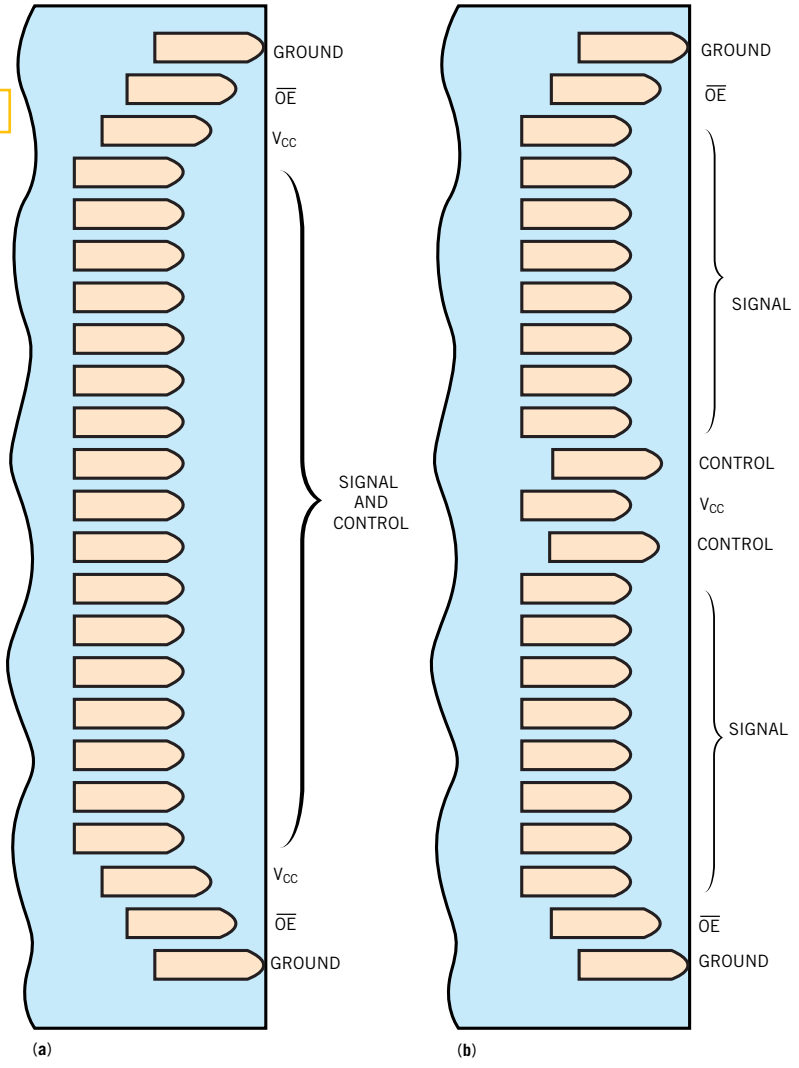

Connection sequences vary by logic manufacturer and technology as shown by this comparison based on Philips (a) and Pericom (b) documentation.

Finally, using specialized logic families between internal functions and module

I/O pins can complicate the design and increase inventory costs in manufacturing. You can alleviate this concern by limiting the functions at the module interface to a short list including transceivers, multiplexers, and perhaps a few other simple functions. But as bus speeds increase and available signal-processing technologies within the module narrow, acceptable hot-swappable logic choices shrink with increasing signal bandwidth and skew requirements.

SWITCHING BUSES MID-DATA STREAM

Separating the logic-isolation function from the logic is one way to untangle issues of hot-swap performance from issues of signal-processing performance. Bus switches allow you to select your logic based on its active performance and still guarantee the isolation you need for hot swapping. Bus switches use a set of NMOS devices designed for low R_{ON} wired as pass-throughs, one in series with each signal (**Figure 6a**). Because the signal looks through the on resistance of a symmetrical MOS structure, bus switches are natively bidirectional without a direction control.

Unlike transmission gates, another natively bidirectional switch formed of a parallel pair of NMOS and PMOS devices (**Figure 6b**), single-FET bus switches have a rising resistance characteristic for signals approaching $V_G - V_{TH}$ due to channel pinch-off (**Reference 5**). For a given geometry, the single-NMOS type has the advantages of presenting a lower stray capacitance to the bus, injecting less charge at turn off.

NMOS bus switches can run into difficulty with bus-undershoot conditions wherein a bus signal dips below ground. Improper termination or poor layout can result in bus undershoot, but some buses, such as the PCI, are designed to operate unterminated and make use of reflected signal energy to get greater bus speed with less power (**Reference 6**). If the bus signal drops a V_{TH} below the substrate, which is held at ground, the switch will turn on. Adding insult to injury, an NMOS bus switch built on a P-type substrate contains a parasitic npn device, which though normally biased off, turns on if you allow the source contact to fall a V_{BE} below the substrate potential (**Figure 7**). The NMOS switch itself, as already mentioned, is comparatively resistive near its threshold. The parasitic npn, however, has a significant β—on the order of 10—and turns on with little substrate current.

