


Nexys 3™ FPGA Board Reference Manual

Revised April 10, 2013

This manual applies to the Nexys 3 rev. B

Overview

The Nexys 3 is a complete, ready-to-use digital circuit development platform based on the Xilinx Spartan-6 LX16 FPGA. The Spartan-6 is optimized for high performance logic, and offers more than 50% higher capacity, higher performance, and more resources as compared to the Nexys 2's Spartan-3 500E FPGA.



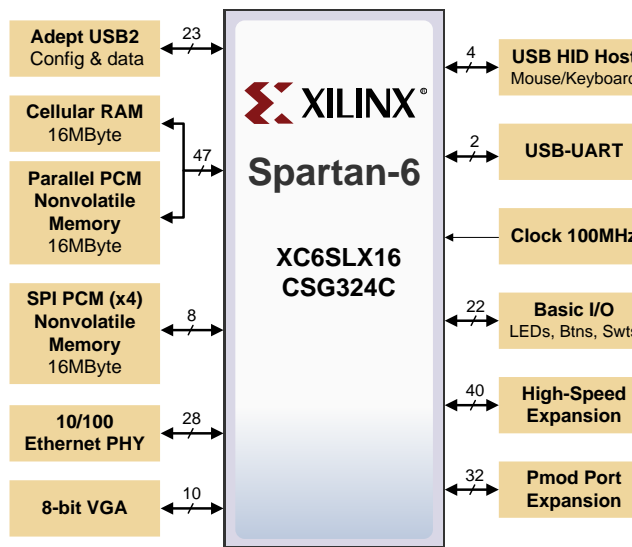
Features include:

- Xilinx Spartan-6 LX16 FPGA in a 324-pin BGA package
- 16Mbyte Cellular RAM (x16)
- 16Mbytes SPI (quad mode) PCM non-volatile memory
- 16Mbytes parallel PCM non-volatile memory
- 10/100 Ethernet PHY
- On-board USB2 port for programming & data xfer
- USB-UART and USB-HID port (for mouse/keyboard)
- 8-bit VGA port
- 100MHz CMOS oscillator
- 72 I/O's routed to expansion connectors
- GPIO includes 8 LEDs, 5 buttons, 8 slide switches and 4-digit seven-segment display
- USB2 programming cable included

In addition to the Spartan-6 FPGA, the Nexys 3 offers an improved collection of peripherals including 32Mbytes of Micron's latest Phase Change nonvolatile memory, a 10/100 Ethernet PHY, 16Mbytes of Cellular RAM, a USB-UART port, a USB host port for mice and keyboards, and an improved high-speed expansion connector. The large FPGA and broad set of peripherals make the Nexys 3 board an ideal host for a wide range of digital systems, including embedded processor designs based on Xilinx's MicroBlaze.

Nexys 3 is compatible with all Xilinx CAD tools, including ChipScope, EDK, and the free WebPack. The Nexys 3 uses Digilent's newest Adept USB2 system that offers FPGA and ROM programming, automated board tests, virtual I/O, and simplified user-data transfer facilities.

A comprehensive collection of board support IP and reference designs, and a large collection of add-on boards are available on the Digilent website. Please see the Nexys 3 page at www.digilentinc.com for more information.

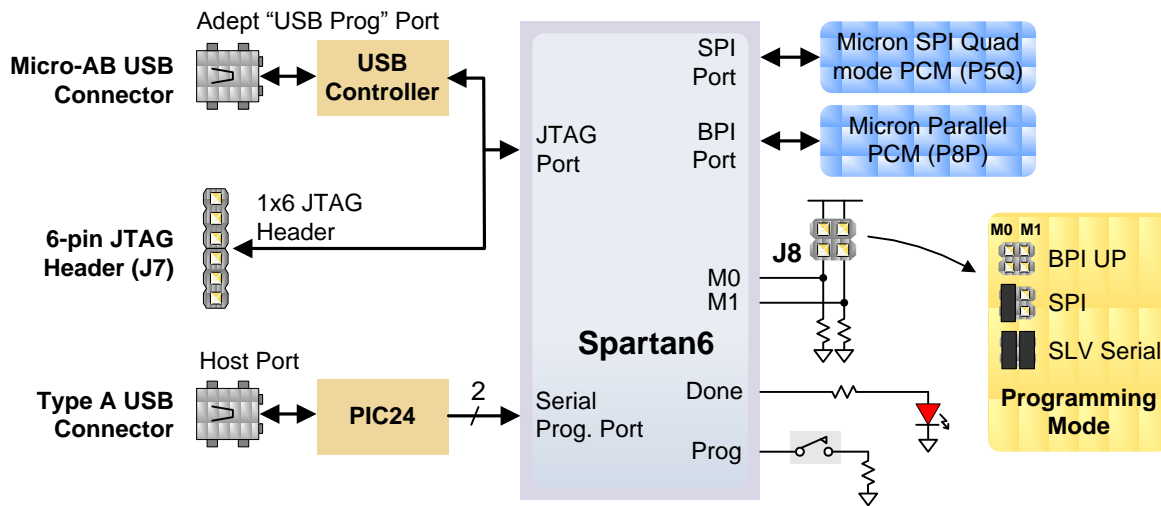


Spartan-6 LX16 features include:

- 2,278 slices each containing four 6-input LUTs and eight flip-flops
- 576Kbits of fast block RAM
- two clock tiles (four DCMs & two PLLs)
- 32 DSP slices
- 500MHz+ clock speeds

1 Configuration

After power-on, the Spartan-6 FPGA board must be configured (or programmed) before it can perform any functions. The FPGA can be configured in one of four ways: a PC can use the Adept "USB Prog" port to program the FPGA any time power is on; a configuration file stored in the non-volatile parallel PCM device can be transferred to the FPGA at power-on using the BPI-UP port; a file stored in the non-volatile serial (SPI) PCM device can be transferred to the FPGA using the SPI port; or a programming file can be transferred from a USB memory stick attached to the USB HID port. An on-board "mode" jumper (J8) selects between the programming modes as shown in the J8 Mode legend in the figure below. JTAG Mode can be accessed at any time without changing jumpers.



Programming files are stored in SRAM-based memory cells within the FPGA. This data defines the FPGA's logic functions and circuit connections, and it remains valid until it is erased by removing board power, by pressing the reset button attached to the PROG input, or by writing a new configuration file using the JTAG port.


FPGA configuration files transferred via the JTAG port use the .bin or .svf file types, files transferred from a USB stick use the .bit file type, and BPI or SPI programming files can use .bit, .bin, or .mcs types. The ISE/WebPack or EDK software from Xilinx can create bit, svf, bin, or mcs files from VHDL, Verilog, or schematic-based source files

(EDK is used for MicroBlaze™ embedded processor-based designs). Digilent's Adept software or Xilinx's iMPACT software can be used to program the FPGA or ROMs using the Adept USB port.

During JTAG programming, a .bit or .svf file is transferred from the PC to the FPGA using the Adept USB port. When programming a non-volatile PCM device, a .bit, .bin, or .mcs file is transferred to the in a two-step process. First, the FPGA is programmed with a circuit that can program PCM devices, and then data is transferred to the PCM device via the FPGA circuit (this complexity is hidden from the user – a simple "program ROM" interface is presented by the programming software. Note the PCM devices are next-generation Flash ROM devices, and they are often referred to as "Flash" or "ROM" memory). After the PCM device has been programmed, it can automatically configure the FPGA at a subsequent power-on or reset event as determined by the J8 jumper setting. Programming files stored in the PCM devices will remain until they are overwritten, regardless of power-cycle events.

The FPGA can be programmed from a memory stick attached to the USB-HID port if the stick contains a single .bit configuration file in the root directory, the J8 Programming Mode jumper is set to JTAG (both jumpers loaded), and board power is cycled. The FPGA will automatically reject any .bit files that are not built for the proper FPGA.

After being successfully programmed, the FPGA will cause the "Done" LED to illuminate. Pressing the Reset button at any time will reset the configuration memory in the FPGA. After being reset, the FPGA will immediately attempt to reprogram itself from one of the PCM devices if the J8 Mode jumper is set to BPI or SPI mode.



