Selecting the Correct CMOS PLD—An Overview of Advanced Micro Devices' CMOS PLDs

Advanced Micro Devices

Application Note

INTRODUCTION

The purpose of this application note is to provide a survey of AMD's CMOS PLDs (Programmable Logic Devices). This includes both PAL (Programmable Array Logic) devices and the more general realm of PLDs to which PAL devices belong. With the proliferation of parts, the selection of the best PLD for your application may seem difficult. If you are a new PLD user, this overview will guide you through the wide variety of different device architectures, speed, and power grades. This tutorial should increase your understanding of the basic characteristic features that make a device appropriate for a given application.

Figure 1 shows a "CMOS PAL Selection Route Map." This can be used as a convenient model of the discussion throughout this paper. It can also be used as a reference guide when you are selecting a PLD for a particular application.

The Benefits of AMD's CMOS PLDs

Before addressing individual products, it is important to understand why CMOS technology is used in PLDs. There are two universal benefits of AMD's CMOS PLDs: electrical erasability and low power.

Electrical Erasability

Because PLDs are programmable, electrical erasability is probably the most important advantage that CMOS technology can bring. AMD's electrically-erasable CMOS has benefits that make it superior to both UV-erasable CMOS and bipolar technologies. The most important advantage is the ability to erase the device electrically in a matter of seconds as opposed to hours for UV-based technologies, and not at all in the case of bipolar. The chief benefit to the user is a very high quality device. This is realized through the ability to erase and reprogram the device many times at various test points in the factory. In fact, the quality is so good that programming and post-programming functional rejects are virtually non-existent.

A second major benefit of electrical erasability is the ease of reprogramming the device by the user. This saves time when prototyping, and allows for recovery of large volumes of devices that may need to be reprogrammed for a variety of reasons, such as mid-production bug fixes. Erasure takes place automatically by the device programmer, and is completely transparent to the user.

Low Power

The most well-known attribute of CMOS is low power consumption. All PALCE (Programmable Array Logic CMOS Electrically-Erasable) devices are offered in half-power versions; they require at most half the supply current of their first-generation bipolar counterparts. Some devices are also offered in quarter-power versions. This is achieved by taking the latest process technology and designing to favor even lower power over the fastest speed possible.

CMOS uses less current than bipolar because most of the current flow only takes place while the transistors are actually switching. With bipolar, current flows through the transistors all the time. AMD's half- and quarter-power CMOS devices take advantage of this as much as possible. However, in order to achieve high speed it is necessary to operate some transistors on the chip in the linear region. Because this circuitry is essentially always switching, the power consumption does not go to zero as it would in a conventional CMOS device.

Lower power requirements are ideal for applications which have tight power budgets, such as mobile telecommunications. Smaller power supplies also reduce cost and lowers heat dissipation. This results in smaller cooling fans, or perhaps no fan at all. It also allows the system designer to pack everything even tighter since less empty space is needed for air circulation. This can make the circuit board fit in a smaller package, again reducing cost.

