



Considerations for the RS-485

Despite its widespread use, RS-485 is not as well understood as it should be. However, if you invest a little time on familiarizing yourself with the bus and pay attention to 10 aspects of your application, you'll find that designing rock-solid implementations is easy.

Recommended Standard 485 (RS-485) has become the industry's workhorse interface for multipoint, differential data transmission. RS-485 is unique in allowing multiple nodes to communicate bi-directionally over a single twisted pair. No other standard combines this capability with equivalent noise rejection, data rate, cable length, and general robustness. For these reasons, a variety of applications use RS-485 for data transmission. The list includes automotive radios, hard-disk drives, LANs, cellular base stations, industrial programmable logic controllers (PLCs), and even slot machines. The standard's widespread acceptance also results from its generic approach, which deals only with the interface's electrical parameters. RS-485 does not specify a connector, cable, or protocol. Higher level standards, such as the ANSI's SCSI standards and the Society of Automotive Engineers' (SAE's) J1708 automotive-communication standard, govern these parameters and reference RS-485 for the electrical specifications.

Although RS-485 is extremely popular, many system designers must learn how to address its interface issues. You should review 10 areas before you design an RS-485 interface into a product. Understanding the issues during system design can lead to a trouble-free application and can reduce time to market.

RS-485 addresses a need beyond the scope of RS-422, which covers buses with a single driver and multiple receivers. RS-485 provides a low-cost, bidirectional, multipoint interface that supports high noise rejection, fast data rates, long cable, and a wide common mode range. The standard specifies the electrical characteristics of drivers and receivers for differential multipoint data transmission but does not specify the protocol, encoding, connector mechanical characteristics, or pin-out. RS-485 networks include many systems that the general public uses daily. These applications appear wherever a need exists for simple, economical communication among multiple nodes. Examples are gas-station pumps, traffic and railroad signals, point-of-sale equipment, and aircraft passenger seats.

The Electronic Industries Association (EIA) Technical Recommendation Committee, TR30, made RS-485 a standard in 1983. The Tele-communications Industry Association (TIA) is now responsible for revisions. RS-485 is currently being revised. After successful balloting, the revised standard will become "ANSI TIA/EIA-485-A." The 10 considerations that you should review early in a system design are:

- Mode and nodes,
- Configurations,
- Interconnect media,
- Data rate vs cable length,




- Termination and stubs,
- Unique differential and RS-485 parameters,
- Grounding and shielding,
- Contention protection,
- Special-function transceivers, and
- Fail-safe biasing.

Mode and Nodes

In its simplest form, RS-485 is a bidirectional half-duplex bus comprising a transceiver (driver and receiver) located at each end of a twisted-pair cable. Data can flow in either direction but can flow only in one direction at a time. A full-duplex bus, on the other hand, supports simultaneous data flow in both directions. RS-485 is mistakenly thought to be a full-duplex bus because it supports bidirectional data transfer. Simultaneous bidirectional transfers require not one but two data pairs, however. RS-485 allows for connection of up to 32 unit loads (ULs) to the bus. The 32 ULs can include many devices but commonly comprise 32 transceivers. *Figure 1* illustrates a multipoint bus. In this application, three transceivers—two receivers receivers and one driver—connect to the twisted pair. You must observe the 32-UL limitation, because the loads appear in parallel with each other and add to the load that the termination resistors present to the driver. Exceeding 32-UL loads excessively limits the drivers and attenuates the differential signal, thus reducing the differential noise margin.


RS-485 drivers are usually called “60 mA drivers.” The name relates to the allowable loading. Developing 1.5V across the 60 Ω termination load (120 Ω at each end of the bus) requires 25 mA. The worst-case input current of a UL is 1 mA (at extreme common mode, explained later). *Figure 2* shows the loading curve of a full UL. The worst-case UL input resistance is 10.56 k Ω , although a frequently quoted incorrect value is 12 k Ω . Thus, 32 ULs require 32 mA drive capability. Adding this current to the 25 mA for the terminations yields 57 mA, which rounds up to an even 60 mA. A driver that cannot supply the full 60 mA violates the standard and reduces the bus’s performance. The resulting problems include reduced noise margin, reduction in the number of unit loads or allowable cable length, and limited common-mode voltage tolerance.