Digilent's Adept software offers a simplified programming interface and many additional features as described below. The Adept USB port is fully compatible with all Xilinx tools, including the iMPACT programming software. The Adept features are always available, regardless of how the FPGA was programmed.

1.1 Adept System


Digilent's Adept high-speed USB2 system can be used to program the FPGA and PCM devices, run automated board tests, add PC-based virtual I/O devices (like buttons, switches, and LEDs) to FPGA designs, and exchange register-based and file-based data with the FPGA. Adept automatically recognizes the Nexys 3 board and presents a graphical interface with tabs for each of these applications. Adept also includes public APIs/DLLs so that users can write applications to exchange data with the Nexys 3 board at up to 38Mbytes/sec. The Adept application, an SDK, and reference materials are freely downloadable from the Digilent website.

1.2 Programming Interface

To program the Nexys 3 board using Adept, first set up the board and initialize the software:

- plug in and attach the power supply
- plug in the USB cable to the PC and to the USB port on the board
- start the Adept software
- turn ON Nexys 3's power switch
- wait for the FPGA to be recognized.

Use the browse function to associate the desired .bit file with the FPGA, and click on the Program button. The configuration file will be sent to the FPGA, and a dialog box will indicate whether programming was successful. The configuration "done" LED will light after the FPGA has been successfully configured.



Before starting the programming sequence, Adept ensures that any selected configuration file contains the correct FPGA ID code – this prevents incorrect .bit files from being sent to the FPGA.

In addition to the navigation bar and browse and program buttons, the Config interface provides an Initialize Chain button, console window, and status bar. The Initialize Chain button is useful if USB communications with the board have been interrupted. The console window displays current status, and the status bar shows real-time progress when downloading a configuration file.

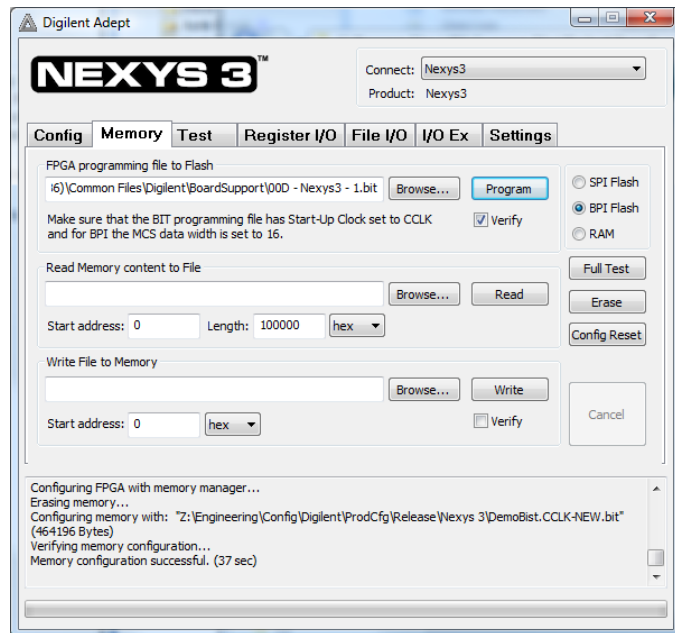
1.3 Memory Interface

The Memory tab allows .bin, .bit, and .mcs configuration files to be transferred to the on-board BPI (parallel) or SPI (serial) PCM devices for subsequent FPGA programming, and allows user data files to be transferred to/from the PCM devices or RAM memories at user-specified addresses. The target memory is selected by clicking one of the three radio buttons in the upper-right corner.

The configuration tool supports programming from any valid ROM file produced by the Xilinx tools. After programming, board power can either be cycled or the Reset button can be pressed to program the FPGA from the PCM device selected by the J8 mode jumper. If programming with a .bit file, the startup clock must be set to CCLK.

All three memory devices (the PCM's and the cellular RAM) can be fully tested by clicking the Full Test button. They can also be completely erased by clicking the Erase button.

The Read/Write tools allow data to be exchanged between files on the host PC and specified address ranges in the memory devices.




1.4 Test Interface

The test interface provides an easy way to verify many of the board's hardware circuits and interfaces. These are divided into two major categories: on-board memory (RAM and Flash) and peripherals. In both cases, the FPGA is configured with test and PC-communication circuits, overwriting any FPGA configuration that may have been present.

Clicking the Run RAM/Flash Test button will identify the CellularRam, SPI Flash, and BPI Flash memory by reading out and verifying the IDCODE on each memory. The memory contents will not be modified. To run a full test on a particular memory device, refer to the Full Test in the Memory Tab.

Clicking the Start Peripherals Test button will initialize GPIO and user I/O testing. Once the indicator near the Start Peripherals Test button turns green, all peripheral tests can be run.

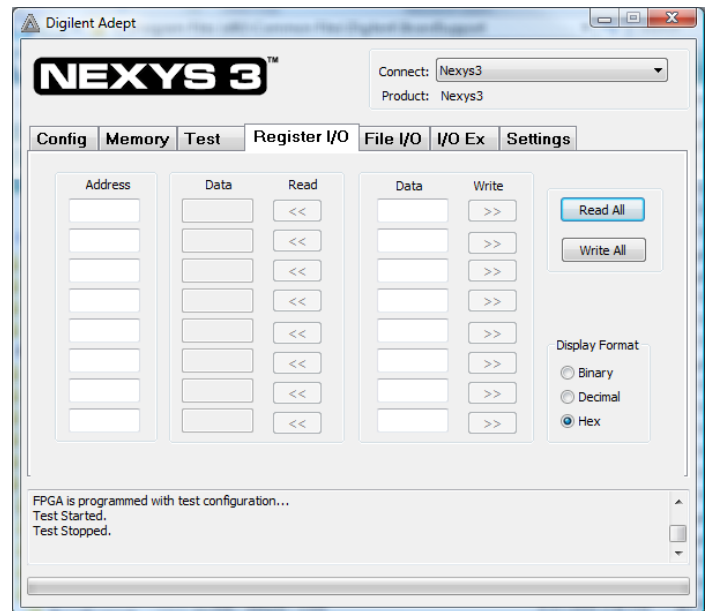
The Test Shorts feature checks all discrete I/O's for shorts to Vdd, GND, and neighboring I/O pins. The switches and buttons graphics show the current states of those devices on the Nexys 3 board. Connect a VGA monitor and USB mouse to visually test the J2 VGA port and J4 USB port respectively.



1.5 Register I/O

The register I/O tab requires that a corresponding IP block, available in the Parallel Interface reference design (DpimRef.vhd) on the Adept page of the Digilent website, is included and active in the FPGA. This IP block provides an EPP-style interface, where an 8-bit address selects a register, and data read and write buttons transfer data to and from the selected address. Addresses entered into the address field must match the physical address included in the FPGA IP block.

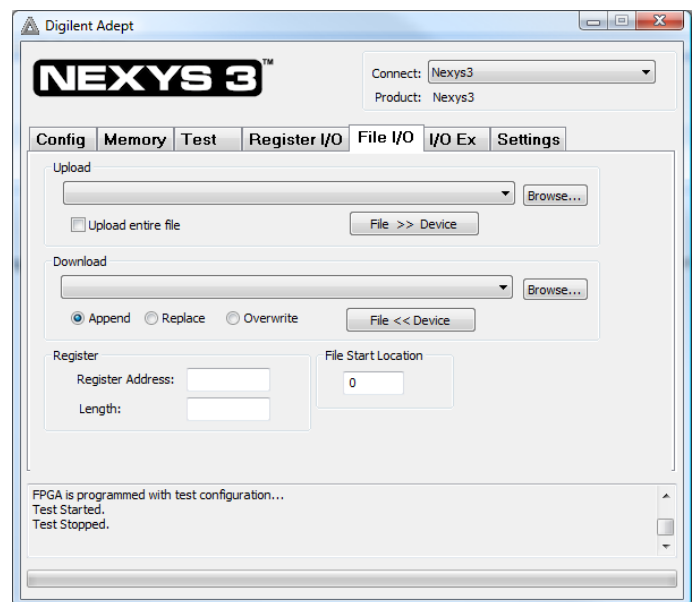
Register I/O provides an easy way to move small amounts of data into and out of specific registers in a given design. This feature greatly simplifies passing control parameters into a design, or reading low-frequency status information out of a design.