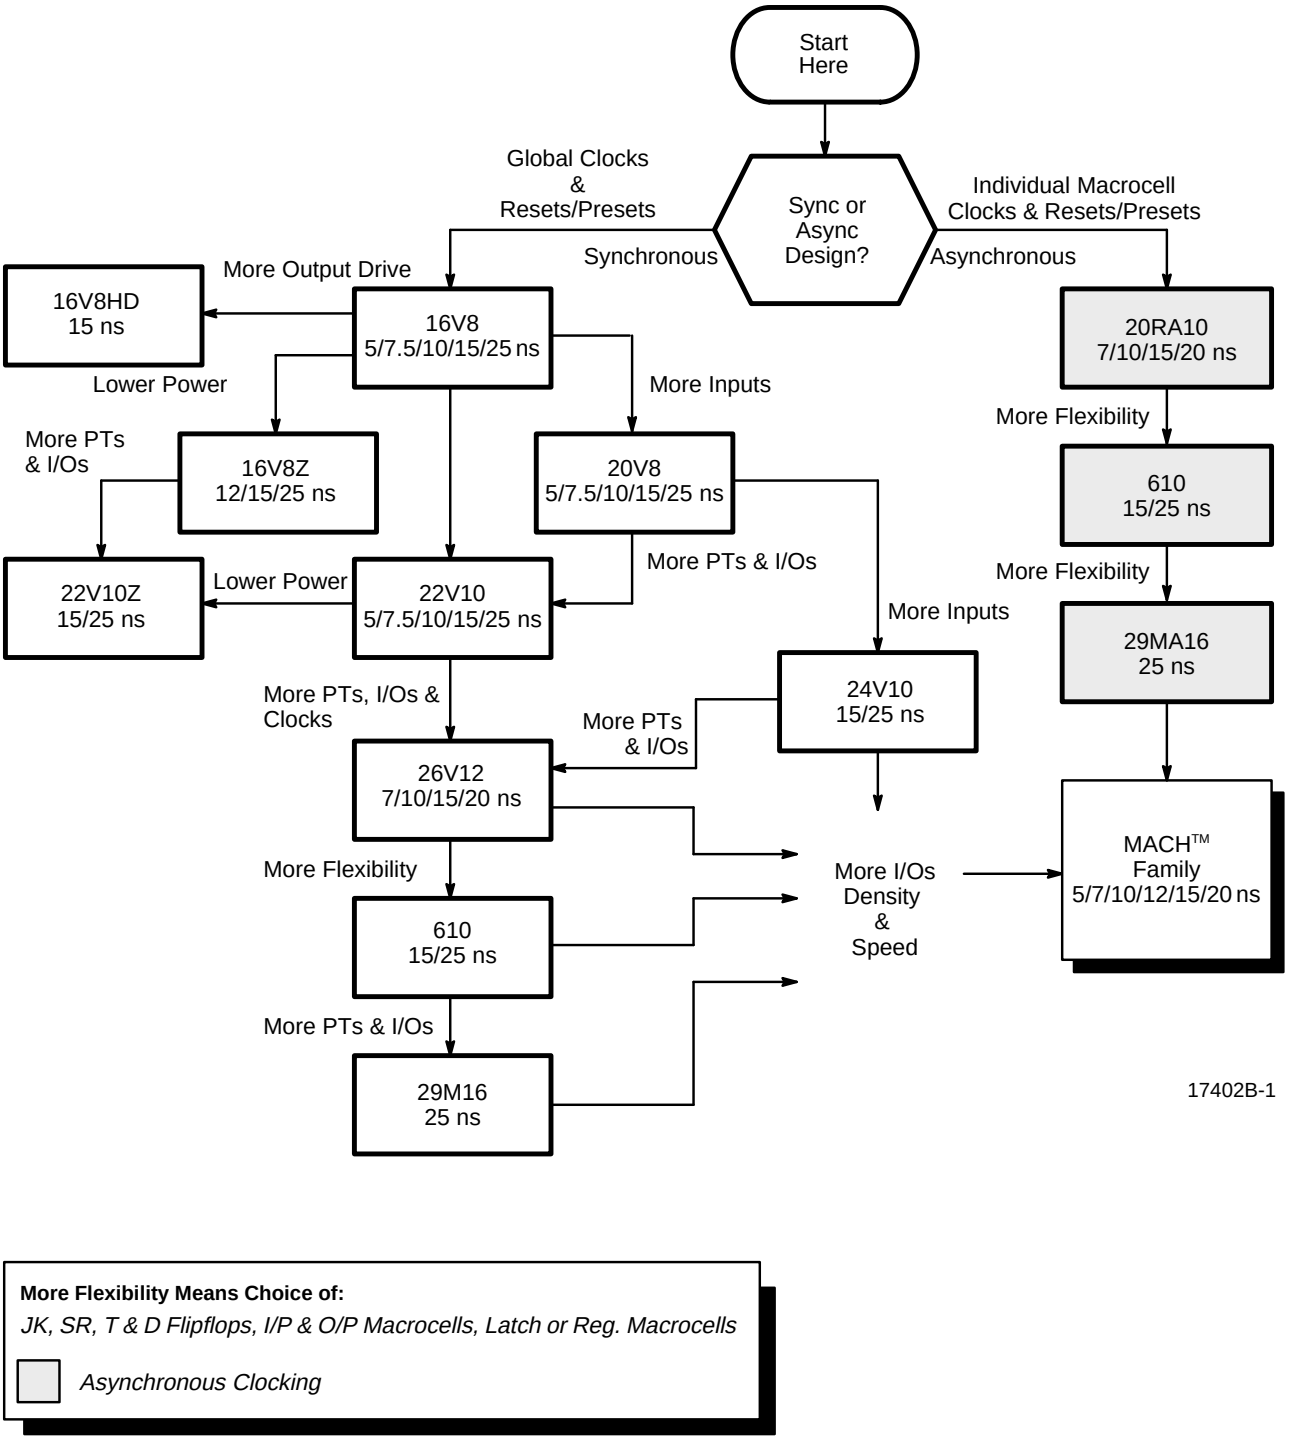

Figure 1. CMOS PAL Selection Route Map

Zero-Power

Some devices are also offered in zero-standby power versions. Instead of *always* operating certain transistors in the linear region to achieve high speed, this circuitry can shut down and the device goes into standby mode. Standby mode is defined as anytime the inputs do not switch for an extended period of time (typically 50 ns). When this happens the current consumption will almost go to zero (I_{CC} < 15 or 30 μA). The outputs will maintain the current state held while the device is in standby mode.

When any input switches, the internal circuitry is fully enabled and the power consumption returns to normal. This feature results in considerable power savings for operation at low to medium frequencies.

Zero-standby devices are desirable for a number of reasons. In portable and field-installed equipment that rely on batteries, battery life is extended. In solar powered systems, fewer solar cells are required.

The Selection Process

When selecting a CMOS PLD, start by determining both the functional and size requirements for your application.

Functional Requirements

The functional requirements of a given application are what determine which device *architecture* should be used. The functional criteria include such issues as the clocking scheme, macrocell flexibility and output drive.

The clocking scheme can be synchronous, where all registers within a device use the same clock signal, or asynchronous, where each register can be clocked individually using any logic signal or combination of logic signals available to the device. These two alternatives are illustrated in Figure 2.

