

By Mark Hoke

INTRODUCTION

High-speed packet processing requires the movement of packets within a system product with minimal latency and overhead. The system product connects to the data communication network (i.e. Internet or corporate backbone) through a line interface module. The line interface module is comprised of the Physical (PHY), Serial De-serialize (SERDES) and MAC (Media Access Control) layers of the OSI model. Some high-speed product designs require a local processor on the interface module to reduce the burden of down stream processing modules. The data path between the PHY/MAC circuitry and the processor requires a bus and/or rate matching, packet-buffering device. To minimize latency and maximize performance the data path should contain a minimum number of devices. This application brief describes how the IDT multi-queue flow-control devices (IDT72T515XX) can be used in a high-speed packet processing architecture (Figure 1).

HIGH-SPEED PACKET PROCESSING OVERVIEW

High-speed packet processing systems can connect to line interfaces ranging from OC-12 up to OC-192 (10 Gbps) and have internal system clock frequencies operating at 200 Mhz and higher. In addition to high-speed input/output line rates and high frequency system clocks, system requirements include low system latency, streamline data paths and minimal system software overhead.

As the line rate and system clock frequencies increase, the data cells become smaller (the data cell is the inverse proportion of the clock frequency). Refer to Figure 2, *Data Cell*. Small data cells require high-speed components, such as the IDT flow-control products, to handle throughput demands.

UNIQUE MULTI-QUEUE FUNCTION

One area affecting performance of a system product or box in a network is the ability to provide system resources to the right data at the right time. As more and more traffic is time sensitive, such as voice and video, the timely delivery of this type of data is mandatory. Bottle necks in a system can hinder or stop the time sensitive data stream. There is also the need for different levels of service and policing in a system.

One way to handle these types of data is to provide a prioritization algorithm that enables the highest priority data access to network resources first. Other issues of starvation and fairness are also applied. This prioritization function can be implemented by either a large, shared memory or multiple smaller queues, one or more for each priority. Using an FPGA or ASIC for this function is complex and filled with performance issues. IDT has now taken that function from the custom world and provided it in an off-the-shelf device.

The IDT multi-queue flow-control devices, perform specialized functions in the data path. Such functions were previously only available in custom, "home-grown" solutions, but are now available off the shelf. These functions off-load the FPGA or ASIC by reducing gate count, I/Os and allowing for smaller, less expensive, custom devices. By bringing data-path functions off of a programmable device, performance is often increased as well.

The multi-queue function of the flow-control devices provide the ability to take a single bus, with multiple virtual channels and rearrange their order on the wire based on a system algorithm. We will now discuss the function of the multi-queue devices to show how the reordering of data is achieved, allowing first access to the highest priority data.