Designers frequently ask, “What is the maximum number of transceivers the bus allows?” The standard does not specify a maximum number of transceivers, but it does specify a maximum of 32 ULs. If a transceiver imposes one unit load, the maximum number of transceivers is also 32. You can now obtain transceivers with 1/2- and 1/4-UL ratings, which allow 64 and 128 transceivers. However, these fractional-UL devices, with their high-impedance input stages, typically operate much more slowly than do single-UL devices. The lower speed is acceptable for buses operating in the low hundreds of kilobits per second, but it may not be acceptable for a 10 Mbps bus. A solution exists for high-speed buses: You can use RS-485 repeaters to connect multiple buses end to end. In this setup, each bus must have no more than 32 loads. Directional control of the repeaters is complex, but hardware can handle it (Reference 1). Therefore, a conservative estimate is that, without using special transceivers, a bus can include 32 transceivers.



An RS-485 bus supports two-way data transfer over a single pair of wires. A typical bus includes multiple nodes. Each transceiver includes a differential driver, D, and a differential receiver, R. The stub length is I . The bus is terminated only at the ends—not at each node.

FIGURE 1.



The loading of a transceiver must remain within the shading region to be one unit load.

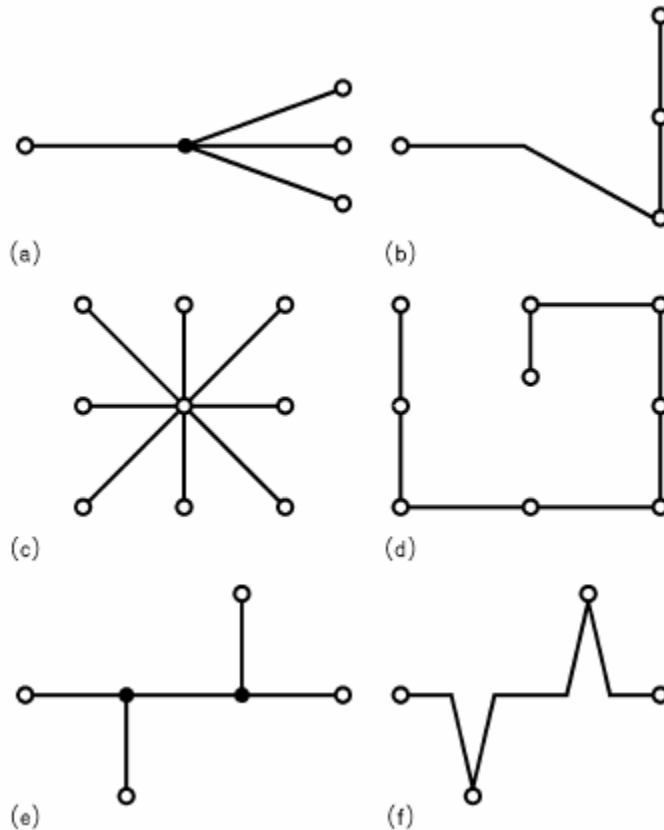


FIGURE 3.

Although you can make RS-485 buses with all of these configurations work, you should avoid the ones in (a), (c), and (e). Those in (b), (d), and (f), offer superior transmission-line performance.

Warning: A final word of caution: Do not connect too many transceivers to the bus or lump too many transceivers too close together. The AC loading, which results mainly from the devices' pin I/O capacitance (usually about 15 pF), can alter the interconnect-medium impedance and cause transmission-line problems.

Configurations

Because RS-485 allows connecting multiple transceivers, the bus configuration is not as straightforward as in a point-to-point bus (RS-232C, for example). In a point-to-point bus, a single driver connects to one receiver alone. The optimal configuration for the RS-485 bus is the daisy-chain connection from node 1 to node 2 to node 3 to node n. The bus must form a single continuous path, and the nodes in the middle of the bus must not be at the ends of long branches, spokes, or stubs. *Figure 3a*, *Figure 3c*, and *Figure 3e* illustrate three common but *improper* bus configurations. (If you mistakenly use one of these



configurations, you can usually make it work but only through substantial effort and modification.) *Figure 3b*, *Figure 3d*, and *Figure 3f* show equivalent daisy-chained configurations. Connecting a node to the cable creates a stub, and, therefore, every node has a stub. Minimizing the stub length minimizes transmission-line problems. For standard transceivers with transition times around 10 ns, stubs should be shorter than 6 in. A better rule is to make the stubs as short as possible. A “star” configuration (*Figure 3c*) is a special case and a cause for concern. This configuration usually does not provide a clean signaling environment even if the cable runs are all of equal length. The star configuration also presents a termination problem, because terminating every endpoint would overload the driver. Terminating only two endpoints solves the loading problem but creates transmission-line problems at the un-terminated ends. A true daisy-chain connection avoids all these problems.