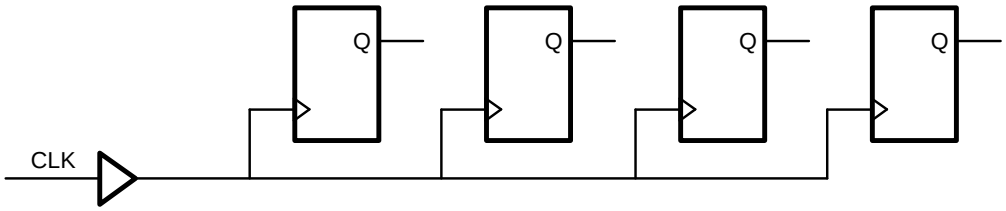

a. Registers with Synchronous Clocking

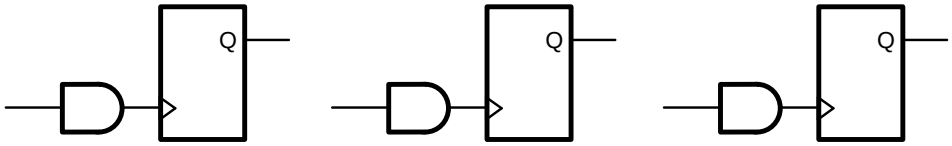

b. Registers with Asynchronous Clocking

17402B-2

Figure 2. Basic Clocking Schemes

Macrocell flexibility refers to the ability to configure the output in various ways. A macrocell (Figure 3) takes the basic sum-of-products logic and adds functionality through features like storage elements, optional path controls, polarity, and feedback. This concept will be further illustrated as each device macrocell is explained throughout the selection process.

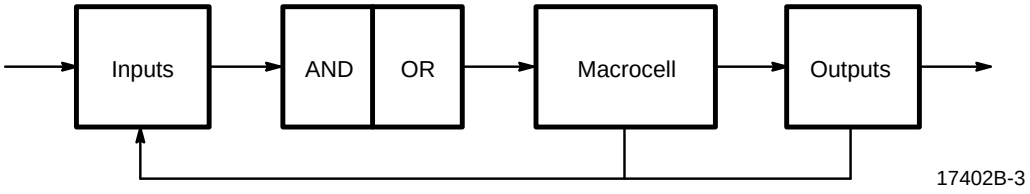

17402B-3

Figure 3. PLD Block Diagram with the Macrocell Included

Size Requirements

The combination of number of inputs and outputs required determines the device *size* that should be used. The number of product terms required to implement a particular design also factor into this decision.

The best approach to selecting the appropriate device is to begin with the simplest and smallest devices and upgrade as necessary to accommodate your application.

Combinatorial and Synchronous Applications

Starting from the top of the flow chart in Figure 1 and taking the path for synchronous designs leads one to those devices best suited for simple state machines, encoders, decoders, muxes, and similar logic applications. For these applications, D-type registered or combinatorial (non-registered) logic is needed. The first choice is the **PALCE16V8, PALCE20V8, or PALCE24V10,** depending on the number of inputs or outputs needed. The macrocells (Figure 4) can be configured to use combinatorial or D-type registered outputs in any combination. The D-type register operates very simply; data presented at the D input of the register will be clocked into the register on the rising edge of the clock signal. The output can be configured as active low- or active high-output depending on the requirements of the downstream devices and the efficiency of the logic.

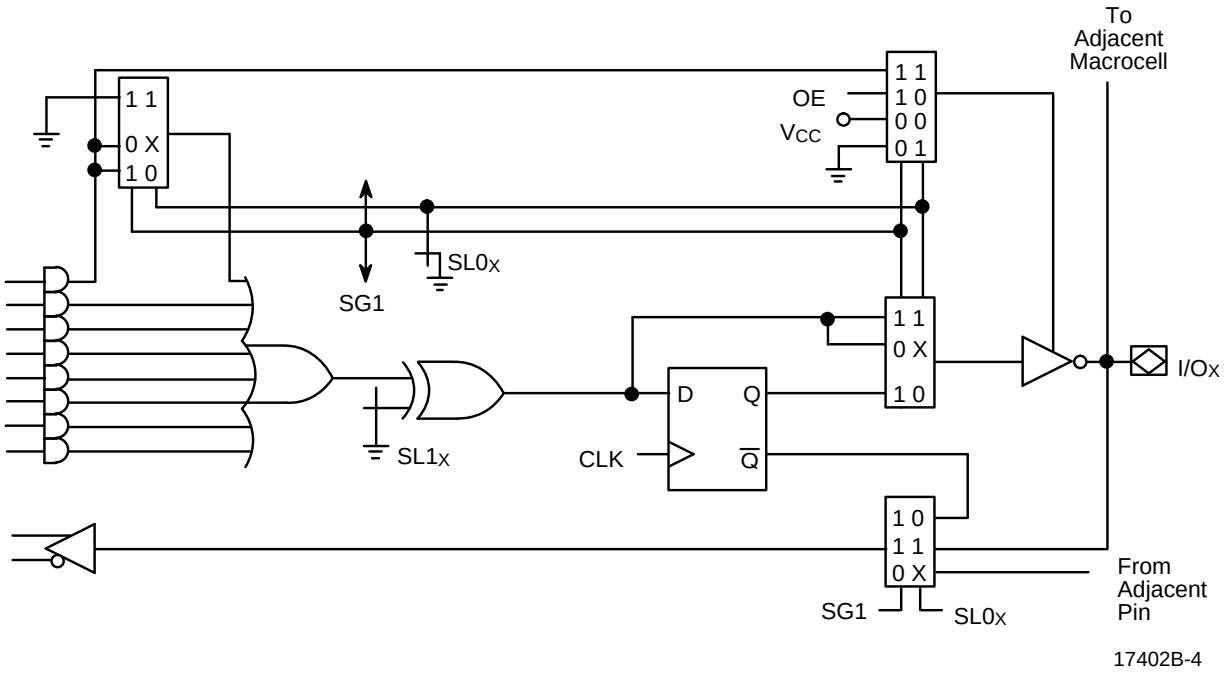

Figure 4. PALCE16V8 Macrocell

More Product Terms and I/O's

If the application requires the features of the PALCE16V8 macrocell, but requires more programmable gate functions, move further down the flowchart. The next device, the **PALCE22V10** (Figure 5) has varied product term distribution. The number of product terms varies among outputs, with up to 16 product terms on some outputs. In addition, it has global synchronous preset and asynchronous reset product terms. These are connected to all macrocells configured as registers, facilitating easy power-up and system reset. This versatility contributes significantly to the 22V10 being the world's most popular PAL architecture. When combined with the PALCE16V8 family, these two device families will likely handle about 80% of PLD applications.