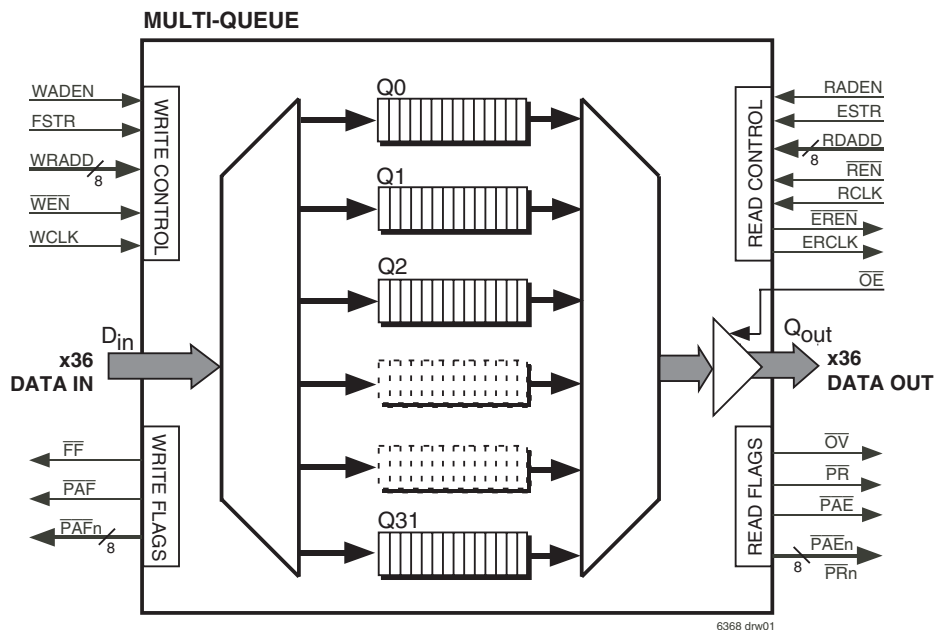


Figure 1. Block Diagram - Multi-queue Flow-control Device

IDT72T515XX MULTI-QUEUE PRODUCT INFORMATION

PRODUCT HIGHLIGHTS

- Configurable from 1 to 32 queues
- Independent read and write access per queue
- User selectable I/O: 2.5V LVTTTL, 1.5V HSTL, 1.8V eHSTL
- Default multi-queue device configurations
 - IDT72T51546 : 1,024 x 36 x 32Q
 - IDT72T51556 : 2,048 x 36 x 32Q
- 100% bus utilization
- 200 MHz maximum clock frequency
- 3.6ns access time
- Multiple bus matching options
 - x36in to x36out
 - x18in to x36out
 - x9in to x36out
 - x36in to x18out
 - x36in to x9out

OPERATIONAL OVERVIEW

The IDT72T515XX multi-queue flow-control devices are single ICs containing up to 32 discrete queues. To minimize printed circuit board trace routing, the multi-queue devices use common data buses to communicate to/from the queues. The data path is structured as a common Din (36:0) input bus, (write port) and a common Qout (36:0) output bus, (read port). Information grams (encapsulated voice, video or data) are written into queue via the write port and read from the read port. Queue writes and reads can be performed at frequencies up to 200 Mhz. Data write and read operations are totally independent of each other, any queue may be selected on the write port and any queue can be accessed on the read port. Queues support simultaneous write and read operations.

For packetized information the multi-queue flow-control devices can be configured for packet mode operation. Packet mode provides an efficient method to mark the start and end of packets stored in the queues.

The multi-queue ICs supports a variety of industry standard input/output specifications such as 2.5V LVTTTL, 1.5V HSTL or 1.8V eHSTL. Refer to Figure 1, *Block Diagram - Multi-queue Flow-control Device* for additional product information.




Figure 2. Data Cell

TYPICAL IMPLEMENTATION

As illustrated in Figure 3, *Typical High-Speed Line Interface Module Block Diagram*, a typical high speed interface module implementation includes the use of a custom ASIC (Application Specific Integrated Circuit) or an FPGA (Field Programmable Gate Array) device in the data path for bus matching, rate matching and packet buffering between the input/output module and the processing interface.

Analysis of Figure 3 reveals several issues to be considered with an implementation that uses either a custom ASIC or the FPGA in the main data path. One of the issues is the circuit complexity required to be designed inside the ASIC/FPGA including the design and verification/test time is very costly, largely due to the vast space required (i.e., gate count) within the device. Another issue flexibility. When the packet buffer is designed as an integral element of an ASIC/FPGA, the buffer is a fixed size, therefore limiting flexibility

and not adaptable to various incoming line rates. If it is an inadequate size, latency and performance are not optimized. The flow-control devices allow flexibility in the system as well as system simplicity (low cost, fast time-to-market).

FLOW-CONTROL MANAGEMENT IMPLEMENTATION

As shown in Figure 4, *High-Speed Packet Processing Using IDT Multi-queue Flow-control Devices* the IDT flow-control devices can provide efficient rate matching between various line rates, comprehensive bus matching and flexible packet buffering. The flow-control products support operating frequencies up to 200 Mhz and transfer rates up to 16 Gb/s. Utilizing flow-control management devices for the packet buffer from a product architect standpoint adds flexibility and extensibility. The packet buffering queues can be configured to optimize overall system throughput.