Interconnect Media

The standard specifies only the driver-output and receiver-input characteristics—not the interconnection medium. You can build RS-485 buses using twisted-pair cables, flat cable, and other media, even backplane pc traces. However, twisted-pair cable is the most common. You can use a range of wire gauges, but designers most frequently use 24 AWG. The characteristic impedance of the cable should be 100Ω to 120Ω. A common misconception is that the cable’s characteristic impedance (Z_0) must be 120Ω, but 100Ω works equally well in most cases. Moreover, the 120Ω cable’s higher Z_0 presents a lighter load, which can be helpful if the cable runs are extremely long. Twisted pair offers noise benefits over flat or ribbon cables. In flat cable, a noise source (usually a conductor carrying an unrelated signal) can be closer to one member of the conductor pair than to the other over an entire wiring-run length. In such cases, more noise capacitively couples to the closer conductor than to the more distant one, producing a differential noise signal that can be large enough to corrupt the data. When you use the twisted pair, the noise source is closest to each of the conductors for roughly half of the wiring-run length. Therefore, the two conductors pick up roughly equal noise voltages. The receiver rejects these voltages because they appear mainly as common mode.

A special ribbon cable that is useful for noise reduction intermixes relatively long twisted sections with short flat sections. This cable provides the advantages of twisted pair between the flat sections and allows the use of insulation displacement connectors at the flat points.

Data Rate Vs Cable Length

You can transmit data over an RS-485 bus for 4000 ft (1200m), and you can also send data over the bus at 10 Mbps. But, you cannot send 10 Mbps data 4000 ft. At the maximum cable length, the maximum data rate is not obtainable: The longer the cable, the slower that data rate, and vice versa. *Figure 4a* shows a conservative curve of data rate vs cable length for RS-422 and RS-485. The two slopes result from different limitations. The maximum cable length is the result of the voltage divider that the cable’s



DC loop resistance and the termination resistance create. Remember, for differential buses, the loop resistance is twice as high as you might expect, because both conductors in the pair equally contribute.

The curve's sloped portion results from AC limitations of the drivers and the cable.

Figure 4b shows four limits for a DS3695 transceiver that drives a common twisted-pair cable. Notice that the data rate vs cable length depends significantly on how you determine the necessary signal quality. This graph includes two types of criteria. The first is a simple ratio of the driver's transition time to the unit interval. A curve showing the results for the common 30% ratio defines the most conservative set of operating points.





FIGURE 4

As you increase the cable length, the maximum data rate decreases. The more jitter you can accept, the greater is the allowable data rate for a given length of cable.

A second method of determining the operating points uses eye-pattern (jitter) measurements. To make such measurements, you apply a pseudo random bit sequence (PRBS) to the driver's input and measure the resulting eye pattern at the far end of the



cable. The amount of jitter at the receiver's threshold vs the unit interval yields the data point. Less jitter means better signal quality. Common operating curves use 5, 10, or 20% jitter. Above 50%, the eye pattern starts to close, and error-free data recovery becomes difficult (Reference 2). The key point is that you can't obtain the maximum data rate at the maximum cable length. But, if you operate the bus within the published, conservative curves, you can expect an error-free installation.