The **PALCE26V12**, a 28-pin version of the 22V10 increases versatility by adding more inputs and outputs, and by adding another global clock. Any macrocell can use one of the two clocks. This allows the logic to be partitioned giving greater design flexibility. In addition, registered outputs can be configured as bidirectional pins on the 26V12.

Since historically most applications in bipolar technology had been done in PAL16R8, PAL20R8, and PAL22V10 families, these types of applications can easily have their power lowered with half-, quarter-, and zero-power versions of the PALCE16V8, PALCE20V8, or PALCE22V10.

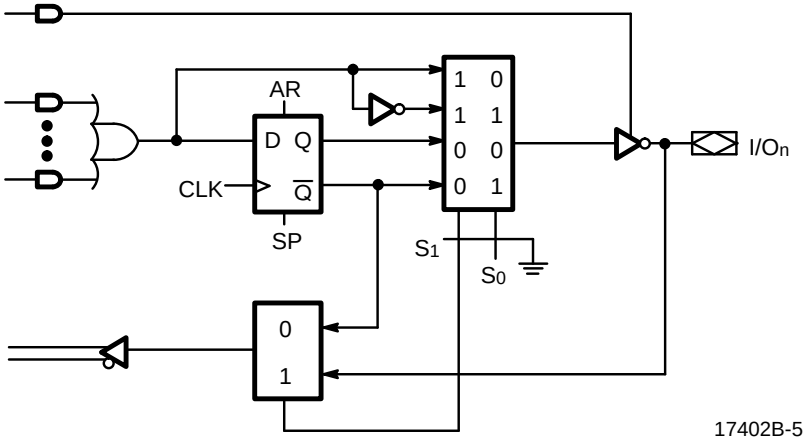

Figure 5. PALCE22V10 Macrocell Diagram

Table 1. Summary of the PALCE16V8 and PALCE22V10 Family Architectures

Device	Macrocell Type	Dedicated Inputs	I/O's	Clocks	Product Terms
PALCE16V8	16V8	8–10	8	1	7–8
PALCE20V8	16V8	10–12	8	1	7–8
PALCE24V10	16V8	14	10	1	7–8
PALCE22V10	22V10	12	10	1	8–16
PALCE26V12	22V10	14	12	2	8–16

Table 2. Summary of the Speed and Power Requirements for the PALCE16V8, PALCE20V8, and PALCE22V10 Families of Devices

Device and Speed Grade	I_{CC}
PALCE16V8H-5	125 mA static
PALCE16V8H-7/-10	115 mA at 25 MHz
PALCE16V8H-15/-25	90 mA at 15 MHz
PALCE16V8Q-15/-25	55 mA at 15 MHz
PALCE16V8Z-25	15 μA in standby mode
PALCE20V8H-5	125 mA static
PALCE20V8H-7/-10	115 mA at 25 MHz
PALCE20V8H-15/-25	90 mA at 15 MHz
PALCE20V8Q-15/-25	55 mA at 15 MHz
PALCE22V10H-5	140 mA at 25 MHz
PALCE22V10H-7	140 mA at 25 MHz
PALCE22V10H-10	120 mA at 25 MHz
PALCE22V10H-15/25	90 mA static
PALCE22V10Q-25	55 mA static
PALCE22V10Z-15/25	30 μA in standby mode

Bus Applications

For applications that require bus interaction, the **PALCE16V8HD** features 64 mA low-output and 16 mA high-output drive capability. This compares to 24 mA low-output and 3.2 mA high-output drive for other PALCE16V8 devices. The PALCE16V8HD also has some unique macrocell features. Because this device is designed to drive a bus, it can be configured with open-drain outputs. Open-drain (open-collector) configuration is sometimes used in bus applications because it provides controlled V_{OH}, termination, and wire-NOR capability. Because the PALCE16V8HD is designed to take inputs directly from a noisy bus, all inputs have 200 mV input threshold hysteresis to improve noise immunity. The inputs can be configured as direct or latched, making additional buffering devices unnecessary. Additionally, the macrocell can be configured as either a D- or T-type register. The T-type register is more efficient for counter applications because fewer product terms are consumed as hold states.

Complex Functions

For those applications that require D, T, J-K, or S-R register capability, the **PALCE610** (Figure 6) has this flexibility. This device has 16 macrocell outputs and four dedicated inputs. The J-K, S-R, and T registers allow easy implementation of counters and larger state machines.

