# PowerPC 405GP System Design Considerations

February 21, 2001 VERSION 1.1

Abstract - As processor speeds increase more attention must be paid to system design issues. The topics of thermal management, terminations, decoupling, and debug are discussed

## 1. Thermal considerations

As frequencies increase and peripheral integration increases, chip power also increases. As power increases, thermal management is an increasingly important part of the system design. The ultimate thermal constraint is the silicon temperature, which is commonly referred to as the junction temperature. However, since most parts are encapsulated in packages the junction temperature is not observable. The next best point to observe temperature is the surface of the package, also referred to as the case. Therefore, for encapsulated parts the temperature range is specified in terms of case temperature.

In some data sheets we have included a table with Case to Ambient thermal impedance, $\theta_{CA}$. These are for reference only since the JEDEC measurement technique for measuring case to ambient thermal impedance requires that the part be mounted on a 101.5mm X 114.5 mm 2S2P (2 signal, 2 power planes) printed circuit board. The board and part are horizontal with no air flow restrictions. For our PBGA (plastic ball grid array) packages, heat is dissipated from the parts not only through convection, but also conduction. The center balls on the PBGA package are referred to as thermal balls. These conduct heat from the center of the package, where the die is mounted, to the ground plane. The ground plane and effectively the whole board become a heat sink. As shown in the following table for the 25mm PBGA package, the size of the board, on which the part is mounted, affects the ability to remove heat from the part.

| Card Size | $\theta_{JC}$ (C/W) | $\theta_{CA}$ (C/W) | | |
|---|---|---|---|---|
| mm | average | natural convection | 100 ft/min | 200 ft/min |
| 30 x 30 | 1.5 | 34 | 27 | 23 |
| 50 x 50 | 1.5 | 24 | 19 | 16 |
| 75 x 75 | 1.5 | 20 | 16 | 14 |
| JEDEC 2S2P card 101.5 x 114.5 mm | 1.5 | 17 | 15 | 13 |

Typically the board is much larger than 30mm X 30mm, but there are also other parts generating and conducting heat to the ground plane. The area represented by the card area is sometimes called the shadow area; which is the effective area of a larger board that is available to dissipate heat from the part.

For example if a 25mm 405GP were mounted on a standard PCI card with some SDRAM, an Ethernet PHY and an FPGA for glue logic, presumably the area around the processor would be open for 40mm on a side. This would yield a shadow area of 105mm X 105mm,close to the JEDEC test standard board. On the other hand, if a board has multiple 25mm 405GPs placed within 20 mm of each other the shadow area would be 45mm x 45mm. In the case of the second board there is not only more power to dissipate, but also the power from one part makes it more difficult to cool the adjacent part. This is seen in the increase of $\theta_{CA}$ from 17 °C/W to 24 °C/W for natural convection. Even accounting for shadow area does not include all the factors that affect the ability to cool the part. The numbers in the table assume a JEDEC testing environment that has the board lying horizontal with the part facing up and no adjacent boards. Orientations like those shown in Figure 1 greatly affect the convective contribution of the heat transfer.

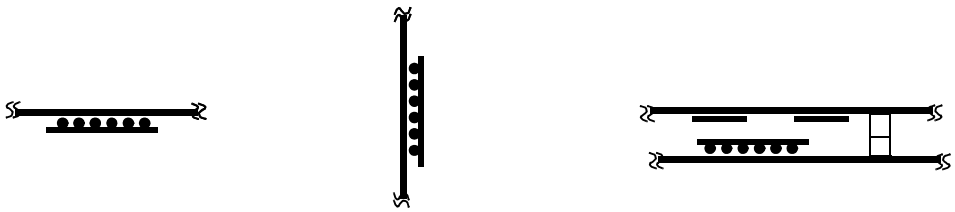

Figure 1

The airflow rates and rise above ambient shown in the table assume that these parameters are measured in the proximity of the part. For example, assume that the system consists of a box with a motherboard and a couple of small riser cards. Placing a thermocouple in the middle of the box and recording the average temperature rise above outside ambient does not account for any localized heating in the area of the processor. Also, measuring inlet or outlet airflow velocity does not account for dead air spaces like those shown in figure 2.

