

What is a bank-switchable dual-port?

A bank-switchable dual-port (or BSDP) is an innovative multi-port memory solution that allows simultaneous access to common memory from two separate ports at higher storage densities and lower cost-per-bit than traditional dual-ports. In order to support larger memory arrays without increasing cost and package size and sacrificing performance, the device relies on SRAM core memory architecture with additional internal logic, instead of the more complex traditional dual-port architecture.

Perhaps the simplest way to think of the BSDP is as a single-chip multiplexed SRAM solution, but with the advantage of multiple banks of memory available. Both ports are able to operate simultaneously in separate banks, doubling the effective bandwidth of a typical multiplexed SRAM where only one port can access the memory at a time.


What is the difference between a bank-switchable dual-port SRAM and IDT's standard (or traditional) dual-ports?

To understand this, we must first consider the difference between a dual-port SRAM and a single-port (or

standard) SRAM. Figure1 depicts the typical cell architecture for each type of memory: the dual-port cell requires extra transistors in order to support the second port's row and bit lines, making the value stored in that cell available to both ports.

Since it has a more complex architecture, the dual-port cell occupies more space, which means that the overall die size for a given amount of storage capacity will be inherently larger than that of a corresponding standard SRAM. As we increase the storage densities of our dual-port products, this larger die will typically increase the cost of manufacturing the device and will also increase the size of the package for the die.

For those applications where high performance and a large amount of memory are needed but where board space and system cost are at a premium, IDT developed the bank-switchable architecture in order to provide the highest possible density of storage at the best possible performance while offering a very low cost-per-bit. Instead of using a dual-port cell, the BSDP uses the standard SRAM cell, and then provides additional logic in order to permit both ports to operate simultaneously in separate banks within the SRAM array. The BSDP ends up sharing most of the key features and benefits of true dual-port devices, as shown in Figure 2.



5646 drw 01

Figure 1. Comparison of Cell Architectures




Figure 2. Comparison of BSDP and Standard Dual-Port


Both devices provide simultaneous access to a common memory array: in the case of the BSDP, the simultaneous accesses occur to separate banks, while in the standard dual-port the simultaneous access can take place at the cell level.

What is the difference between a bank-switchable dual-port SRAM device and a multiplexed SRAM solution?

As a single-chip solution, the BSDP offers many advantages when compared to a multiplexed SRAM solution.

Figure 3 below compares the two types of solution using synchronous components (similar advantages are also available using asynchronous devices). The BSDP, as a single chip solution, typically occupies less board space and provides a reduced component count. This can also mean a shorter design cycle for the customer, with a corresponding improvement in time-to-market.

Perhaps the most significant advantage of the BSDP over a normal multiplexed solution is the improvement in performance. With two separate ports, each with a separate and independent clock input, the BSDP can execute simultaneous accesses at maximum frequency on both busses – at 200 MHz on each port in a 36-bit device,



5646 drw 03

Figure 3. Comparison of Multiplexed SRAM and BSDP

this equates to 14.4 Gb/sec of bandwidth. The conventional multiplexed SRAM, with the inherent derating associated with the additional logic registers and buffers, may occasionally have to sacrifice access cycles in order to support switching between the busses, and so may run at an effective frequency less than the rated 200 MHz. The simultaneous access provided by the BSDP effectively doubles the bandwidth of the device compared to a multiplexed SRAM.

What determines the choice between a bank-switchable device and a traditional dual-port?

Ultimately, this decision is driven by the desired granularity in the customer's arbitration of the device. By arbitration, we mean the process of ensuring that the operations of the two ports do not conflict with one another (i.e., a port does not try to read something that is being written by the opposite port, or both ports do not try to write to the same address at the same time). In a great many applications, customers allow for arbitration simply by sectioning the memory into logical blocks based on address space, then allowing a particular block of addresses to be accessed by only one port at a time, so as to preclude any chance of collision between the ports. This arbitration, sometimes referred to as "ping-ponging" or "paging", may be administered through an external FPGA or ASIC, or handled within system software. If the application never allows both ports to access the same address at the same time, then the bank-switchable architecture is likely the ideal candidate, offering higher density and lower cost-per-bit than the traditional dual-port, while continuing to support simultaneous access from both ports. In the case that the application needs to allow for simultaneous access from both ports to the same address (this would typically be valid only if both ports were reading), then the traditional dual-port is the only viable choice – it alone has the extra circuitry at cell level to support the simultaneous access. The greater complexity of the cell imposes a restriction on the density of storage and a higher cost-per-bit compared to the BSDP.

What controls the arbitration of the banks?