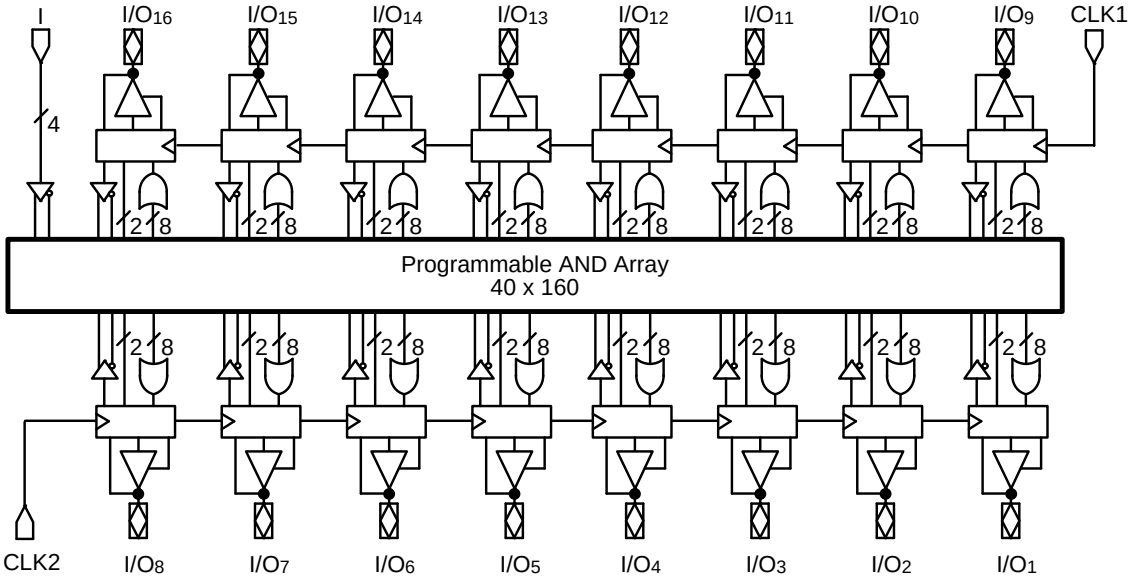

17402B-6

Figure 6. PALCE610 Block Diagram

More Flexibility

Moving down to the bottom of the flow chart reveals a more flexible device than the PALCE610. Applications that make use of embedded, or buried, registers can take advantage of the **PALCE29M16** (Figure 7). Buried register operation is very useful when a state machine uses state bits that do not need to be brought outside the chip. This allows the pin associated with the macrocell to be available as an input. Eight of the 16 macrocells have *dual feedback* capability. This means that these macrocells have two independent feedback paths: one from the register and one from the I/O pin. The other eight macrocells have *single* feedback, where both paths are available but, only one or the other can be used. Since almost every pin is an I/O pin this device has 29 available inputs to the programmable AND array.

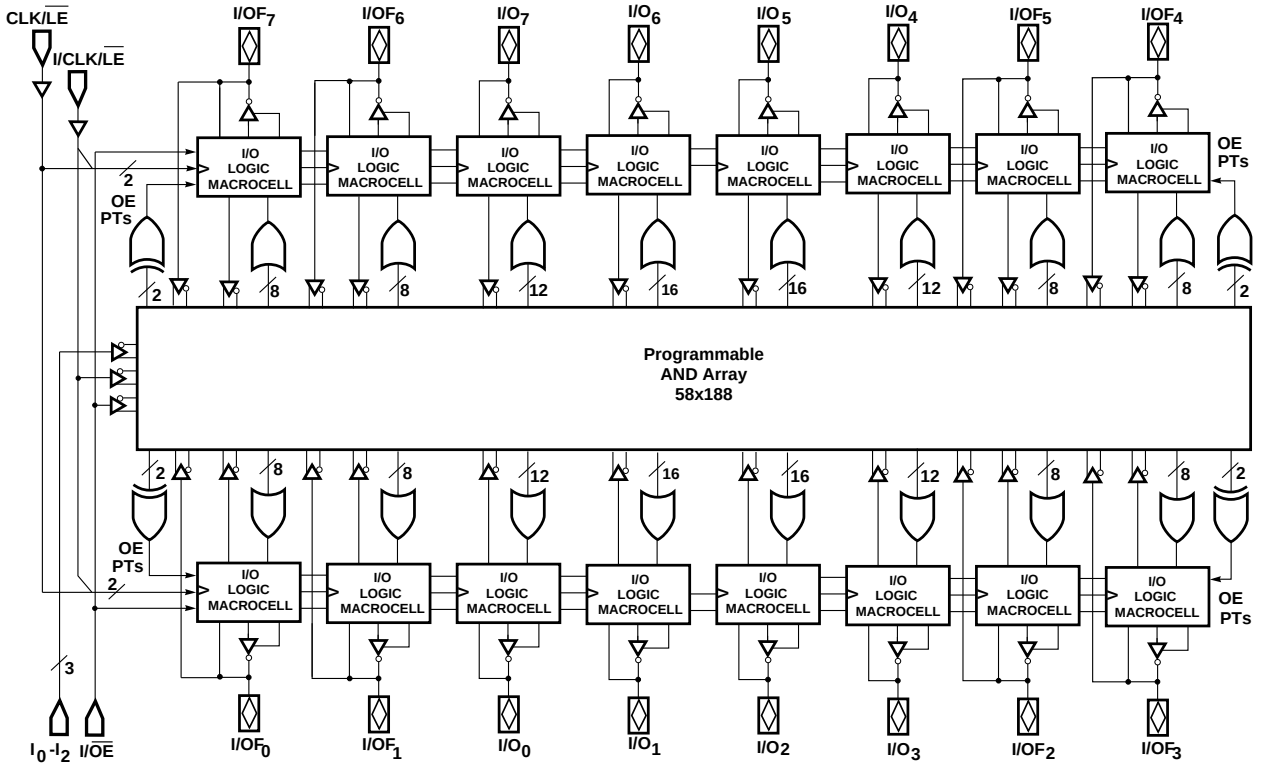

17402B-7

Figure 7. PALCE29M16 Block Diagram

The macrocells (Figure 8) can be configured as latches or registers. Latch operation allows the flip-flop in the macrocell to become transparent while the latch is enabled. When the latch is disabled, the flip-flop will hold the current state. This kind of operation is useful in sample and hold applications. Also, the register or latch may be used with the macrocell input pin for synchronizing signals. The active level of the latch enable, as well as the clock edge (rising or falling) are both programmable.