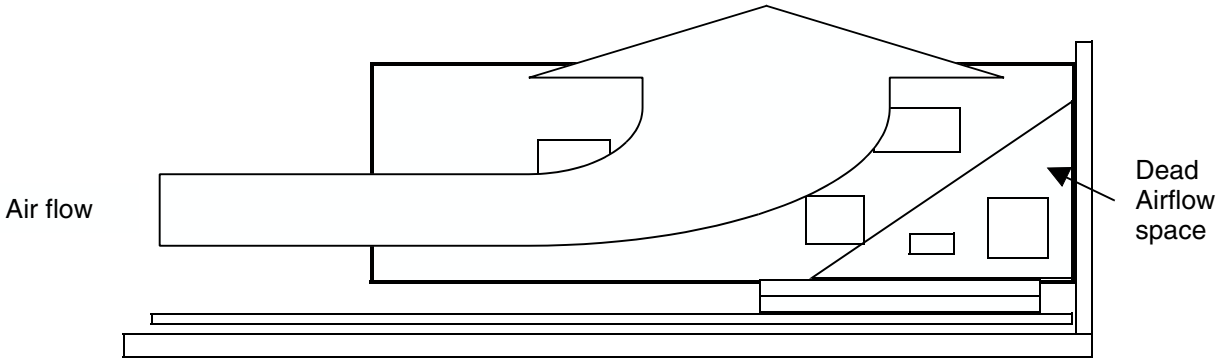


Figure 2 (Side view of system)

Maximum power is determined by two main factors, the logic switching power and the I/O switching power. The logic switching power is based on the switching frequency, the logic transition rate, and the percent of logic circuits switching at the same time. The percent of logic circuits switching simultaneously is also called the switching factor (SF). The currents specified in the data sheet as maximums are determined while running test patterns that produce a high SF and high I/O switching frequency. Maximum currents are also determined at maximum voltages and maximum temperatures and worst case process to account for the effect of transition rate. Some designers perform detailed power measurements, and determine that their application's maximum power will actually be lower than the specified maximum. This can be due to tighter regulation of the supply voltage, or application code that produces a lower switching factor or I/O switching rate than were used to specify maximum power.

Since most of the time the lower power determination is done experimentally it is important to remember that there are three factors that affect the transition rate of the logic: temperature, voltage and process. During a measurement the temperature and voltage can be accounted for, but the process variation cannot.

There can be approximately a 10% increase in logic power due to the affect of process variation on the transition rate. This must be factored into any determination of power based on measurements. According to the data sheet, the maximum case temperature is 85°C at 2 Watts of maximum power with a $\theta_{JC}$ of 1.5°C/W. This means that the junction temperature maximum is 88 °C. Therefore, if the designer knows their maximum power is 1 watt, can they run the part at 86.5 °C; since the junction temperature will still be 88 °C maximum? The official answer is no, the specification that we guarantee is 85 °C case temperature. However, the analysis is correct. In this particular case the difference is small enough that it is not really an issue.

2. Terminations

When are terminations required? And, what value should they be? These are two common questions.

Terminations are required when signal reflections cause unacceptable overshoot and undershoot transients. This is stating the obvious, but there really is no other general answer. As cycle times are decreasing the slew rates on I/O signals are increasing. With the high slew rates of high performance interfaces, such as SDRAM, any impedance mismatch along the signal net causes reflections, except for very short nets.

The AC characteristics of the driver, the number and distribution of loads, the signal net topology, and the circuit board stack (also referred to as the lay-up) affect reflections and therefore the shape of the propagated signal. The best way to determine the need for, and type of, termination is to use signal analysis software. PCB design packages offer analysis software that is aware of the board stack, signal net topology, and loads (based on component profiles).

The last piece of necessary information is the driver AC characteristics, which is commonly provided by component vendors as Spice or IBIS models. IBIS models for most of the IBM embedded processors can be found on the web site at: http://www.chips.ibm.com/techlib/products/powerpc/models.html.

The three most common types of termination are series, parallel, and AC. A schematic representation of each type is shown in Figure 3.