This differs between the async and sync bank-switchable devices. In the async devices, there are 4 banks internally, and there are 4 pins on the device designated as "Bank Select" pins. Each of these pins is associated with a specific bank. Inputting a logic "LOW" on a given pin assigns its associated bank to the left port, while setting it to a logic "HIGH" assigns the associated bank to the right port. The pins are typically controlled by an external controller (often one of the devices connected to the dual-port) that has responsibility for arbitrating access requests for the two ports. Two of the address pins on each port are

designated for bank access, while the others specify the address within the bank. On the async device, we have special mailbox functions implemented to facilitate communication between the ports in order to manage the access requests as efficiently as possible.

On the sync devices, we set up the internal multiplexing architecture so that six of the address pins on each port control the bank access, while the rest of the address pins control the access within the specific bank. This means that accessing a given bank is as simple as asserting the appropriate bank address, and there is no need for an external controller as in the async devices. Because of the focus on supporting high frequency operations for the sync devices, there is no on-chip arbitration or special functions such as mailboxes – these functions typically reduce the operating speed of a device.

The datasheets for both types of device discuss the specifics of controlling bank accesses in much greater detail – please refer to these documents, available on the IDT website at <http://www.idt.com>.

What types of bank-switchable solutions does IDT support?

Table 1 indicates the range of BSDP solutions available from IDT. These include a range of densities, interface types, bus widths, operating voltages, speeds, and package options.

What packages are the bank-switchable devices available in?

Table 2 indicates the appropriate package information for the package codes listed in Table 1.

When using a bank-switchable device, what happens if both ports try to access the same bank at the same time?

This varies slightly for the async and the sync devices. For the async devices, in the event that a port asserts a bank address for a bank that has not been assigned to it via the bank select pin, the access is invalid (no data is written, or in the case of a read, random data is output). The other port, to whom the bank has been allocated, is able to proceed without interruption.

For the sync devices, if both ports assert the same bank address at the same time, then each interferes with the other, with the end result that both accesses are invalid (random data is written, or in the case of a read, random data is output).

Width	Memory Array	Number of Banks	Bank Config (each)	Array Density (# Bits)	Part Number and Type	Core Voltage (V)	I/O Voltage (V)	Max Speed	Packages
x36	128Kx36	64	2Kx36	4M	70V7599 (sync)	3.3	3.3 or 2.5 selectable per port	200MHz	BF-208 BC-256 DR-208
	256Kx36	64	4Kx36	9M	70V7519 (sync)	3.3	3.3 or 2.5 selectable per port	200MHz	BF-208 BC-256 DR-208
x18	256Kx18	64	4Kx18	4M	70v7319 (sync)	3.3	3.3 or 2.5 selectable per port	200MHz	BF-208 BC-256 DD-144
	512Kx18	64	8Kx18	9M	70V7339 (sync)	3.3	3.3 or 2.5 selectable per port	200MHz	BF-208 BC-256 DD-144
x16	64Kx16	4	16Kx16	1M	707288 (async)	5	5	15ns	PN-100
	64Kx16	4	16Kx16	1M	70V7288 (async)	3.3	3.3	15ns	PN-100
	32Kx16	4	8Kx16	512K	707278 (async)	5	5	15ns	PN-100
	32Kx16	4	8Kx16	512K	70V7278 (async)	3.3	3.3	15ns	PN-100

5646 tbl 01

Table 1. Configuration Options

Pin	Type	PKG	Pitch (mm)	Area (mm ²)	Length (mm)	Width (mm)	Height (mm)
208	fpBGA	BF-208	0.8	225	15	15	1.4
256	BGA	BC-256	1.0	289	17	17	1.4
208	PQFP	DR-208	0.5	784	28	28	3.5
144	TQFP	DD-144	0.5	400	20	20	1.4
100	TQFP	PN-100	0.5	196	14	14	1.4

5646 tbl 02

Table 2. Package Information

What are some other advantages of IDT's bank-switchable devices?

- ❑ IDT's x36 and x18 sync BSDPs share common packages and footprints with our sync traditional dual-ports, allowing for 5 levels of density within the same footprint. This permits the greatest flexibility in supporting upgrades or in offering different board configurations with varying amounts of memory without significant design revisions.
- ❑ IDT offers JTAG (compliant with IEEE 1149.1) in the fpBGA(BF-208), BGA(BC-256), and PQFP(DR-208) packages in order to support debug and test operations and improve manufacturability and board yield.
- ❑ IDT's BSDPs offer selectable 3.3V and 2.5V operations on the I/Os for the sync devices, supporting the ability to share data between busses at different voltages.
- ❑ IDT's devices offer x9 byte controls on each port, allowing customers to share data between busses at different widths.

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 Jan 2024)

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/.