Figure 3. Typical High-Speed Line Interface Module Block Diagram

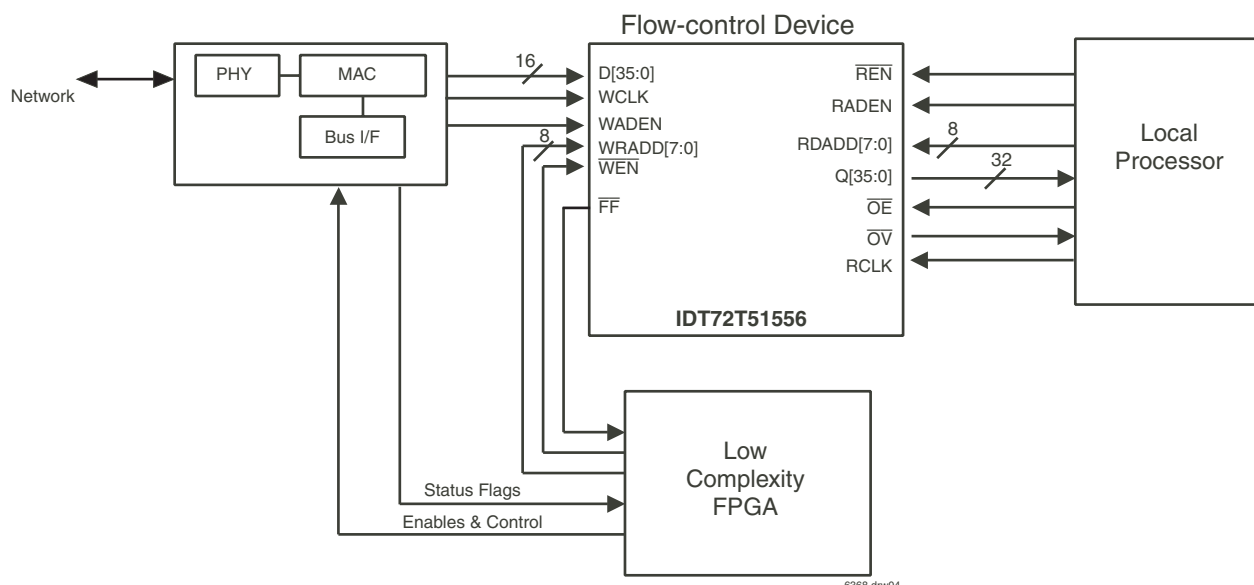


Figure 4. High-Speed Packet Processing Using IDT Multi-queue Flow-control Devices

TABLE 1 — BENEFITS OF THE IDT72T515XX FLOW-CONTROL DEVICES

FUNCTIONALITY	FCM IMPLEMENTATION	TYPICAL IMPLEMENTATION
Queues	Flexible, up to 32 queues per device can be defined to the application	Rigid, uses an embedded monolithic memory array
Device Cascading	Available	Not Available
Buffer Architecture	Enables a consolidated buffering architecture to reduce system latency	Multiple buffers & memory controllers fragmented across several devices, which increase system latency
Buffer Size	Adjustable buffer size to minimize processor/system overhead	Small fix sized buffers, requiring more processor overhead/system time to manage the data path
System Design time	Decreased design time by reducing complexity of data path, memory controllers and FPGA, hence shorter design time	Increased design time due to multiple memory controllers, complex data path and FPGA control functions.

BENEFITS

The benefits of using IDT flow-control devices are numerous. Benefits include: (1) packet-buffering flexibility; the queues used to buffer packets can be adjusted in both depth and width to accommodate the incoming line rate and packet size (2) design flexibility, the number of packet buffering queues are user definable (3) using flow-control devices reduce the overall design risk and complexity of the custom ASIC or FPGA, thereby reducing time to market (4) a smaller ASIC or FPGA may be used, which translates into a direct material cost savings.

OPERATIONAL INFORMATION

WRITE QUEUE SELECTION & WRITE OPERATION

The IDT72T515XX multi-queue ICs can be configured with up to 32 queues that data can be written into via a common write port using the data inputs, Din, Write Clock (WCLK) and Write Enable (WEN). A queue is selected using the write address bus (WRADD), the rising edge of WCLK and Write Address Enable (WADEN). The queue selection occurs on a single WCLK cycle, the selected queue will remain selected until another queue is selected.

The write port is designed to achieve 100% bus utilization. This means that data can be written into the device on every WCLK rising edge including the cycle that a new queue is being addressed.

All subsequent writes will be written to that queue until a new queue is selected. A minimum of 3 WCLK cycles must occur between queue selections on the write port. On the next WCLK rising edge the write port discrete full flag will update to show the full status of the newly selected queue. On the second rising edge of WCLK, data present on the data input bus, Din can be written into the newly selected queue provided that WEN is LOW and the new queue is not full. The cycle of the queue selection and the next cycle will continue to write data present on the data input bus, Din into the previous queue provided that WEN is active LOW.