1.6 File I/O

The File I/O tab can transfer files between the PC and the Nexys 3 FPGA. A number of bytes (specified by the Length value) can be streamed into a specified register address from a file or out of a specified register address into a file. During upload and download, the file start location can be specified in terms of bytes.

As with the Register I/O tab, File I/O also requires specific IP to be available in the FPGA. This IP can include a memory controller for writing files into the on-board Ram and Flash memories.

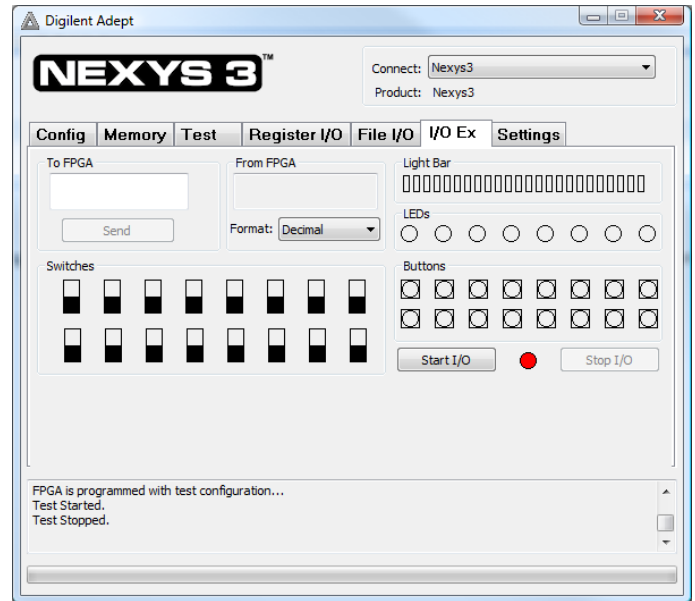


1.7 I/O Expand

The I/O Expand tab works with an IP block in the FPGA to provide additional simple I/O beyond the physical devices found on the Nexys 3 board. Virtual I/O devices include a 24-LED light bar, 16 slide switches, 16 push buttons, 8 discrete LEDs, a 32-bit register that can be sent to the FPGA, and a 32-bit register that can be read from the FPGA.

The IP block, available in the Adept I/O Expansion reference design (AdeptIOExpansion.zip) on the Adept page of the Digilent website, provides a simple interface with well-defined signals. This IP block can easily be included in, and accessed from, user-defined circuits.

For more information, see the Adept documentation available at the Digilent website.



2 Power Supplies

The Nexys 3 board can receive power from the Adept USB port or from an external power supply. Jumper JP1 (near the power jack) determines which source is used.


The USB port can deliver enough power for the vast majority of designs. It is possible that a very demanding application, including an application that drives many peripheral boards, might require more power than can be delivered by the USB port. Some applications may also need to run without being connected to a PC's USB port. In either case, an external power supply or battery pack can be used by setting JP1 to "Wall".

The main regulator on the Nexys 3 can accommodate input voltages up to 5.5VDC. An external DC wall-plug supply should provide at least five watts of input power, and use a coax center-positive 2.1mm internal-diameter plug. An external battery pack can also be used by connecting the battery's terminal leads to connector J11 (J11 is in parallel with the wall-plug power jack, so if a battery pack is connected, a wall plug should not be). An external battery pack should also be limited to 5.5VDC, and should be capable of delivering adequate power for the application.

Voltage regulator circuits from Linear Technology create the required 3.3V, 2.5V, 1.8V, and 1.2V supplies from the main power input. The table below provides additional information (typical currents depend strongly on FPGA configuration and the values provided are typical of medium size/speed designs).

Supply	Circuits	Device	Amps (max/typ)
3.3V	FPGA I/O, USB ports, Clocks, ROM & RAM I/O, Ethernet	IC13: LTC3633	3A / 200mA
2.5V	Optional voltage for Bank0 and VHDC connector	IC14: LTC3619	800mA / 0mA
1.2V	FPGA Core	IC13: LTC3633	3A / 0.2 to 1.0A
1.8V	RAM and ROM core	IC14: LTC3619	400mA / 0.1 to 0.3A

Table 1. Nexys 3 Power Supplies.



The Nexys 3 power supplies are enabled (or turned on) by a logic-level Power switch (SW8). A power-good LED (LD8), driven by the wired-OR of the "power good" outputs on the supplies, indicates that all supplies are operating within 10% of nominal.

The VU output of the main power jumper (JP1) is available to the VHDC expansion connector if jumper JP4 is loaded. Care must be taken to ensure the VUEXP delivered to any attached expansion board is the correct voltage – since VU is driven directly from an attached supply, this means a supply of the proper voltage must be used (e.g., 5V).

3 Memory

The Nexys 3 board contains three external memories, all from Micron: a 128Mbit Cellular RAM (pseudo-static DRAM); a 128Mbit parallel non-volatile PCM (phase-change memory); and a 128Mbit serial PCM device. The Cellular RAM and parallel PCM device share a common bus, and the serial PCM is on a dedicated quad-mode (x4) SPI bus. The non-volatile PCM memories are byte and bit alterable without requiring a block erase, so they are faster and more versatile than conventional Flash in most applications.

The 16Mbyte Cellular RAM (Micron part number M45W8MW16) has a 16-bit bus that supports 8 or 16 bit data access. It can operate as a typical asynchronous SRAM with read and write cycle times of 70ns, or as a synchronous memory with an 80MHz bus. When operated as an asynchronous SRAM, the Cellular RAM automatically refreshes its internal DRAM arrays, allowing for a simplified memory controller (similar to any SRAM controller). When operated in synchronous mode, continuous transfers of up to 80MHz are possible.

The parallel PCM device (Micron part number NP8P128A13T1760E) is organized as 8Mwords of 16bits each. It contains 128 individually erasable 64K-blocks, one of which is subdivided into four 16K "parameter blocks" that can offer increased data protection. Normally, a device with a parameter block at the high end of the address space is loaded (a "top parameter" block device). The parallel PCM memory offers 115ns read cycle times, with 25ns page-mode reads within blocks. It has an internal 64-byte write buffer that can be written with 50ns cycle times, and the 64-byte buffer can be transferred to the Flash array in 120us (typical). The parallel PCM device also contains an SPI

port for serial data transfer, but that function is not enabled on the Nexys 3 board (the ADV and WAIT signals are connected between the FPGA and PCM, but they serve no functions).