Termination and Stubs

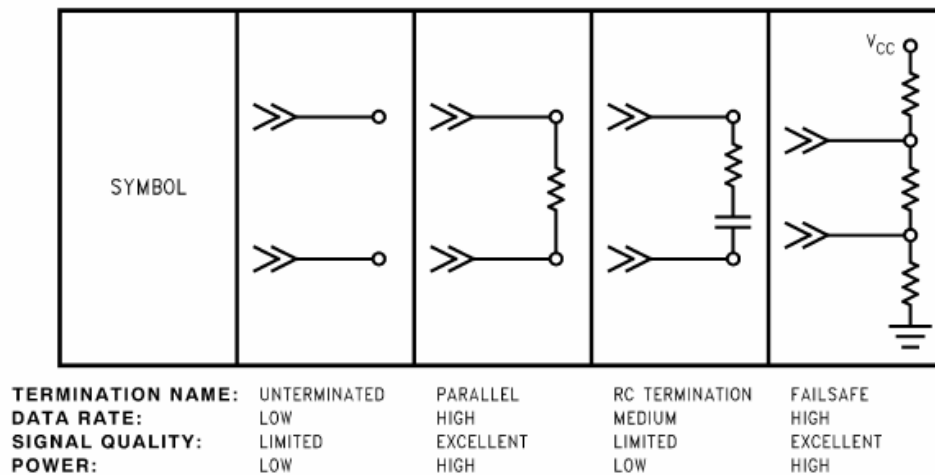
Most RS-485 buses require termination because of fast transitions, high data rates, or long cables. The purpose of the termination is to prevent adverse transmission-line phenomena, such as reflections. Both ends of the main cable require termination. A common mistake is to connect a terminating resistor at each node—a practice that causes trouble on buses that have four or more nodes. The active driver sees the four termination resistors in parallel, a condition that excessively loads the driver. If each of the four nodes connects a 100 Ω termination resistor across the bus, the active driver sees a load of 25 Ω instead of the intended 50 Ω . The problem becomes substantially worse with 32 nodes. If each node includes a 100 Ω termination resistor, the load becomes 3.12 Ω . You can include provisions for termination at every node, but you should activate the termination resistors only at the end nodes (by using jumpers, for example). Stubs appear at two points. The first is between the termination and the device behind it. The second is between the main cable and a device at the middle of the cable. *Figure 1* shows both stubs. The symbol "l" denotes the stub length. Keep this distance as short as possible. Keeping a stub's electrical length below one-fourth of the signal's transition time ensures that the stub behaves as a lumped load and not as a separate transmission line. If the stub is long, a signal that travels down the stub reflects to the main line after hitting the input impedance of the device at the end of the stub. This impedance is high compared with that of the cable. The net effect is degradation of signal quality on the bus. Keeping the stubs as short as possible avoids this problem. Instead of adding a long branch stub, loop the main cable to the device you wish to connect. If you must use a long stub, drive it with a special transceiver designed for the purpose.

Termination Options

You have several options for terminating an RS-485 bus. The first option is no termination. This option is feasible if the cable is short and if the data rate is low. Reflections occur, but they settle after about three round-trip delays. For a short-cable, the round-trip delay is short and, if the data rate is low, the unit interval is long. Under these conditions, the reflections settle out before sampling, which occurs at the middle of the bit interval. The most popular termination option is to connect a single resistor across the conductor pair at each end. The resistor value matches the cable's differential-mode characteristic impedance. If you terminate the bus in this way, no reflections occur, and the signal fidelity is excellent. The problem with this termination option is the power dissipated in the termination resistors. If you must minimize power dissipation, an RC termination may be the solution. In place of the single resistor, you use a resistor in series



with a capacitor. The capacitor appears as a short circuit during transitions, and the resistor terminates the line. Once the capacitor charges, it blocks the DC loop current and presents a light load to the driver. Low-pass effects limit use of the RC termination to lower data-rate applications, however (Reference 3). Another popular option is a modified parallel termination that also provides a fail-safe bias. A detailed discussion of fail-safe biasing occurs later in the article. *Figure 5* compares the four popular termination methods. The main point to remember is that, if you use termination, you should locate the termination networks at the two extreme ends of the cable, not at every node.



RS-485 buses can use four methods of termination. Achieving the best electrical performance requires accepting higher power dissipation in the termination resistors.

FIGURE 5.

Grounding and Shielding

Although the potential difference between the data-pair conductors determines the signal without officially involving ground, the bus needs a ground wire to provide a return path for induced common-mode noise and currents, such as the receivers' input current. A typical mistake is to connect two nodes with only two wires. If you do this, the system may radiate high levels of EMI, because the common-mode return current finds its way back to the source, regardless of where the loop takes it. An intentional ground provides a low-impedance path in a known location, thus reducing emissions.

