

# Runway Bus Electrical Design Considerations

The Runway bus's high bandwidth is a result of the strategies adopted for its electrical design. These included an efficient data transfer scheme, a simple clock system with low clock skew, a compact bus topology, and a termination strategy that eliminates dead cycles when changing bus masters.

## Data Transfer Scheme

The simple data transfer strategy, shown in Fig. 1, allows most of the cycle to be used to transfer data. An edge-triggered Runway pad driver is enabled by the rising edge of the on-chip Runway clock, RCK, causing the data to be driven onto the external bus. This driven data is then latched one cycle later at the receiving devices by the next rising edge of the receiver's on-chip Runway clock. On each Runway VLSI chip, the same physical clock edge is used to trigger the signal driver and latch the data from the previous cycle in the signal receiver.

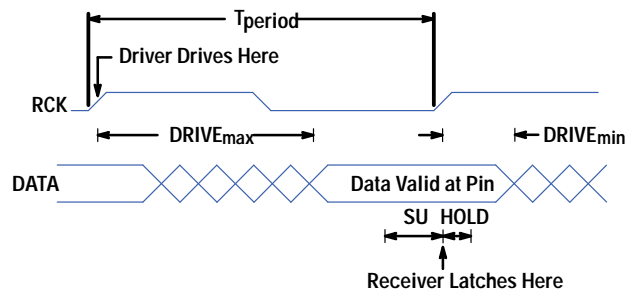
The following two equations express timing constraints that must be met for proper operation:

Setup time equation:  $DRIVE_{max} + SKEW + SU \leq T_{period}$

Hold time equation:  $SKEW + HOLD \leq DRIVE_{min}$

where DRIVE is the delay from the rising edge of RCK at the driver to the time when the data is valid at the receiver, SU is the receiver setup time, HOLD is the receiver hold time, SKEW is the maximum skew of the clock signal (RCK) between the driver of one chip and the receiver of another, and  $T_{period}$  is the clock period.

**Fig. 1.** The data transfer timing strategy for the Runway bus allows most of the cycle to be used to transfer data.



## Clock Path

The Runway clock orchestrates the transfer of information among the components on the bus. The path of the Runway clock to the driver and receiver circuits can be divided into three components: on-board clock generation and distribution to the VLSI chip inputs, on-chip clock reception and buffering, and on-chip clock distribution to the Runway driver and receiver circuits.

Skew can be introduced by any of these components. Inspection of the setup and hold time equations reveals that it is desirable to reduce skew to as small a value as possible.

The clock path begins at the custom VLSI clock generation chip. This chip generates several differential pairs of clock outputs, one per Runway VLSI chip. By using one chip as the source of the clock signals in this system, the output-to-output skew was kept very small.

Each dedicated clock pair is carefully routed on the printed circuit board to its Runway VLSI chip. The traces are adjusted in length so that the arrival time of each clock at the pins of the Runway VLSI chips can be accurately placed with respect to the others. Because of known timing differences in the paths of the clock from input pin to driver and receiver circuits for each type of Runway VLSI chip, it is useful to be able to tune the clocks in this manner. More will be said about this later.

Each differential clock signal is received at each chip by a receiver/buffer circuit that transforms the signal into a single-ended signal RCK with normal CMOS voltage levels. This RCK signal then fans out to all the Runway signal driver and receiver circuits located at the pads and the associated interface circuitry located in the core of the chip. Since the interface circuitry is similar on all three types of Runway VLSI chips, the capacitive loading on RCK is nearly identical for all three types, which ensures that the delay through the clock buffer is similar for all Runway VLSI chips.

The RCK signal is routed using various techniques to reduce the distribution delay and thus the variation in delay. The clock receiver/buffer bit slice is centrally placed in the interface so that the total distance the RCK signal must travel on-chip to the farthest signal bit slice is minimized. This clock is routed in wide metal so that the delay along this line is low. The signal pad

ordering in the Runway interfaces for all of the Runway VLSI chips is nearly identical. This ensures that the distance from the clock buffer to a signal pad is the same for all of the Runway VLSI chips.

The goal in the design of the Runway clock system was to have the on-chip clock RCK arrive at the corresponding signal driver or receiver at the same time at each Runway device. Since the CPU is fabricated in CMOS14, a faster technology than the CMOS26 process used for the I/O adapter and memory controller chips, the on-board clock signal to the CPU is delayed to account for this known timing difference. Thus the clock skew is only a function of the CMOS26 parameters, which keeps the skew to a minimum.

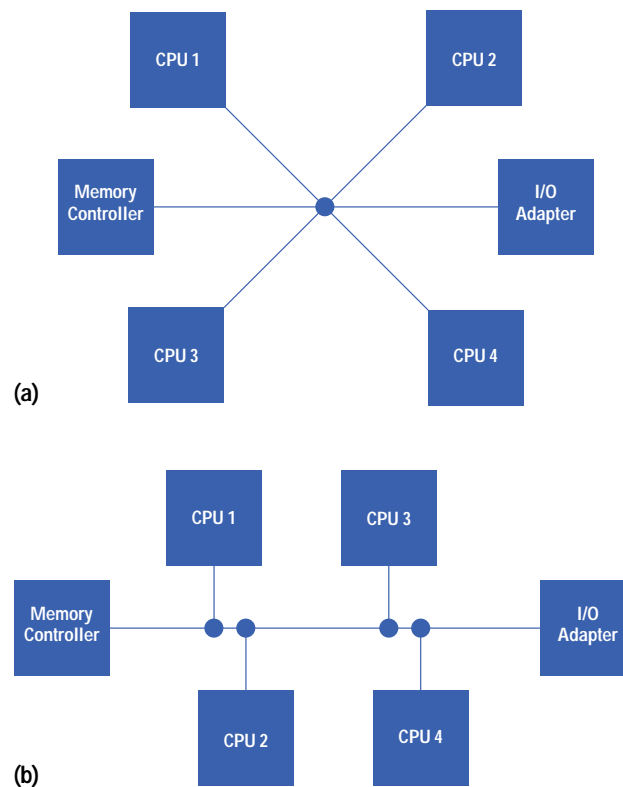
Overall, the total chip-to-chip clock skew on RCK at the signal driver and receiver circuits is under 1.1 ns worst-case.