Like the PALCE22V10, there is varied product term distribution among the macrocells. Also, preload capability is available using a global product term to define the preload condition. Preload capability allows arbitrary states to be loaded directly into the register, making it unnecessary to cycle through long, complex vector sequences to get the device in a particular state. This is normally only performed by the device programmer when it performs functional tests. The preload product term allows preload to be engaged by hand, without the need for supervoltages (voltages above V_{CC} needed to engage preload on most PLDs).

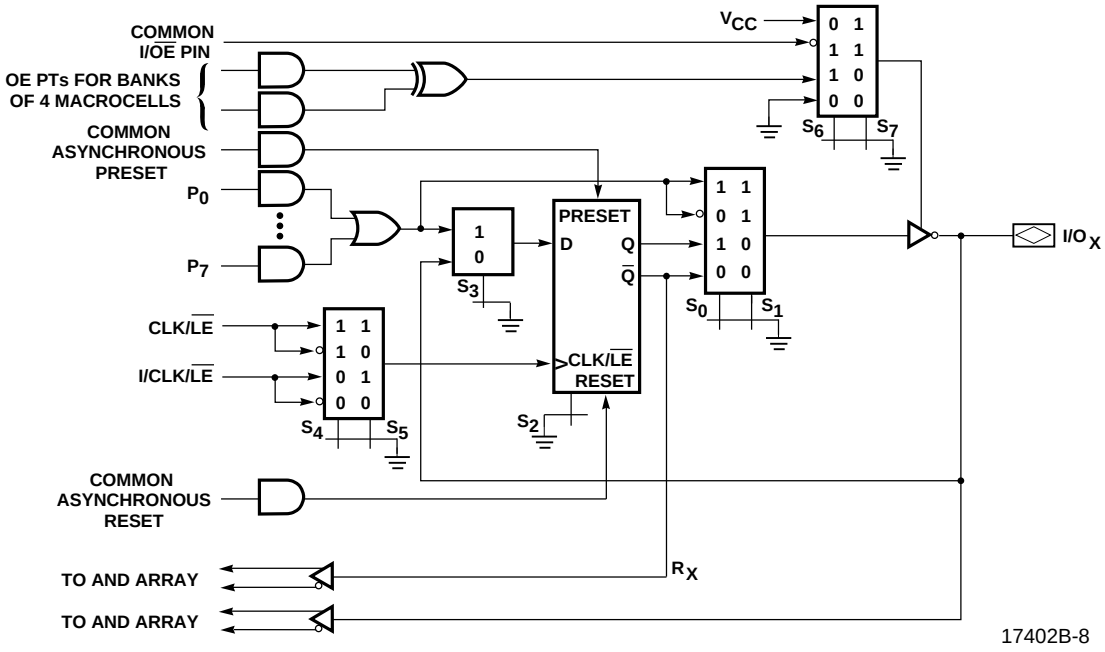

Figure 8. PALCE29M16 Macrocell Diagram

Asynchronous Applications

Starting at the top of the flow chart again, but this time taking the path for asynchronous applications, you will find **PALCE20RA10.** The PALCE20RA10 is the simplest of the asynchronous devices. Each macrocell (Figure 9) is clocked individually using one product term per macrocell. Also, reset, preset, and output enable are controlled individually with one product term each. There is also a global output enable pin which is combined with the product term enable to determine if the output is enabled or disabled.

A dedicated global preload pin allows all registered macrocells to be preloaded simultaneously using normal TTL levels.

The PALCE610 has the distinction of bridging the gap between synchronous and asynchronous register clocking. The PALCE610 macrocells can be clocked via individual product terms for each macrocell, or the macrocells can be clocked in banks of eight via two dedicated clocks. This is done by using a clock/output enable mux (Figure 10). If the macrocell is configured as combinatorial or as a synchronous register, output enable/disable is controlled by a product term. If asynchronous register mode is desired, the same product term is used as a clock and the macrocell is always enabled.

As mentioned above, the PALCE610 can act as a synchronous or asynchronous PAL device. As shown, a special function product term can be steered to control either the output enable or the macrocell clock. In the latter mode, each register can be individually clocked.

If your application requires the basic features of the PALCE29M16, except with individual macrocell control, the **PALCE29MA16** (Figure 11) should be considered. Four product terms in each macrocell are dedicated to control clocking, output enable, asynchronous reset, and asynchronous preset. These functions are controlled either globally or in blocks on the PALCE29M16. A common clock pin and output enable pin are still maintained, but the user has a choice of using either the common control pin or individual macrocell control via the control product term.