If the newly selected queue is full at the point of its selection, then writes to that queue will be prevented, a full queue cannot be written into. In the 32 queue multi-queue devices the WRADD address bus is 8 bits wide. The least significant 5 bits are used to address one of the 32 available queues within a single multi-

queue device. The most significant 3 bits are used when a device is connected in expansion mode, up to 8 devices can be connected in expansion. Refer Figure 5, *Write Operations & First Word Fall Through*.

READ QUEUE SELECTION & READ OPERATION

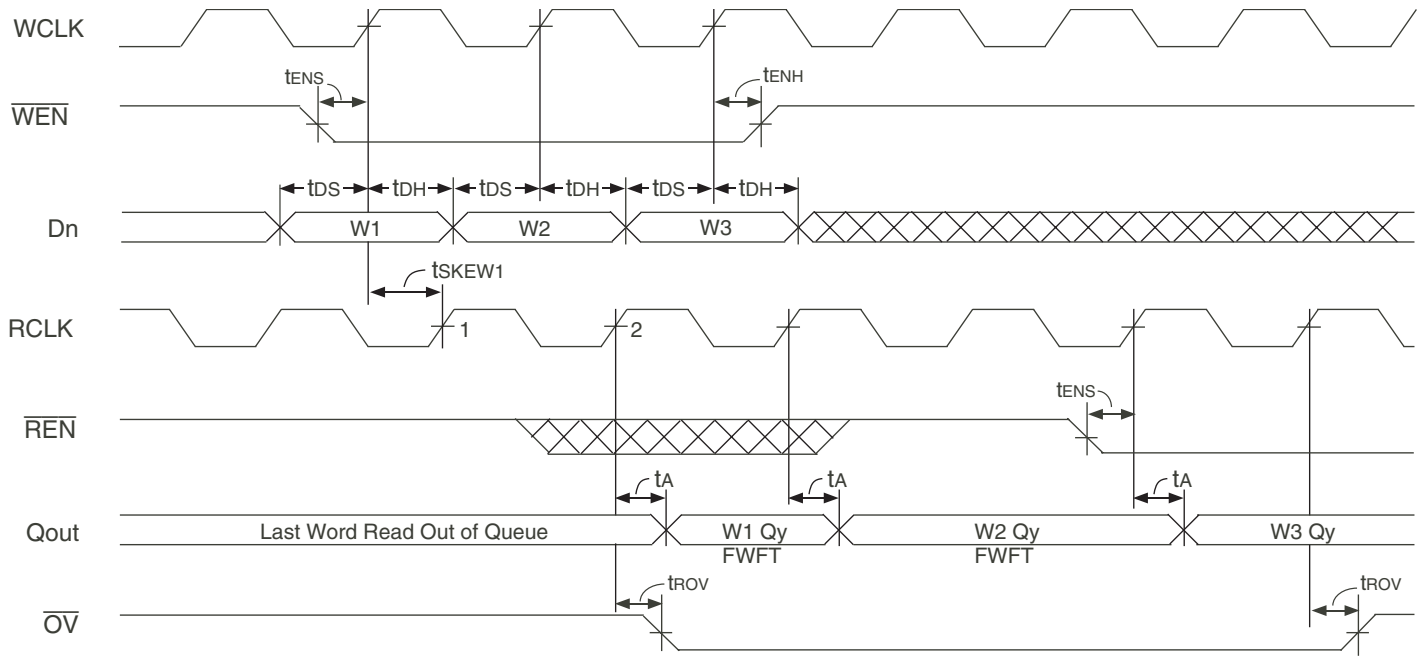
The multi-queue flow-control devices can be configured with up to 32 queues. Each queue is read from via a common read port using the data outputs, Qout, read clock (RCLK) and read enable (\overline{REN}). In a shared bus architecture, an output enable (\overline{OE}) control pin is also provided to place the output drivers in an High-Impedance state.

A queue is selected using the read address bus (RDADD), the rising edge of RCLK and read address enable (RADEN). Selecting a queue for reading requires only one cycle. All subsequent reads will be read from the selected queue until a new queue is selected. A minimum of 3 RCLK cycles must occur between queue selections on the read port. On the same RCLK rising edge that the new queue is selected, data can still be read from the previously selected queue, provided that \overline{REN} is LOW, active and the previous queue is not empty on the following rising edge of RCLK, a word will be read from the previously selected queue regardless of \overline{REN} due to the fall through operation (provided the queue is not empty).