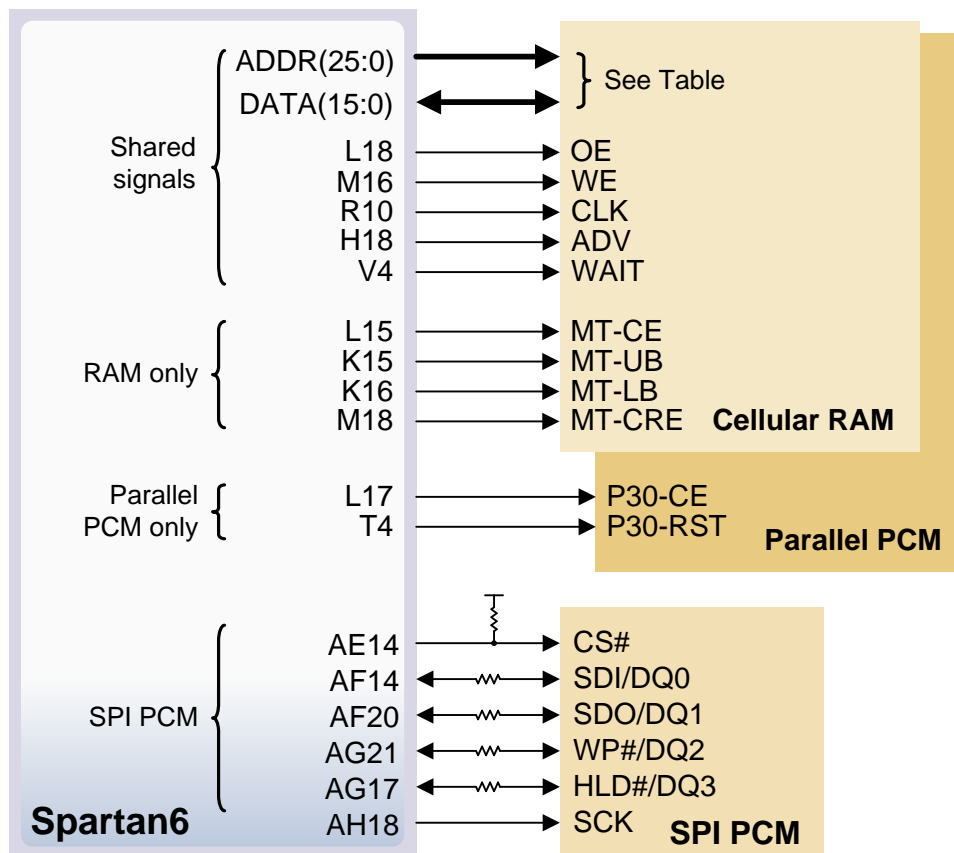
The Cellular RAM and parallel PCM share a common 16-bit data bus and 24-bit address bus. The Cellular RAM is byte addressable using the upper-byte and lower-byte signals (MT-UB and MT-LB), but the P8P PCM is configured for 16 byte operations only (it is not byte addressable). The output enable (OE) and write enable (WE) signals are shared by both devices, but each device has individual chip enable (CE) signals. Additionally, the Cellular RAM has clock (MT-CLK), wait (MT-WAIT), address valid (MT-ADV) and control register enable (MT_CRE) signals available to the FPGA for use with synchronous transfers, and the PCM device has Reset (RP#). With simple changes, the Nexys 3 board can accommodate either the PCM device or the older "P33" Flash device. The signal names in the Nexys 3 schematic reference the P33 device; the P33 ADV and WAIT signals are not used in by the PCM device.

The 16Mbyte serial PCM device (Micron part number NP5Q128A13ESFC0E) is also bit alterable without requiring an erase cycle. It supports the legacy SPI protocol as well as the newer Quad I/O and Dual I/O protocols, at bus speeds up to 50MHz.

FPGA configuration files can be written to both PCM devices, and mode settings are available to cause the FPGA to automatically read a configuration from one of these devices at power on. A Spartan-6 LX16 configuration file requires about 512Kbytes of memory, leaving about 97% of the PCM devices available for user data.

Both PCM devices are loaded with configuration files at the factory. The SPI PCM device contains a file that configures the FPGA to test the Nexys 3 memory devices during manufacturing, and this file isn't needed after the board test is complete. The BPI PCM device contains a file that configures the Nexys 3 with a basic user demonstration program, and it can be used to verify board functions. This same .bit file is available for download from the Digilent website. If Mode jumper J8 is set to BPI mode and power is applied, the user demo configuration will be loaded. The demo drives a counter on the 7-segment display, drives the user LEDs on and off when the user switches are toggled, turns off digits on the 7-segment display when user buttons are pressed, and drives an image out the VGA port. A USB mouse can be connected to J4 for a simple visual demonstration.

Please refer to the manufacturer's data sheets and the reference designs posted on Digilent's website for more information about the memory devices.



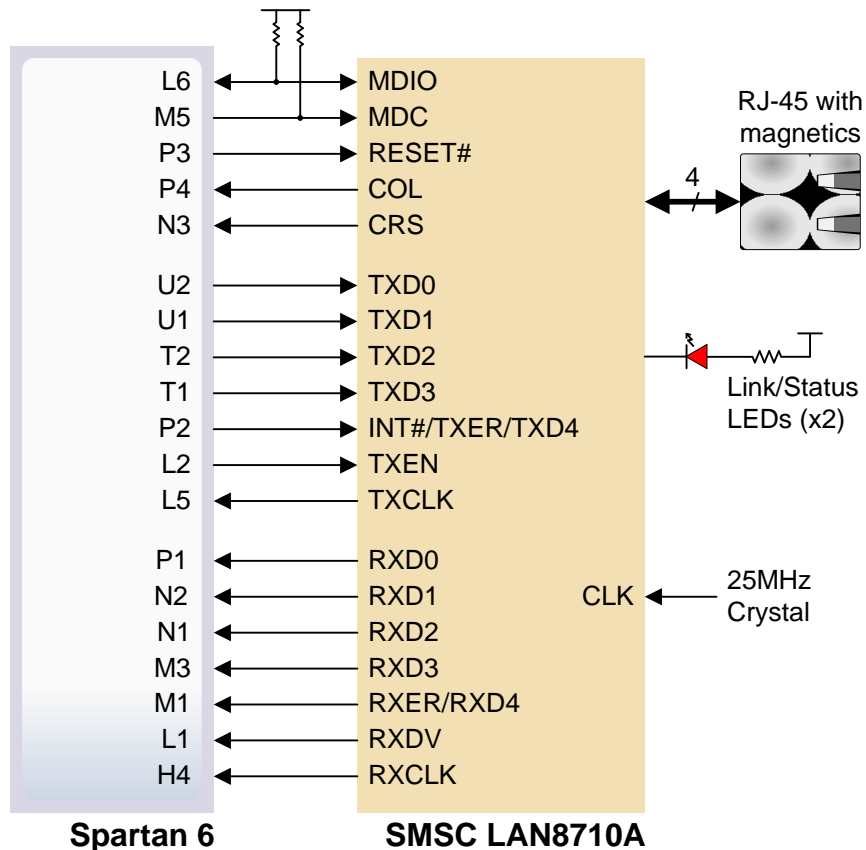
Address Bus			Data Bus	
ADDR25: F15	ADDR16: G13	ADDR7: H15	DATA15: T8	DATA6: T3
ADDR24: F16	ADDR15: E16	ADDR6: H16	DATA14: R8	DATA5: R3
ADDR23: C17	ADDR14: E18	ADDR5: G16	DATA13: U10	DATA4: V5
ADDR22: C18	ADDR13: K12	ADDR4: G18	DATA12: V13	DATA3: U5
ADDR21: F14	ADDR12: K13	ADDR3: J16	DATA11: U13	DATA2: V14
ADDR20: G14	ADDR11: F17	ADDR2: J18	DATA10: P12	DATA1: T14
ADDR19: D17	ADDR10: F18	ADDR1: K17	DATA9: P6	DATA0: R13
ADDR18: D18	ADDR9: H13	ADDR0: K18	DATA8: N5	
ADDR17: H12	ADDR8: H14		DATA7: R5	

4 Ethernet PHY

The Nexys 3 board includes an SMSC 10/100 Ethernet PHY (SMSC part number LAN8710) paired with an RJ-45 Ethernet jack with integrated magnetics. EDK-based designs can access the PHY using either the `xps_ethernetlite` IP core or the `xps_ll_temac` IP core. The Nexys 3 Base System Builder (BSB) support package automatically generates a test application for the Ethernet MAC; this can be used as a reference for creating custom designs. ISE designs can use the IP Core Generator wizard to create an Ethernet MAC controller IP core. The SMSC PHY uses the MII interface and supports 10/100 Mb/s. At power-on reset, the PHY is set to the following defaults:

- MII mode interface
- 100Base-TX with auto negotiation enabled, advertising half-duplex, with CRS active during receive
- PHY address = 000

Refer to the LAN8710A data sheet on the www.smsc.com website for further information.



5 Oscillators/Clocks

The Nexys 3 board includes a single 100MHz CMOS oscillator connected to pin V10 (V10 is the GCLK0 input in bank 2). The input clock can drive any or all of the four clock management tiles in the Spartan-6. Each tile includes two Digital Clock Managers (DCMs) and four Phase-Locked Loops (PLLs).