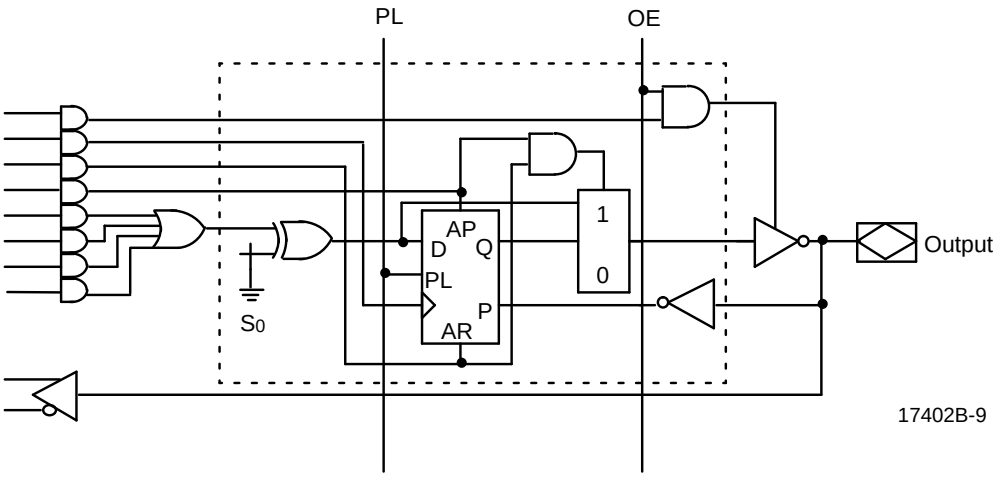

Figure 9. PALCE20RA10 Macrocell Diagram

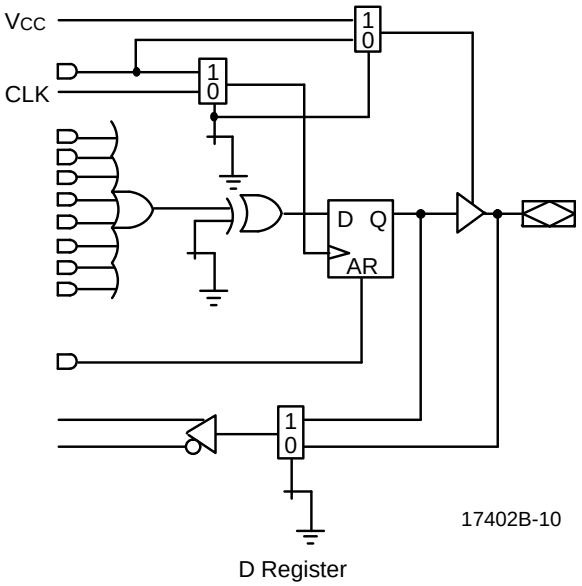

Figure 10. PALCE610 Macrocell (Configured as a D-Register) with the Output Enable/Clock Mux

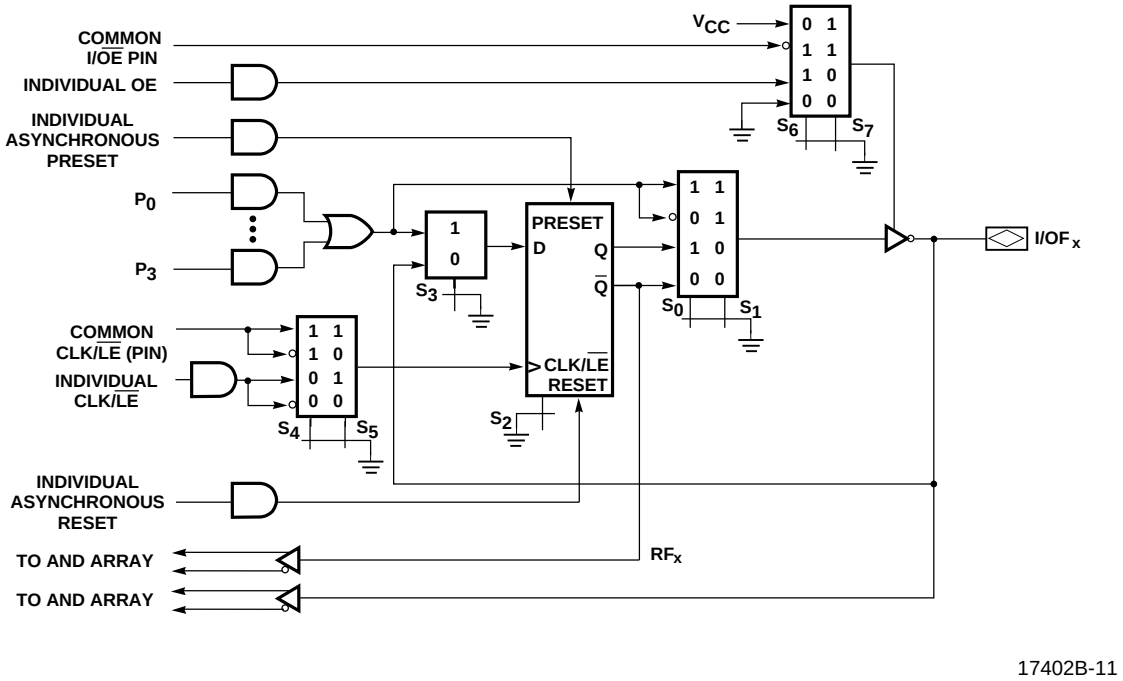

Figure 11. PALCE29MA16 Macrocell

MACH Devices

At the bottom of the flow chart, both the synchronous and asynchronous design paths converge into the MACH Family. MACH devices extend AMD's PLD offerings into the realm of what is referred to as mid-density. Mid-density devices allow multiple smaller PLD designs to be consolidated into one device. Mid-density covers replacement of just a couple of smaller PAL devices, all the way up to designs that would traditionally be done with small gate arrays. These devices span from 900 to 3600 gates, with 32 to 64 macrocells, and are offered in 44- to 84-pin packages.

MACH devices use PAL blocks that are interconnected using a switch matrix. The members of the families are differentiated by the number of pins, macrocells, clocks, and the amount of interconnect. The MACH 1 family has output macrocells; the MACH 2 family has output and buried macrocells. All signals, whether registered or combinatorial can be buried. The basic macrocell, common to both families resembles the PALCE22V10 macrocell with the additional choice of using D- or T-type registers. The MACH211 is illustrated in Figure 12.

