


EXPANSION CARDS PART 2:

AGP

The first in this series of Tech Tips on expansion cards took a look at the PCI slot, and the variety of devices that may find their home in one. Graphics cards are one of the many items that may be used in a PCI slot, but the demands of fast-paced video games require more speed and greater bandwidth than the PCI Bus can provide. Thus, the AGP slot was born, providing a dedicated interface to transfer graphics data only.

The letters ‘AGP’ stand for Accelerated Graphics Port, and it is the term used to describe a dedicated, point-to-point interface that connects a video card directly to the system’s memory and processor.



AGP was first introduced by Intel in 1996, and is based off of their previous work in developing the PCI bus. Despite being based on PCI technology, the AGP and PCI slots on a motherboard are not interchangeable, so an AGP card can not be installed into a PCI slot, and vice versa.

The initial release of AGP saw a sizeable performance boost over PCI, and the few revisions to the standard helped increase this even more as years went by. Other than having a dedicated path to the system's memory and processor, several other design features help AGP outperform PCI when it comes to graphics performance. Three of the other advancements: pipelining, side band addressing and graphics address remapping table are described below.

Data transfer is improved through 'pipelining', a term used to describe the ability of an AGP graphics card to receive, and act upon, multiple instructions simultaneously. PCI data transfers require each piece of necessary information to be received separately before acting on any of it.

Something called 'side band addressing' also provides AGP with a

performance boost. Basically, additional lines of data are included with each packet to instruct the system as to where this data is to be used. PCI data transfers do not have this addressing information, and the system must look at the data itself in order to determine its destination. This is an obvious time saver, as well as a resource saver since the processor doesn't have to analyze all data just to determine the address.



AGP allows the operating system to store texture maps in the system's memory which allows for more space, and perhaps faster access, rather than being limited to the use of graphics card memory only. Graphics art address remapping table, also known as GART, is a term used to describe a process that maps physical memory as virtual memory for the storage of texture maps. Basically, GART takes the system memory it is allowed to use to store texture maps and re-addresses it so that the system thinks these maps are now

actually being stored in the frame buffer, or virtual memory. This might not sound like anything special, but this re-addressing requires that the texture map be written to memory only once and it is locked into place right where the AGP card can find it quickly.

AGP can be broken down into different groups based on revisions to the specification (AGP 1.0, AGP 2.0, and AGP 3.0), as well as by the general speeds (1x, 2x, 4x, and 8x). There is overlap between the various categories, with AGP 1.0 supporting 1x and 2x, AGP 2.0 supporting 1x, 2x, and 4x, and AGP 3.0 supporting 4x and 8x. For a complete break down of all the combinations available, please visit [this page](#).

Before taking a look at the specifications of AGP, let's have a refresher as to what was available on PCI prior to the birth of AGP.

The standard PCI bus has a width of 32-bit, operates at 33 MHz, provides a maximum bandwidth of 132 MB/s (which has to be shared by all devices connected), and operates on 3.3V (or 5V on the original standard).

The first version released was AGP 1.0 with a speed of 1x, which offered the following specifications: 32-bit bus width, operating at

66 MHz, providing a maximum bandwidth of 266 MB/s, and utilizing 3.3V. So, it can be seen that right out of the gate, AGP offered double the bandwidth of PCI.




Each speed increase over 1x provided double the bandwidth as well as double the clock speed through the use of special signaling. So, AGP 2x offers a maximum bandwidth of 533 MB/s at a speed of 133 MHz, AGP 4x offers a maximum bandwidth of 1066 MB/s at a speed of 266 MHz, and AGP 8x offers 2.1 GB/s at a speed of 533 MHz.



Given the timeline of the evolution of these cards, AGP 8x cards dominate today's marketplace. Finding some cards that are backwards compatible is possible, but the tricky part may be ensuring that the slot on the motherboard will accept them. Comparing the connector on this 128MB Apollo GeForce FX6600 GT card, to the connector on this 64MB Hercules 3D

Prophet Ultra II card, and to the connector on this 256MB Chaintech GeForce FX5200 card shows that the first one is obviously different than the second two. The Apollo card is 8x only, the Hercules card is 4x/2x compatible, and the Chaintech card is 8x/4x, which results in different notches in the connector.



AGP 1.0 only features a 3.3V connection, the release of AGP 2.0 saw the availability of both a 3.3V and 1.5V connector, and AGP 3.0 uses the same 1.5V, but only requires 0.8V for signaling. In order to protect cards of different voltages/formats, special keyed connectors were designed so that only the correct card could be installed on any motherboard. A universal connector was eventually released for AGP 1.0/2.0 which allowed cards of either voltage to be installed. For a schematic of the various connectors, please visit [this page](#). Although AGP 3.0 can share in the use of a universal connection, many motherboards now only support 4x/8x cards based on the AGP 3.0 standard.

Another specification for AGP was released between 2.0 and 3.0, and was referred to as AGP Pro. AGP Pro was intended to be the new standard to meet the demands of high powered graphics

workstations, but it never really garnered widespread acceptance. Speeds of 1x, 2x, and 4x were supported with AGP Pro, and it utilized either a 3.3V, 5V, or a universal connector, similar to AGP 2.0. But, the AGP Pro connector was not the same size as the 'standard' AGP connector (see schematic at link above), meaning there were now three more possible connections to consider. An AGP Pro connection is longer than a standard AGP connection, and depending on the connector type, it could accept AGP 1.0 and 2.0 cards.

Modern motherboards supporting AGP will specify what type of card is compatible with the board, so the guess work is eliminated when trying to match one with the other. For example, this Socket 754 Chaintech motherboard specifies that it has one AGP 4x/8x slot and this Biostar LGA 775 motherboard specifies that it has one 8x AGP slot.

Source: <http://www.geeks.com/techtips/2005/techtips-030305.htm>