DCMs provide the four phases of the input frequency (0°, 90°, 180°, and 270°), a divided clock that can be the input clock divided by any integer from 2 to 16 or 1.5, 2.5, 3.5... 7.5, and two antiphase clock outputs that can be multiplied by any integer from 2 to 32 and simultaneously divided by any integer from 1 to 32.

PLLs use VCOs that can be programmed to generate frequencies in the 400MHz to 1080MHz range by setting three sets of programmable dividers during FPGA configuration. VCO outputs have eight equally-spaced outputs (0°, 45°, 90°, 135°, 180°, 225°, 270°, and 315°) that can be divided by any integer between 1 and 128.

Please refer to the Spartan-6 data sheet at www.xilinx.com for further information on the clock management tiles.

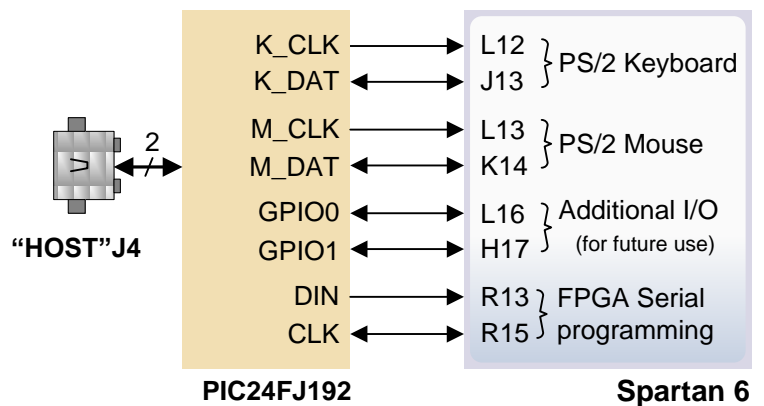
6 USB-UART Bridge (Serial Port)

The Nexys 3 includes an FTDI FT232 USB-UART bridge to allow PC applications to communicate with the board using standard Windows COM port commands. Free USB-COM port drivers, available from www.ftdichip.com under the "Virtual Com Port" or VCP heading, convert USB packets to UART/serial port data. Serial port data is exchanged with the FPGA using a two-wire serial port (TXD/RXD) and software flow control (XON/XOFF). After the drivers are installed, I/O commands from the PC directed to the COM port will produce serial data traffic on the N17 and N18 FPGA pins.



7 USB HID Host

A Microchip PIC24FJ192 microcontroller provides the Nexys 3 with USB H ID host capability. Firmware in the microcontroller can drive a mouse or a keyboard attached to the type A USB connector at J4 labeled "Host". Hub support is not currently available, so only a single mouse or a single keyboard can be used. The PIC24 drives four signals into the FPGA – two are used as a keyboard port following the keyboard PS/2 protocol, and two are used as a mouse port following the mouse PS/2 protocol.




Two PIC24 I/O pins are also connected to the FPGA's two-wire serial programming port, so the FPGA can be programmed from a file stored on a USB memory stick. To program the FPGA, attach a memory stick containing a single .bit programming file in the root directory, load both M0 and M1 on J8 jumper, and cycle board power. This will cause the PIC processor to program the FPGA, and any incorrect bit files will automatically be rejected. Note the PIC24 reads the FPGA's mode, init, and done pins, and can drive the PROG pin as a part of the programming sequence.

7.1 HID Controller

To access the USB host controller, EDK designs can use the standard PS/2 core (non-EDK designs can use a simple state machine).

Mice and keyboards that use the PS/2 protocol¹ use a two-wire serial bus (clock and data) to communicate with a host device. Both use 11-bit words that include a start, stop, and



Symbol	Parameter	Min	Max
T _{CK}	Clock time	30us	50us
T _{SU}	Data-to-clock setup time	5us	25us
T _{HLD}	Clock-to-data hold time	5us	25us

¹Not all keyboard manufacturers strictly adhere to the PS/2 specifications; some keyboards may not produce the proper signaling voltages or use the standard communication protocols. Compatibility with the USB host may vary between different keyboards.

odd parity bit, but the data packets are organized differently, and the keyboard interface allows bi-directional data transfers (so the host device can illuminate state LEDs on the keyboard). Bus timings are shown in the figure. The clock and data signals are only driven when data transfers occur, and otherwise they are held in the idle state at logic '1'. The timings define signal requirements for mouse-to-host communications and bi-directional keyboard communications. A PS/2 interface circuit can be implemented in the FPGA to create a keyboard or mouse interface.

7.2 Keyboard

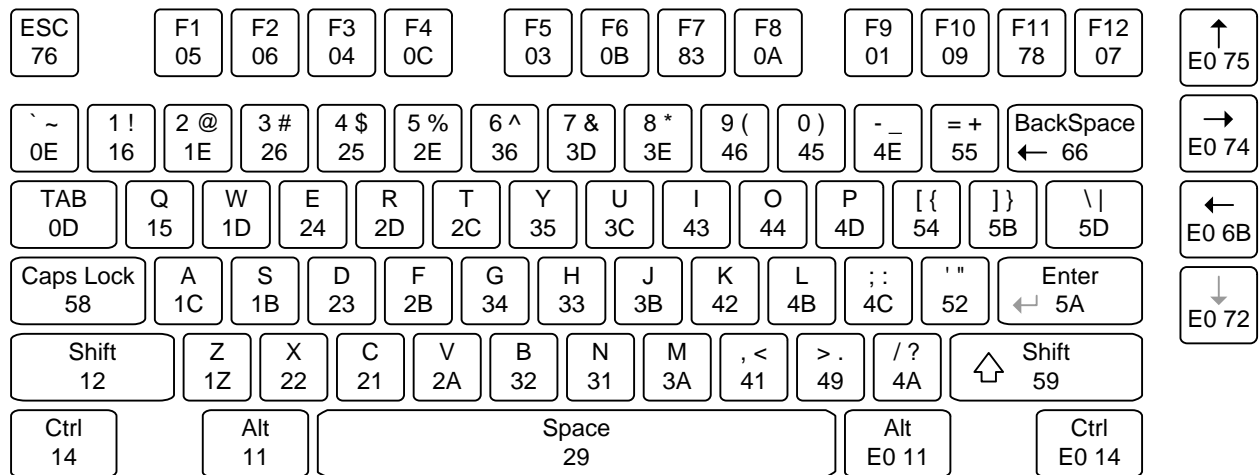
The keyboard uses open-collector drivers so the keyboard, or an attached host device, can drive the two-wire bus (if the host device will not send data to the keyboard, then the host can use input-only ports).

PS/2-style keyboards use scan codes to communicate key press data. Each key is assigned a code that is sent whenever the key is pressed. If the key is held down, the scan code will be sent repeatedly about once every 100ms. When a key is released, an F0 key-up code is sent, followed by the scan code of the released key. If a key can be shifted to produce a new character (like a capital letter), then a shift character is sent in addition to the scan code, and the host must determine which ASCII character to use. Some keys, called extended keys, send an E0 ahead of the scan code (and they may send more than one scan code). When an extended key is released, an E0 F0 key-up code is sent, followed by the scan code. Scan codes for most keys are shown in the figure. A host device can also send data to the keyboard. Below is a short list of some common commands a host might send.

- ED Set Num Lock, Caps Lock, and Scroll Lock LEDs. Keyboard returns FA after receiving ED, then host sends a byte to set LED status: bit 0 sets Scroll Lock, bit 1 sets Num Lock, and bit 2 sets Caps lock. Bits 3 to 7 are ignored.
- EE Echo (test). Keyboard returns EE after receiving EE.
- F3 Set scan code repeat rate. Keyboard returns F3 on receiving FA, then host sends second byte to set the repeat rate.
- FE Resend. FE directs keyboard to re-send most recent scan code.
- FF Reset. Resets the keyboard.

The keyboard can send data to the host only when both the data and clock lines are high (or idle). Since the host is the bus master, the keyboard must check to see whether the host is sending data before driving the bus. To facilitate this, the clock line is used as a "clear to send" signal. If the host pulls the clock line low, the keyboard must not send any data until the clock is released. The keyboard sends data to the host in 11-bit words that contain a '0' start bit, followed by 8-bits of scan code (LSB first), followed by an odd parity bit and terminated with a '1' stop bit. The keyboard generates 11 clock transitions (at 20 to 30KHz) when the data is sent, and data is valid on the falling edge of the clock.