Electromagnetic-compatibility and application requirements determine whether you need a shield. A shield both prevents the coupling of external noise to the bus and limits emissions from the bus. Generally, a shield connects to a solid ground (normally, the metal frame around the system or subsystem) with a low impedance at one end and a series RC network at the other. This arrangement prevents the flow of DC ground-loop currents in the shield.



Contention Protection

Because RS-485 allows for connecting multiple drivers to the bus, the standard addresses the topic of contention. When two or more drivers are in contention, the signal state on the bus is not guaranteed. If two drivers are on at the same time and if they are driving the same state, the bus state is valid. However, if the drivers are in opposite states, the bus state is undetermined, because the differential voltage on the bus drops to a low value within the receiver's threshold range. Because you do not know the driver states, you must assume the worst—namely, that the data on the bus is invalid. Contention can also damage the ICs. If several drivers are in one state, a single driver in the opposite state sinks a high current (as much as 250 mA). This large current causes excessive power dissipation. A difference in ground potential between nodes only aggravates this dissipation. In this situation, the driver's junction temperature can increase beyond safe limits. The RS-485 standard recommends the use of special circuitry, such as a thermal shutdown circuit, to prevent such damage. Most RS-485 devices use this technique. The shutdown circuit disables the driver outputs when the junction temperature exceeds 150°C and automatically re-enables the outputs when the junction cools. If the fault is still present, the device cycles into and out of thermal shutdown until someone clears the fault. Besides thermal shutdown, other current limiting is required to prevent accidental damage. If an active output is shorted to any voltage within the -7V to +12V range, the resulting current must not exceed 250 mA. In addition, the outputs of a driver must not sustain damage if they are shorted together indefinitely. (Entering thermal shutdown is allowed, of course.) Lastly, RS-485 drivers must source and sink large currents (60 mA). This situation requires outputs of rather large geometry, which provide robust ESD protection.

Fail-Safe Biasing

The need for fail-safe operation is both the principal application issue and most frequently encountered problem with RS-485. Fail-safe biasing provides a known state in which there are no active drivers on the bus. Other standards do not have to deal with this issue, because they typically define a point-to-point or multidrop bus with only one driver. The one driver either drives the line or is off. Because there is only one source on the bus, the bus is off when the driver is off. RS-485, on the other hand, allows for connection of multiple drivers to the bus. The bus is either active or idle. When it is idle with no drivers on, a question arises as to the state of the bus. Is it high, low, or in the state last driven?


The answer is any of the above. With no active drivers and low-impedance termination resistors, the resulting differential voltage across the conductor pair is close to zero, which is in the middle of the receivers' thresholds. Thus, the state of the bus is truly undetermined and cannot be guaranteed. Some of the functional protocols that many applications use aggravate this problem. In an asynchronous bus, the first transition indicates the start of a character. It is important for the bus to change states on this leading edge. Otherwise, the clocking inside the UART is out of sync with the character



and creates a framing error. The idle bus can also randomly switch because of noise. In this case, the noise emulates a valid start bit, which the UART latches. The result is a framing error or, worse, an interrupt that distracts the CPU from other work.

The way to provide fail-safe operation requires only two additional resistors. At one end of the bus (the master node, for example), connect a pull-up and pull-down resistor (*Figure 10*). This arrangement provides a simple voltage divider on the bus when there are no active drivers. Select the resistors so that at least 200 mV appears across the conductor pair. This voltage puts the receivers into a known state. Values that can provide this bias are 750Ω for the pull-up and pull-down resistors, 130Ω across the conductor pair at the fail-safe point, and a 120Ω termination at the other end of the cable. For balance, use the same value for the pull-up and pull-down resistors. Reference 4 provides extensive details on this issue.

Forethought into these 10 areas before production greatly reduces the likelihood of problems. RS-485 is unique in its capabilities and requirements. Fully understanding these 10 issues leads to a rock-solid, trouble-free, multipoint differential interface that maximizes the benefits of RS-485 and provides the application with robust, rugged, highly noise-tolerant data communication.



Unless you do something to keep the situation from occurring, when no driver is driving the bus, the receivers cannot determine the bus state. Fail-safe biasing is a bus-termination method, which ensures that, even when no driver is active, a differential voltage large enough to unambiguously determine the bus state appears at the receiver inputs.