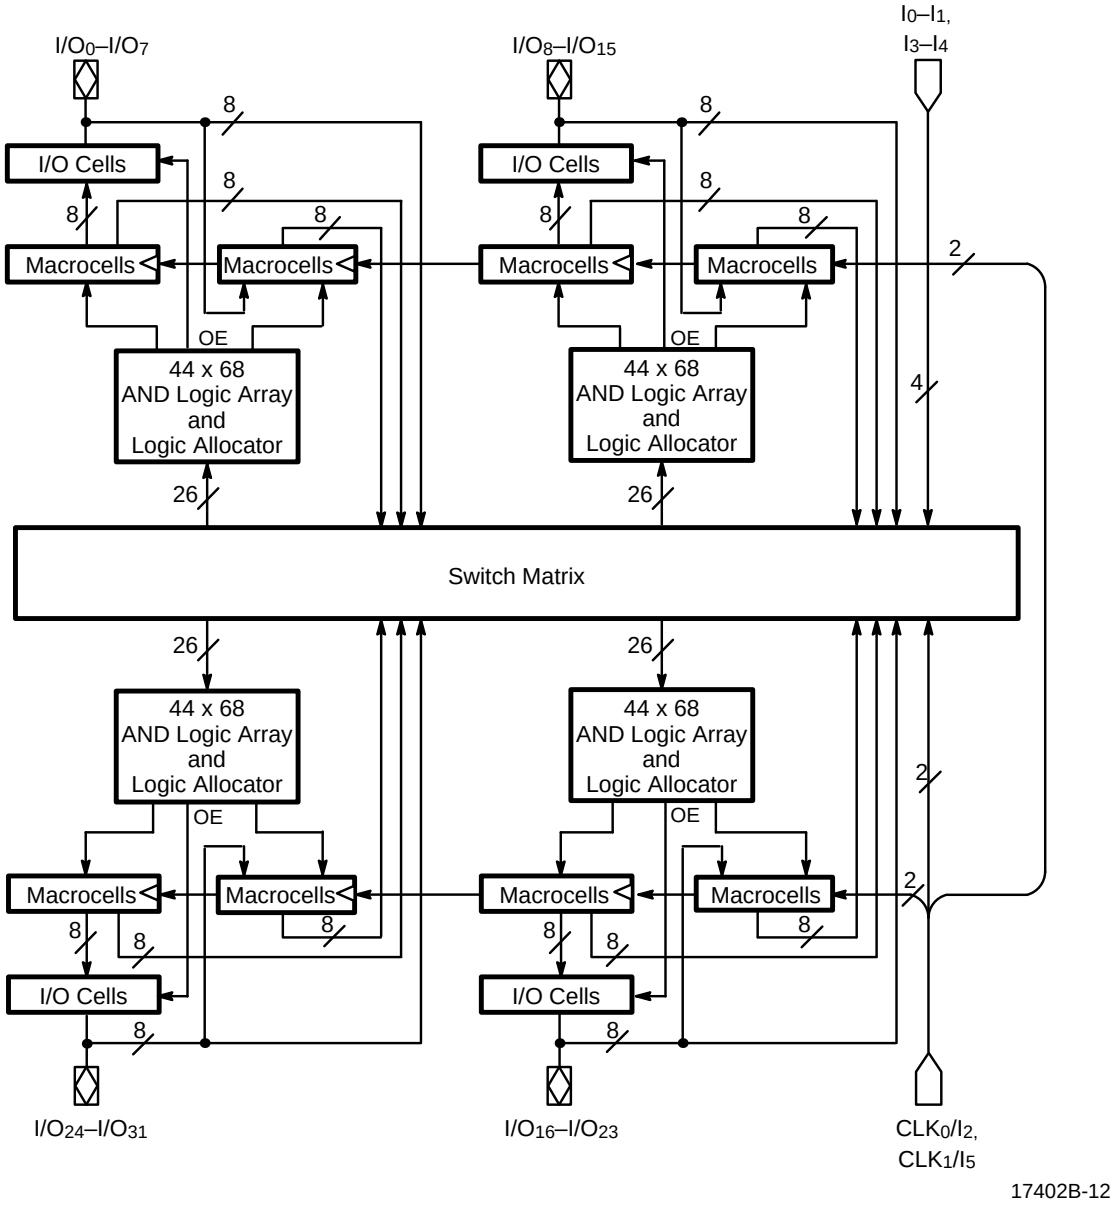

Figure 12. MACH211 Block Diagram

Both synchronous and asynchronous versions are available. The asynchronous MACH215 macrocell looks more like a PALCE20RA10 macrocell, rather than PALCE22V10 type macrocells. The synchronous devices are better suited to structured designs; the asynchronous MACH215 is better suited for random logic collection.

All MACH devices have the advantage of fast (5, 7, 10, 12, 15, and 20 ns), predictable timing, which is a unique advantage when compared to other mid-density PLDs.

SUMMARY

Selecting the appropriate PLD for a given application involves matching your requirements with various device capabilities. Following the guidelines in this article along with the "CMOS PAL Selection Route Map" makes it easier.

AMD's wide array of electrically-erasable CMOS PLDs, combined with strong third party support through our FusionPLD^{SM} relationships, means an excellent selection of devices, software, and programming hardware. This gives you a virtual toolbox of solutions for your system logic requirements, along with the strong technical support you expect.