Scan codes for most PS/2 keys are shown in the figure below.




PS/2 keyboard scan codes.

7.3 Mouse

The mouse outputs a clock and data signal when it is moved, otherwise, these signals remain at logic '1'. Each time the mouse is moved, three 11-bit words are sent from the mouse to the host device. Each of the 11-bit words contains a '0' start bit, followed by 8 bits of data (LSB first), followed by an odd parity bit, and terminated with a '1' stop bit. Thus, each data transmission contains 33 bits, where bits 0, 11, and 22 are '0' start bits, and bits 11, 21, and 33 are '1' stop bits. The three 8-bit data fields contain movement data as shown in the figure above. Data is valid at the falling edge of the clock, and the clock period is 20 to 30KHz.

The mouse assumes a relative coordinate system wherein moving the mouse to the right generates a positive number in the X field, and moving to the left generates a negative number. Likewise, moving the mouse up generates a positive number in the Y field, and moving down represents a negative number (the XS and YS bits in the status byte are the sign bits – a '1' indicates a negative number). The magnitude of the X and Y numbers represent the rate of mouse movement – the larger the number, the faster the mouse is moving (the XV and YV bits in the status byte are movement overflow indicators – a '1' means overflow has occurred). If the mouse moves continuously, the 33-bit transmissions are repeated every 50ms or so. The L and R fields in the status byte indicate Left and Right button presses (a '1' indicates the button is being pressed).



Mouse data format.

8 VGA Port

The Nexys 3 board uses 10 FPGA signals to create a VGA port with 8-bit color and the two standard sync signals (HS – Horizontal Sync, and VS – Vertical Sync). The color signals use resistor-divider circuits that work in conjunction with the 75-ohm termination resistance of the VGA display to create eight signal levels on the red and green VGA signals, and four on blue (the human eye is less sensitive to blue levels). This circuit, shown in figure 13, produces


video color signals that proceed in equal increments between 0V (fully off) and 0.7V (fully on). Using this circuit, 256 different colors can be displayed, one for each unique 8-bit pattern. A video controller circuit must be created in the FPGA to drive the sync and color signals with the correct timing in order to produce a working display system.

8.1 VGA System Timing

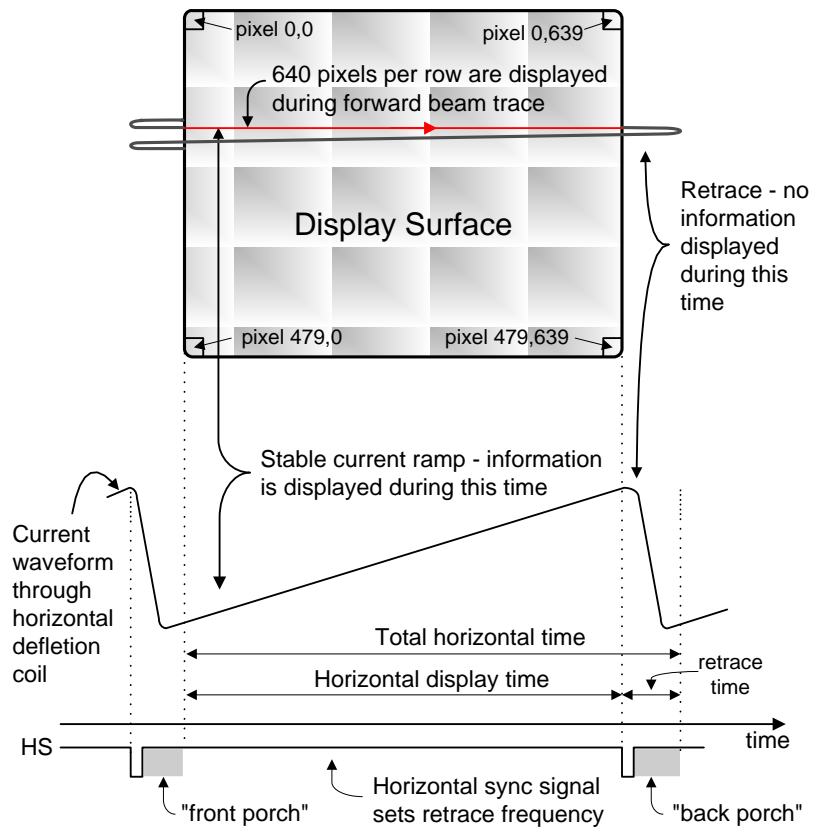
VGA signal timings are specified, published, copyrighted and sold by the VESA organization (www.vesa.org). The following VGA system timing information is provided as an example of how a VGA monitor might be driven in 640 by 480 mode. For more precise information, or for information on other VGA frequencies, refer to documentation available at the VESA website.

CRT-based VGA displays use amplitude-modulated moving electron beams (or cathode rays) to display information on a phosphor-coated screen. LCD displays use an array of switches that can impose a voltage across a small amount of liquid crystal, thereby changing light permittivity through the crystal on a pixel-by-pixel basis. Although the following description is limited to CRT displays, LCD displays have evolved to use the same signal timings as CRT displays (so the "signals" discussion below pertains to both CRTs and LCDs). Color CRT displays use three electron beams (one for red, one for blue, and one for green) to energize the phosphor that coats the inner side of the display end of a cathode ray tube (see illustration). Electron beams emanate from "electron guns" which are finely-pointed heated cathodes placed in close proximity to a positively charged annular plate called a "grid". The electrostatic force imposed by the grid pulls rays of energized electrons from the cathodes, and those rays are fed by the current that flows into the cathodes. These particle rays are initially accelerated towards the grid, but they soon fall under the influence of the much larger electrostatic force that results from the entire phosphor-coated display surface of the CRT being charged to 20kV (or more). The rays are focused to a fine beam as they pass through the center of the grids, and then they accelerate to impact on the phosphor-coated display surface. The phosphor surface glows brightly at the impact point, and it continues to glow for several hundred microseconds after the beam is removed. The larger the current fed into the cathode, the brighter the phosphor will glow.

Between the grid and the display surface, the beam passes through the neck of the CRT where two coils of wire produce orthogonal electromagnetic fields. Because cathode rays are composed of charged particles (electrons), they can be deflected by these magnetic fields. Current waveforms are passed through the coils to produce magnetic fields that interact with the cathode rays and cause them to transverse the display surface in a "raster" pattern, horizontally from left to right and vertically from top to bottom. As the cathode ray moves over the surface of the display, the current sent to the electron guns can be increased or decreased to change the brightness of the display at the cathode ray impact point.




Information is only displayed when the beam is moving in the "forward" direction (left to right and top to bottom), and not during the time the beam is reset back to the left or top edge of the display. Much of the potential display time is therefore lost in "blanking" periods when the beam is reset and stabilized to begin a new horizontal or vertical display pass. The size of the beams, the frequency at which the beam can be traced across the display, and the frequency at which the electron beam can be modulated determine the display resolution. Modern VGA displays can accommodate different resolutions, and a VGA controller circuit dictates the resolution by producing timing signals to control the raster patterns. The controller must produce synchronizing pulses at 3.3V (or 5V) to set the frequency at which current flows through the deflection coils, and it must ensure that video data is applied to the electron guns at the correct time. Raster video displays define a number of "rows" that corresponds to the number of horizontal passes the cathode makes over the display area, and a number of "columns" that corresponds to an area on each row that is assigned to one "picture element" or pixel.



Typical displays use from 240 to 1200 rows and from 320 to 1600 columns. The overall size of a display and the number of rows and columns determines the size of each pixel.

Video data typically comes from a video refresh memory, with one or more bytes assigned to each pixel location (the Nexys 3 uses three bits per pixel). The controller must index into video memory as the beams move across the display, and retrieve and apply video data to the display at precisely the time the electron beam is moving across a given pixel.