## Bus Topology

The components on the bus are designed to be close together to limit the capacitive and inductive load on each Runway signal line. The setup and hold time equations can be used to determine how best to lay out the signal path. The requirements of the two equations sometimes conflict. For example, the setup time equation wants us to minimize  $DRIVE_{max}$  and the hold time equation wants us to maximize  $DRIVE_{min}$ . In plain English, we want an interconnect scheme that minimizes the overall trace length while maintaining the greatest separation between components.

An ideal connection topology would be a star with the devices placed at the tips of the star as shown in Fig. 2a. Because of manufacturing difficulties with this topology, the modified star shape shown in Fig. 2b, which fits comfortably using standard printed circuit board technology, was chosen. As the figure suggests, the main trunk of the Runway bus consists of a standard printed circuit trace running along a backplane with at most four daughter cards attached to the backplane, two per side. Each daughter card will hold one CPU. The memory controller and the I/O adapter reside on the backplane along with the clock generation circuitry. This connection scheme interconnects six Runway devices with less than 9 inches of total printed circuit trace for the longest signal with no two devices farther apart than 4.5 inches.

**Fig. 2.** (a) Star topology. (b) Modified star topology of the Runway bus.



## Termination Strategy

A parallel termination strategy using external resistors is usually used for high-speed buses so that incident wave switching can be employed, that is, the receiver will switch (receive) when it detects the first or incident wave propagating down the transmission line. In this style of termination, the driver usually only needs to drive in one direction. The termination resistor drives the bus the other way when the driver tristates or turns off. While the driver is on, direct current flows constantly. When the driver turns off, the bus is disturbed by the change of current through the inductive traces and bond wires. This disturbance sends a wave propagating down the transmission line in the direction opposite to the direction of propagation when the driver turns on. A special frequency-limiting case is the master changeover, when a driver at the end of the bus starts to drive the same value that was driven by the master at the other end of the bus in the previous cycle. In

this case, constructive interference of the two propagating waves may cause the bus to take a long time to settle. It is not uncommon to insert a dead cycle in the protocol to allow extra time for the bus to settle when the bus changes masters.

On a series-terminated bus, the bus driver has the ability to drive in both directions. The on-impedance of the driver transistor acts as the termination resistor. The driver transistor will turn on and drive the bus to the desired level. Near the end of the cycle when the bus is nearing its final value, the drivers will be sourcing or sinking only a small fraction of their peak currents at the start of the cycle.

Because of this, when the driver is disabled at the end of the cycle, there is very little disturbance in the line. This makes it possible to have another driver master the bus in the very next cycle. However, because the on-impedance of the driver is not well-controlled, the receiver must usually wait to receive the reflected wave, which increases the bus propagation delay.

Runway bus topology is very compact. This means that the time difference between the arrivals of the incident and reflected waves is relatively small compared to the bus cycle time. Had we employed parallel termination on the Runway bus, we might have been able to increase the frequency of the bus by about 20%. Since the dead cycle would have cost us about 20 to 30 percent in bandwidth, we decided to use series termination instead.

Other advantages of series-terminated buses include good tolerance to impedance mismatches and long stubs and no dc power dissipation. Also, we saved valuable board space by not having to place resistors on each Runway bus signal.

### Simulated and Characterized Performance

Early in the Runway bus simulations, it became clear that the result would be dependent not only on which driver drove the bus, but also on the current state of the bus. The state of the bus is precisely determined by the history of drivers driving the bus along with the starting condition of the bus. Since the current state of the bus is mostly determined by who was last driving it, all possible pairs of successive bus transactions by two drivers were simulated. The symmetry of the bus and our ability to predict and eliminate combinations that would not be worst-case helped cut down the number of slow-case simulations to 32. A network of fast HP 9000 Model 720 and 750 workstations was able to run these simulations, each of which normally takes one machine about one hour to run, in about 4 hours.

The SPICE model to simulate the worst case was in constant revision as more and more details from the design were implemented. The final model had artwork-extracted transistor models for the signal driver and receiver for each Runway VLSI chip and a detailed schematic model for each package trace and board connector. The printed circuit traces were modeled using SPICE transmission-line primitives.

The final simulated worst-case bus frequency came in at 152 MHz using a fully loaded Runway bus. The characterized frequency of the bus over the extremes of process, temperature, and voltage showed operation of at least 140 MHz. The maximum characterization frequency was limited to 140 MHz because of the limitations of other system components. These results gave us the confidence to conclude that the Runway bus will work at the 120-MHz frequency goal with sufficient manufacturing margin.

### Acknowledgments

I would especially like to thank Fred Eatock for his countless and thorough Runway bus simulations, Dave "Spreadsheet" Malicoat for his careful design of the clock system which kept skew low, Ken Pomaranski and Tony Chan for their very compact board layouts, Denny Renfrow for suggesting and designing the full-state predriver, John Youden for his incredible understanding of transmission line phenomena, Craig Gleason for his invaluable design suggestions, and all the other great engineers who had a hand in this design.

---

---

**Nicholas S. Fiduccia**  
**Member of the Technical Staff**  
**Systems Technology Division**

---

---

- ▶ [Go to Article 3](#)
- ▶ [Go to Table of Contents](#)
- ▶ [Go to HP Journal Home Page](#)