A number of circuits guard against bus undershoots. The two most common use a charge pump and an active clamp. The charge-pump arrangement forces the substrate and off-state gate voltage below ground (**Reference 7**). The Pericom PI5C3xxxC family makes use of this type of undershoot protection. The active clamp, typified by TI's CBTK family, uses a bias generator to set a voltage slightly above ground. An active clamp referred to this voltage readily turns on during signal undershoots, preventing either the NMOS device or the parasitic npn from turning on. Both options take additional die area to implement, and both can amortize most of that extra die area over a number of bus-switch channels sharing a die. They both result in fast response times and comparably low node capacitances and leakage currents.

Figure 6

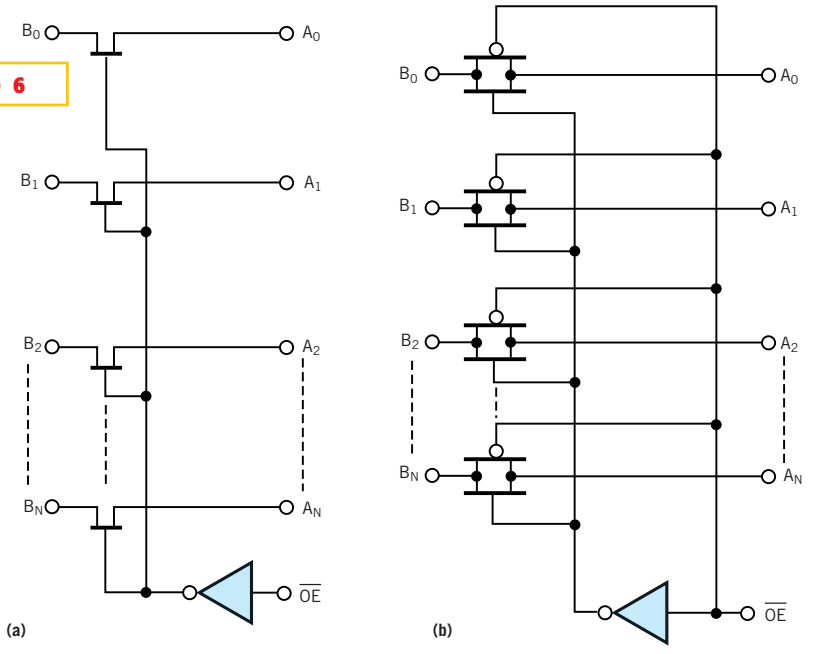

NMOS bus switches (a) can simplify hot-swap designs. An alternative structure uses transmission gates (b), which exhibit more consistent R_{ON} over the signal span but may be less tolerant of overvoltages.

Figure 7

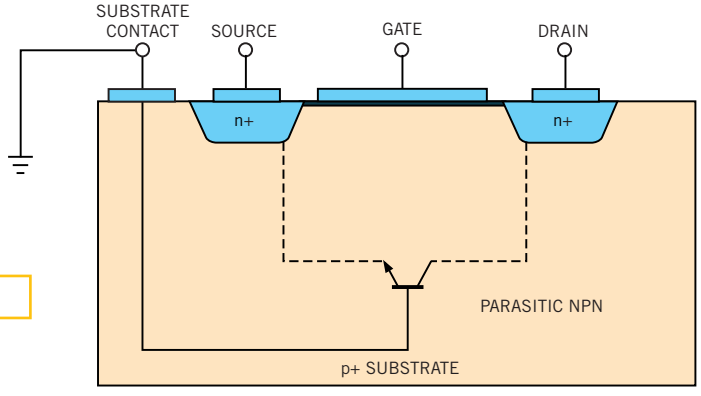

Undershoot on a bus switch can not only turn on the NMOS device but can also activate a parasitic npn.

Like some logic families, bus switches are also available with pre-bias circuitry. A bias supply connects to the module's interface circuitry whenever the bus switch is open and disconnects when the bus switch is closed, minimizing the glitch energy on the bus as circuits are attached. TI's CBTLV series has examples of bus switches with pre-bias circuitry.

Due to the enormous breadth of applications for hot-swappable logic and bus switches and the large number of nongeneric forms they take, it is difficult to call out and compare particular models without establishing application-specific constraints, complete with signal amplitudes, rise and fall times, and bus-load specifications. The next step to apply this information to your application is to consult the literature available from the component vendors listed in the **sidebar**, "For more information."□

References

1. "Live insertion," Application note SDYA012, Texas Instruments, October 1996.

2. Magdaluyo, Mike, "Hot-swap your way to high availability," *EDN*, March 30, 2000.

3. "Live insertion," technical memo, Fairchild Semiconductor, May 2001.

4. "Pericom FCT logic for hot-plug applications," application note 3, Pericom, December 1998.

5. "Zero delay bus switches," application note 1, Pericom, December 1998.

6. Sultana, Nadira and Chris Graves, "Texas Instruments solution for undershoot protection for bus switches," Application report, Texas Instruments, April 2000.

7. Ma, Jimmy, "Undershoot protection when using bus switches," Application note 30, Pericom, May 2001.

You can reach Technical Editor Joshua Israelsohn at 1-617-558-4427, fax 1-617-558-4470, e-mail jisraelsohn@cahners.com.

Author's biography

Never one to power down, Joshua Israelsohn practices safe hot swapping wearing his trusty and fashionable wrist strap. He does, however, wait for the #502 to stop before jumping on or off the bus.