When a queue is selected on the read port, the next word available in that queue (provided that the queue is not empty) will fall through to the output register. In a 32 queue configuration the RDADD address bus is used to address the 32 available queues. The most significant 3 bits are used for addressing devices in expansion mode, up to 8 devices can be cascaded. Refer to Figure 6, *Read Queue Select, Read Operation and \overline{OE} Timing*.

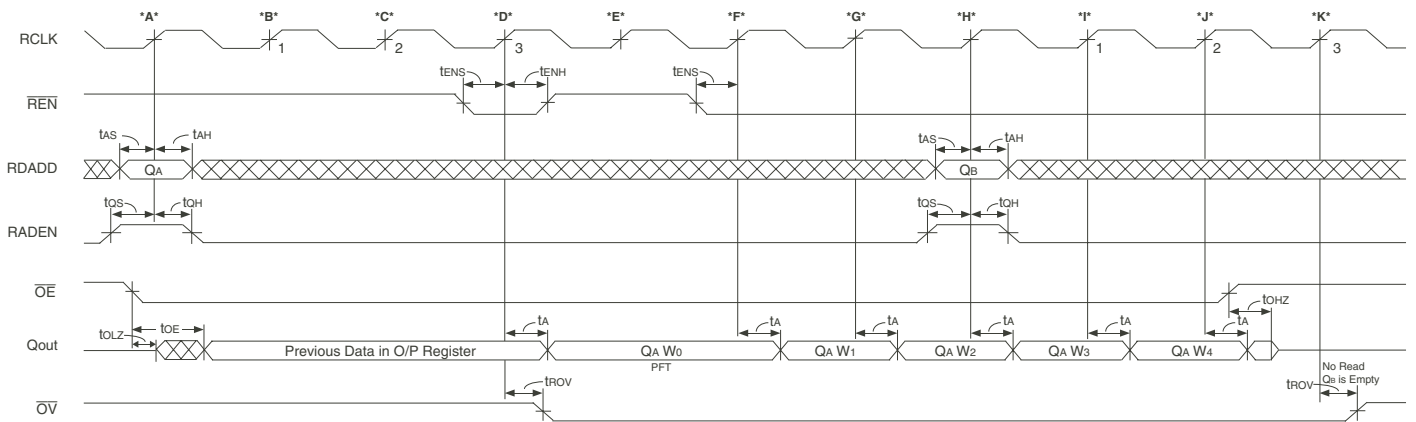
CONCLUSION

Multi-queue flow-control devices provide the most flexibility and cost effectiveness. The multi-queue flow-control devices are an all-in-one solution providing rate matching, bus matching and a flexible buffering queue architecture. By using these devices, functions that were previously only available in "home-grown", custom designs such as FPGAs or ASICs, are now available off the shelf. They come with complete models, fully tested and verified so designers can simplify their designs, shorten design cycles, reduce design and system costs and get to market faster. All of this with full-line rate performance.



6368 drw05

Figure 5. Write Operations & First Word Fall Through



6368 drw06

Figure 6. Read Queue Select, Read Operation and \overline{OE} Timing

AC ELECTRICAL CHARACTERISTICS

(Commercial: $V_{CC} = 2.5V \pm 0.15V$, $T_A = 0^\circ C$ to $+70^\circ C$; Industrial: $V_{CC} = 2.5V \pm 0.15V$, $T_A = -40^\circ C$ to $+85^\circ C$; JEDEC JESD8-A compliant)