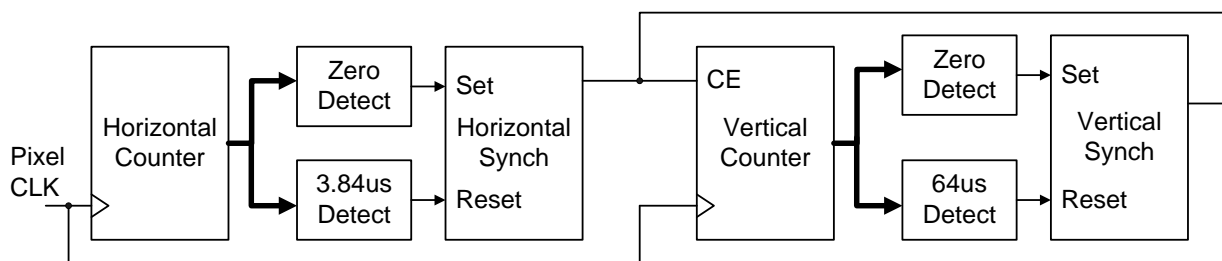
A VGA controller circuit must generate the HS and VS timings signals and coordinate the delivery of video data based on the pixel clock. The pixel clock defines the time available to display one pixel of information. The VS signal defines the "refresh" frequency of the display, or the frequency at which all information on the display is redrawn. The minimum refresh frequency is a function of the display's phosphor and electron beam intensity, with practical refresh frequencies falling in the 50Hz to 120Hz range. The number of lines to be displayed at a given refresh frequency defines the horizontal "retrace" frequency. For a 640-pixel by 480-row display using a 25MHz pixel clock and 60 +/-1Hz refresh, the signal timings shown in the table at right



Symbol	Parameter	Vertical Sync			Horiz. Sync	
		Time	Clocks	Lines	Time	Clks
T_S	Sync pulse	16.7ms	416,800	521	32 us	800
T_{disp}	Display time	15.36ms	384,000	480	25.6 us	640
T_{pw}	Pulse width	64 us	1,600	2	3.84 us	96
T_{fp}	Front porch	320 us	8,000	10	640 ns	16
T_{bp}	Back porch	928 us	23,200	29	1.92 us	48

can be derived. Timings for sync pulse width and front and back porch intervals (porch intervals are the pre- and post-sync pulse times during which information cannot be displayed) are based on observations taken from actual VGA displays.

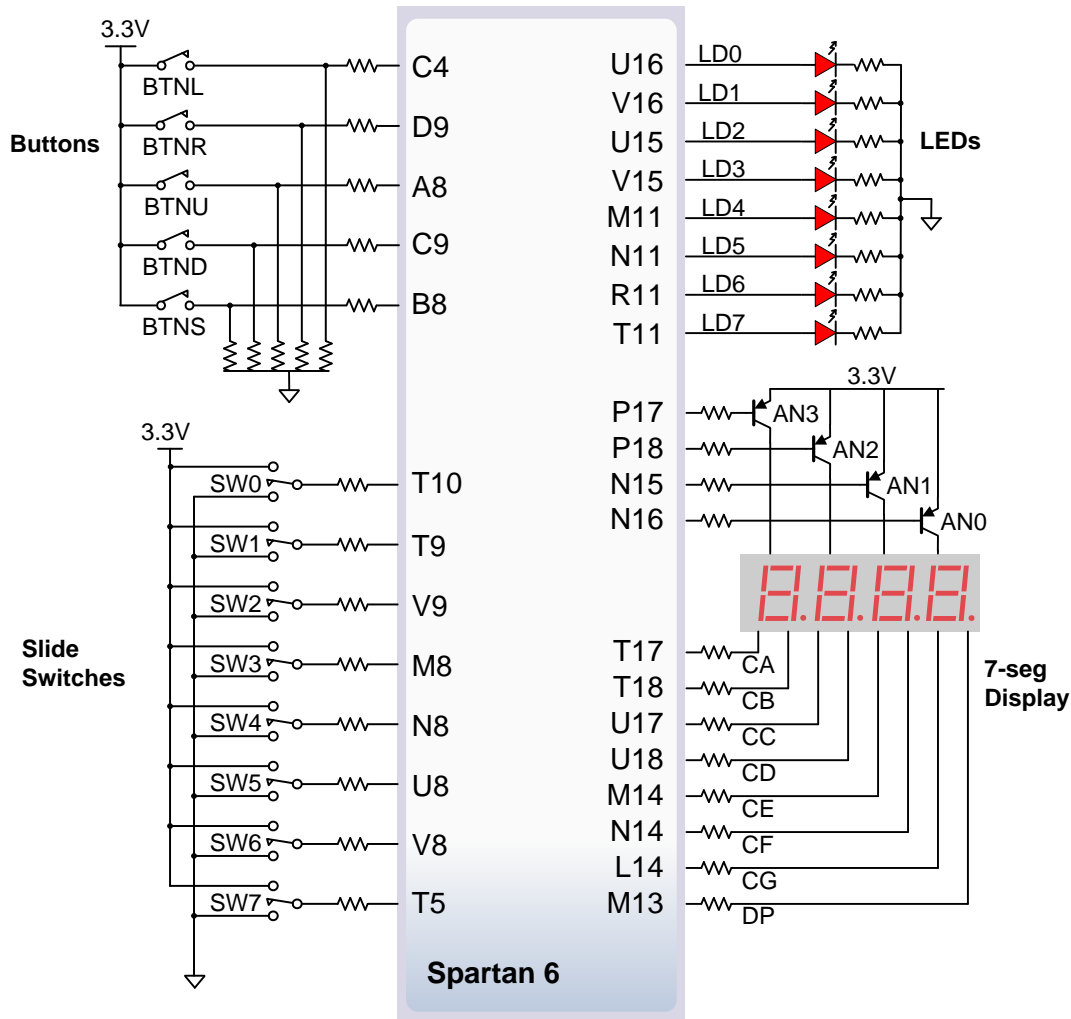
A VGA controller circuit decodes the output of a horizontal-sync counter driven by the pixel clock to generate HS signal timings. This counter can be used to locate any pixel location on a given row. Likewise, the output of a vertical-sync counter that increments with each HS pulse can be used to generate VS signal timings, and this counter can be used to locate any given row. These two continually running counters can be used to form an address into video RAM. No time relationship between the onset of the HS pulse and the onset of the VS pulse is specified, so the designer can arrange the counters to easily form video RAM addresses, or to minimize decoding logic for sync pulse generation.



9 Basic I/O

The Nexys 3 board includes eight slide switches, five push buttons, eight individual LEDs, and a four digit seven-segment display. The pushbuttons and slide switches are connected to the FPGA via series resistors to prevent damage from inadvertent short circuits (a short circuit could occur if an FPGA pin assigned to a pushbutton or slide switch was inadvertently defined as an output). The pushbuttons are "momentary" switches that normally generate a low output when they are at rest, and a high output only when they are pressed. Slide switches generate constant high or low inputs depending on their position.


The eight individual high-efficiency LEDs are anode-connected to the FPGA via 390-ohm resistors, so they will turn on when a logic high voltage is applied to their respective I/O pin. Additional LEDs that are not user-accessible indicate power-on, FPGA programming status, and USB and Ethernet port status.



9.1 Seven-Segment Display


The Nexys 3 board contains a four-digit common anode seven-segment LED display. Each of the four digits is composed of seven segments arranged in a "figure 8" pattern, with an LED embedded in each segment. Segment LEDs can be individually illuminated, so any one of 128 patterns can be displayed on a digit by illuminating certain LED segments and leaving the others dark. Of these 128 possible patterns, the ten corresponding to the decimal digits are the most useful.

The anodes of the seven LEDs forming each digit are tied together into one "common anode" circuit node, but the LED cathodes remain separate. The common anode signals are available as four "digit enable" input signals to the 4-digit display. The cathodes of similar segments on all four displays are connected into seven circuit nodes labeled CA through CG (so, for example, the four "D" cathodes from the four digits are grouped together into a single circuit node called "CD"). These seven cathode signals are available as inputs to the 4-digit display. This signal connection scheme creates a multiplexed display, where the cathode signals are common to all digits but they can only illuminate the segments of the digit whose corresponding anode signal is asserted.