Symbol	Parameter	Commercial		Com'l & Ind'l ⁽¹⁾		Unit
		IDT72T51546L5 IDT72T51556L5		IDT72T51546L6 IDT72T51556L6		
		Min.	Max.	Min.	Max.	
f _s	Clock Cycle Frequency (WCLK & RCLK)	—	200	—	166	MHz
t _A	Data Access Time	0.6	3.6	0.6	3.7	ns
t _{CLK}	Clock Cycle Time	5	—	6	—	ns
t _{CLKH}	Clock High Time	2.3	—	2.7	—	ns
t _{CLKL}	Clock Low Time	2.3	—	2.7	—	ns
t _{DS}	Data Setup Time	1.5	—	2.0	—	ns
t _{DH}	Data Hold Time	0.5	—	0.5	—	ns
t _{ENS}	Enable Setup Time	1.5	—	2.0	—	ns
t _{ENH}	Enable Hold Time	0.5	—	0.5	—	ns
t _{RS}	Reset Pulse Width	30	—	30	—	ns
t _{RSS}	Reset Setup Time	15	—	15	—	ns
t _{RSR}	Reset Recovery Time	10	—	10	—	ns
t _{PRSS}	Partial Reset Setup	1.5	—	2.0	—	ns
t _{PRSH}	Partial Reset Hold	0.5	—	0.5	—	ns
t _{OLZ} ($\overline{OE-Q_n}$) ⁽²⁾	Output Enable to Output in Low-Impedance	0.6	3.6	0.6	3.7	ns
t _{OHZ} ⁽²⁾	Output Enable to Output in High-Impedance	0.6	3.6	0.6	3.7	ns
t _{OE}	Output Enable to Data Output Valid	0.6	3.6	0.6	3.7	ns
f _c	Clock Cycle Frequency (SCLK)	—	10	—	10	MHz
t _{SCLK}	Serial Clock Cycle	100	—	100	—	ns
t _{SCKH}	Serial Clock High	45	—	45	—	ns
t _{SCKL}	Serial Clock Low	45	—	45	—	ns
t _{SDS}	Serial Data In Setup	20	—	20	—	ns
t _{SDH}	Serial Data In Hold	1.2	—	1.2	—	ns
t _{SENS}	Serial Enable Setup	20	—	20	—	ns
t _{SENH}	Serial Enable Hold	1.2	—	1.2	—	ns
t _{SDO}	SCLK to Serial Data Out	—	20	—	20	ns
t _{SENO}	SCLK to Serial Enable Out	—	20	—	20	ns
t _{SDOP}	Serial Data Out Propagation Delay	1.5	3.7	1.5	3.7	ns
t _{SENOP}	Serial Enable Propagation Delay	1.5	3.7	1.5	3.7	ns
t _{PCWQ}	Programming Complete to Write Queue Selection	20	—	20	—	ns
t _{PCRO}	Programming Complete to Read Queue Selection	20	—	20	—	ns
t _{AS}	Address Setup	1.5	—	2.5	—	ns
t _{AH}	Address Hold	1.0	—	1.5	—	ns
t _{WFF}	Write Clock to Full Flag	—	3.6	—	3.7	ns
t _{ROV}	Read Clock to Output Valid	—	3.6	—	3.7	ns
t _{STS}	$\overline{PAE}/\overline{PAF}$ Strobe Setup	1.5	—	2.0	—	ns
t _{STH}	$\overline{PAE}/\overline{PAF}$ Strobe Hold	0.5	—	0.5	—	ns
t _{QS}	Queue Setup	1.5	—	2.0	—	ns
t _{QH}	Queue Hold	1.0	—	0.5	—	ns
t _{WAF}	WCLK to \overline{PAF} flag	0.6	3.6	0.6	3.7	ns
t _{RAE}	RCLK to \overline{PAE} flag	0.6	3.6	0.6	3.7	ns
t _{PAF}	Write Clock to Synchronous Almost-Full Flag Bus	0.6	3.6	0.6	3.7	ns
t _{PAE}	Read Clock to Synchronous Almost-Empty Flag Bus	0.6	3.6	0.6	3.7	ns

NOTES:

1. Industrial temperature range product for the 6ns is available as a standard device. All other speed grades available by special order.
2. Values guaranteed by design, not currently tested.

AC ELECTRICAL CHARACTERISTICS (CONTINUED)

(Commercial: Vcc = 2.5V ± 0.15V, TA = 0°C to +70°C; Industrial: Vcc = 2.5V ± 0.15V, TA = -40°C to +85°C; JEDEC JESD8-A compliant)