A scanning display controller circuit can be used to show a four-digit number on this display. This circuit drives the anode signals and corresponding cathode patterns of each digit in a repeating, continuous succession, at an update rate that is faster than the human eye can detect. Each digit is illuminated just one-quarter of the time, but because the eye cannot perceive the darkening of a digit before it is illuminated again, the digit appears continuously illuminated. If the update or "refresh" rate is slowed to around 45 hertz, most people will begin to see the display flicker.

In order for each of the four digits to appear bright and continuously illuminated, all four digits should be driven once every 1 to 16ms, for a refresh frequency of 1KHz to 60Hz. For example, in a 60Hz refresh scheme, the entire display would be refreshed once every 16ms, and each digit would be illuminated for ¼ of the refresh cycle, or 4ms. The controller must drive the cathodes with the correct pattern when the corresponding anode signal is driven. To illustrate the process, if AN0 is asserted while CB and CC are asserted, then a "1" will be displayed in digit position 1. Then, if AN1 is asserted while CA, CB and CC are asserted, then a "7" will be displayed in digit position 2. If AN0 and CB, CC are driven for 4ms, and then A1 and CA, CB, CC are driven for 4ms in an endless succession, the display will show "17" in the first two digits. An example timing diagram for a four-digit controller is provided.



10 Expansion Connectors

The Nexys 3 board has a 68-pin VHDC connector for high-speed/parallel I/O, and an 8-pin Pmod connector for lower speed and lower pin-count I/O.



10.1 VHDC Connector

The VHDC connector includes 40 data signals (routed as 20 impedance-controlled matched pairs), 20 grounds (one per pair), and eight power signals. This connector, commonly used for SCSI-3 applications, can accommodate data rates of several hundred megahertz on every pin. Both board-to-board and board-to-cable mating connectors are available. Data sheets for the VHDC connector and for mating board and cable connectors can be found on the Digilent website, as well as on other vendor and distributor websites. Mating connectors and cables of various lengths are also available from Digilent and from distributors.

All FPGA pins routed to the VHDC connector are located in FPGA I/O bank0. The FPGA's bank0 I/O power supply pins and the VHDC connector's four Vcc pins are connected together by a small, segregated power supply plane in the PCB. This sub-plane can be connected to 2.5V or 3.3V, depending on the position of jumper JP8. This arrangement allows peripheral boards and the FPGA to share the same Vcc and signaling voltage across the connector, whether it be 3.3V or 2.5V.

The unregulated board voltage VU5V0 (nominally 5V) is also routed to four other VHDC pins, supplying up to 1A of additional current to peripheral boards. A second jumper (JP4) allows the unregulated board voltage to be disconnected from the VHDC connector if desired.

All I/O's to the VHDC connector are routed as matched pairs to support LVDS signaling, commonly powered at 2.5V. The connector uses a symmetrical pinout (as reflected around the connector's vertical axis) so that peripheral boards as well as other system boards can be connected. Connector pins 15 and 49 are routed to FPGA clock input pins.

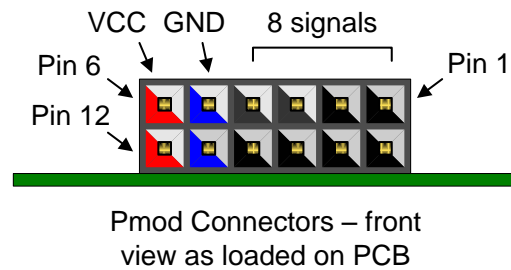
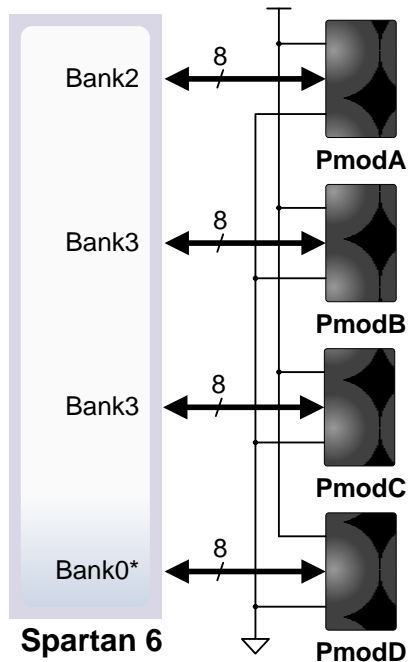


VHDC Connector Pinout			
IO1-P: B2	IO1-N: A2	IO11-P: C10	IO11-N: A10
IO2-P: D6	IO2-N: C6	IO12-P: G9	IO12-N: F9
IO3-P: B3	IO3-N: A3	IO13-P: B11	IO13-N: A11
IO4-P: B4	IO4-N: A4	IO14-P: B12	IO14-N: A12
IO5-P: C5	IO5-N: A5	IO15-P: C13	IO15-N: A13
IO6-P: B6	IO6-N: A6	IO16-P: B14	IO16-N: A14
IO7-P: C7	IO7-N: A7	IO17-P: F13	IO17-N: E13
IO8-P: D8	IO8-N: C8	IO18-P: C15	IO18-N: A15
IO9-P: B9	IO9-N: A9	IO19-P: D14	IO19-N: C14
IO10-P: D11	IO10-N: C11	IO20-P: B16	IO20-N: A16

10.2 Pmod Connectors

Pmod connectors are 2x6 right-angle, 100-mil female connectors that mate with standard 2x6 pin headers available from a variety of catalog distributors. Each 12-pin Pmod connector provides two 3.3V VCC signals (pins 6 and 12), two Ground signals (pins 5 and 11), and eight logic signals. VCC and Ground pins can deliver up to 1A of

current. Pmod data signals are not matched pairs, and they are routed using best-available tracks without impedance control or delay matching.



Pmod Pinouts			
JA1: T12	JB1: K2	JC1: H3	JD1: G11
JA2: V12	JB2: K1	JC2: L7	JD2: F10
JA3: N10	JB3: L4	JC3: K6	JD3: F11
JA4: P11	JB4: L3	JC4: G3	JD4: E11
JA7: M10	JB7: J3	JC7: G1	JD7: D12
JA8: N9	JB8: J1	JC8: J7	JD8: C12
JA9: U11	JB9: K3	JC9: J6	JD9: F12
JA10: V11	JB10: K5	JC10: F2	JD10: E12

Digilent produces a large collection of Pmod accessory boards that can attach to the Pmod and VHDC expansion connectors to add ready-made functions like A/D's, D/A's, motor drivers, sensors, cameras and other functions. See www.digilentinc.com for more information.

11 Built-In Self Test

A demonstration configuration is loaded into the BPI PCM device on the Nexys 3 board during manufacturing. This demo bit file is available for download from the Digilent website. If the demo configuration is present in the BPI PCM device and the Nexys 3 board is powered on in BPI mode, a simple demo project will allow basic hardware verification. The demo drives a counter on the 7-segment display, drives the user LEDs on and off when the user switches are toggled, turns off digits on the 7-segment display when user buttons are pressed, and drives an image out the VGA port. A USB mouse can be connected to J4 for a simple visual demonstration.

If the demo configuration is not present in the BPI device, it can be downloaded from the Digilent website and programmed directly into the FPGA or reloaded into the PCM device using the Adept programming software. The same basic functionality is also available using the automated Test tab in the Adept application. Using the Test tab automatically loads a demo configuration into the FPGA, so no separate bit files or project need to be used.

All Nexys 3 boards are 100% tested during the manufacturing process. If any device on the Nexys 3 board fails test or is not responding properly, it is likely that damage occurred during transport or during use. Typical damage includes stressed solder joints and contaminants in switches and buttons resulting in intermittent failures. Stressed solder joints can be repaired by reheating and reflowing solder and contaminants can be cleaned with off-the-shelf electronics cleaning products. If a board fails test within the warranty period, it will be replaced at no cost. If a

board fails test outside of the warranty period and cannot be easily repaired, Digilent can repair the board or offer a discounted replacement. Contact Digilent for more details.