Symbol	Parameter	Commercial		Com'l & Ind'l ⁽¹⁾		Unit
		IDT72T51546L5 IDT72T51556L5		IDT72T51546L6 IDT72T51556L6		
		Min.	Max.	Min.	Max.	
t _{ERCLK}	RCLK to Echo RCLK Output	—	4.0	—	4.2	ns
t _{CLKEN}	RCLK to Echo $\overline{\text{REN}}$ Output	—	3.6	—	3.7	ns
t _{PAELZ} ⁽²⁾	RCLK to $\overline{\text{PAE}}$ Flag Bus to Low-Impedance	0.6	3.6	0.6	3.7	ns
t _{PAEHZ} ⁽²⁾	RCLK to $\overline{\text{PAE}}$ Flag Bus to High-Impedance	0.6	3.6	0.6	3.7	ns
t _{PAFLZ} ⁽²⁾	WCLK to $\overline{\text{PAF}}$ Flag Bus to Low-Impedance	0.6	3.6	0.6	3.7	ns
t _{PAFHZ} ⁽²⁾	WCLK to $\overline{\text{PAF}}$ Flag Bus to High-Impedance	0.6	3.6	0.6	3.7	ns
t _{FFHZ} ⁽²⁾	WCLK to Full Flag to High-Impedance	0.6	3.6	0.6	3.7	ns
t _{FFLZ} ⁽²⁾	WCLK to Full Flag to Low-Impedance	0.6	3.6	0.6	3.7	ns
t _{OV LZ} ⁽²⁾	RCLK to Output Valid Flag to Low-Impedance	0.6	3.6	0.6	3.7	ns
t _{OV HZ} ⁽²⁾	RCLK to Output Valid Flag to High-Impedance	0.6	3.6	0.6	3.7	ns
t _{FSYNC}	WCLK to $\overline{\text{PAF}}$ Bus Sync to Output	0.6	3.6	0.6	3.7	ns
t _{FXO}	WCLK to $\overline{\text{PAF}}$ Bus Expansion to Output	0.6	3.6	0.6	3.7	ns
t _{ESYNC}	RCLK to $\overline{\text{PAE}}$ Bus Sync to Output	0.6	3.6	0.6	3.7	ns
t _{EXO}	RCLK to $\overline{\text{PAE}}$ Bus Expansion to Output	0.6	3.6	0.6	3.7	ns
t _{PR}	RCLK to Packet Ready Flag	0.6	3.6	0.6	3.7	ns
t _{SKEW1}	SKEW time between RCLK and WCLK for $\overline{\text{FF}}$ and $\overline{\text{OV}}$	4	—	4.5	—	ns
t _{SKEW2}	SKEW time between RCLK and WCLK for $\overline{\text{PAF}}$ and $\overline{\text{PAE}}$	5	—	6	—	ns
t _{SKEW3}	SKEW time between RCLK and WCLK for $\overline{\text{PAF}}[0:7]$ and $\overline{\text{PAE}}[0:7]$	5	—	6	—	ns
t _{SKEW4}	SKEW time between RCLK and WCLK for $\overline{\text{PR}}$ and $\overline{\text{OV}}$	5	—	6	—	ns
t _{SKEW5}	SKEW time between RCLK and WCLK for $\overline{\text{OV}}$ when in Packet Ready Mode	8	—	10	—	ns
t _{XIS}	Expansion Input Setup	1.0	—	1.0	—	ns
t _{XIH}	Expansion Input Hold	0.5	—	0.5	—	ns

NOTES:

1. Industrial temperature range product for the 6ns is available as a standard device. All other speed grades available by special order.
2. Values guaranteed by design, not currently tested.

IMPORTANT NOTICE AND DISCLAIMER

RENESAS ELECTRONICS CORPORATION AND ITS SUBSIDIARIES (“RENESAS”) PROVIDES TECHNICAL SPECIFICATIONS AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES “AS IS” AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT OF THIRD-PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for developers who are designing with Renesas products. You are solely responsible for (1) selecting the appropriate products for your application, (2) designing, validating, and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, or other requirements. These resources are subject to change without notice. Renesas grants you permission to use these resources only to develop an application that uses Renesas products. Other reproduction or use of these resources is strictly prohibited. No license is granted to any other Renesas intellectual property or to any third-party intellectual property. Renesas disclaims responsibility for, and you will fully indemnify Renesas and its representatives against, any claims, damages, costs, losses, or liabilities arising from your use of these resources. Renesas' products are provided only subject to Renesas' Terms and Conditions of Sale or other applicable terms agreed to in writing. No use of any Renesas resources expands or otherwise alters any applicable warranties or warranty disclaimers for these products.

(Disclaimer Rev.1.01)

Corporate Headquarters

TOYOSU FORESIA, 3-2-24 Toyosu,
Koto-ku, Tokyo 135-0061, Japan
www.renesas.com

Trademarks

Renesas and the Renesas logo are trademarks of Renesas Electronics Corporation. All trademarks and registered trademarks are the property of their respective owners.

Contact Information

For further information on a product, technology, the most up-to-date version of a document, or your nearest sales office, please visit www.renesas.com/contact-us/.