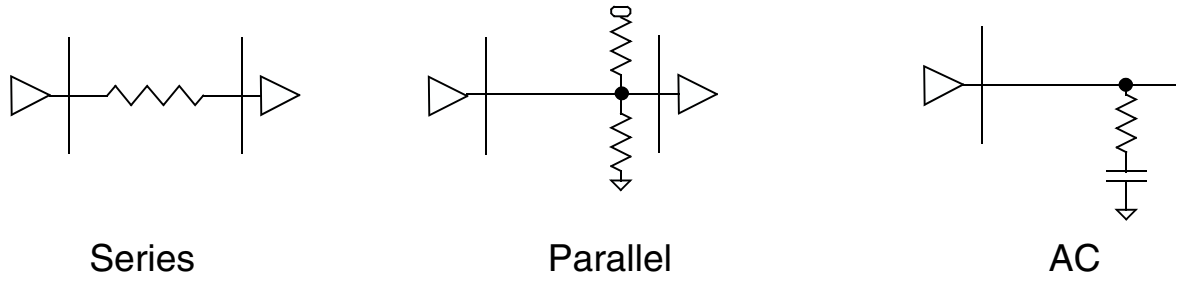


Figure 3

Series termination is typically used when a net consists of one driver and one load. Parallel termination has the disadvantage of having a constant DC power draw. The AC termination is typically used when there is more than one load. The net topology for a multiple load net with an AC termination is shown in Figure 4. Note that the connection is point to point.

If the proper constraints are specified, this type of routing can be achieved with PCB autorouters. Routing a net as shown in Figure 5 will result in multiple reflections that will be very difficult to terminate.

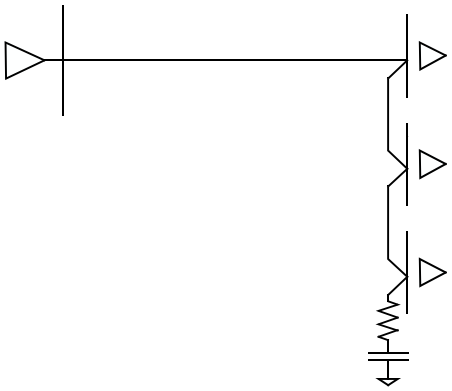

Figure 4. (AC Termination)

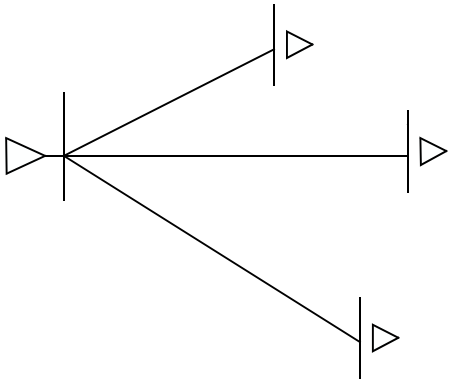

Figure 5.

We provide schematics of the evaluation boards that we design as a beginning reference point. The style and value of the terminations shown on these boards should only be taken as a possible starting point. At times things are done on the evaluation card for test and characterization purposes only. The time and effort spent on proper signal analysis before committing a design to copper may make the difference between a board that functions and one that does not.

3. Decoupling

Proper decoupling in high-speed high performance systems is imperative. Decoupling provides the instantaneous current required by the processor that cannot be supplied by the power supplies due to parasitic inductance. The 405GP requires decoupling of both the 2.5V core voltage and the 3.3V I/O voltage. Even though the I/O current only accounts for 10-20% of the total chip current, decoupling of the I/O voltage is potentially the most critical. Of the 300 I/O on the 405GP, 247 can be outputs. Simultaneously switching large numbers of I/O into large capacitive loads requires significant instantaneous current.

Obviously, the more decoupling the better, but there is a limit to the cost and space available to effectively decouple. A useful rule of thumb is one 0.1uF capacitor for every 8-10 I/O. This recommendation should not be used to replace any standing guidelines companies have developed. In addition there should be a bulk decoupling capacitor for every 5-10 smaller decoupling capacitors. The quantity of capacitors is only one part of effective decoupling. The 0.1uF capacitors should be high frequency monolithic ceramic capacitors. The use of capacitors with greater stability such as those with X7R dielectric is recommended.

Placement is one of the most critical aspects of decoupling. To minimize inductance the cap must be placed at the power and ground pins. Figure 6 shows the placement of the decoupling capacitors on the 405GP evaluation board. These capacitors are placed under the part; which is roughly the area shown, to minimize track length. The ideal configuration for bottom side mounted parts, that will be flow soldered, is shown in figure 7. The bulk decoupling capacitors should be high frequency tantalum capacitors in the range of 33uF. These should also be located right at the part but the constraints of the 0.1" X 0.125" mounting pads forced the placement to about 0.5" from the package to avoid all the signal vias. Both the 0.1uf and the bulk decoupling capacitors use surface mount technology that also provides better high frequency response.

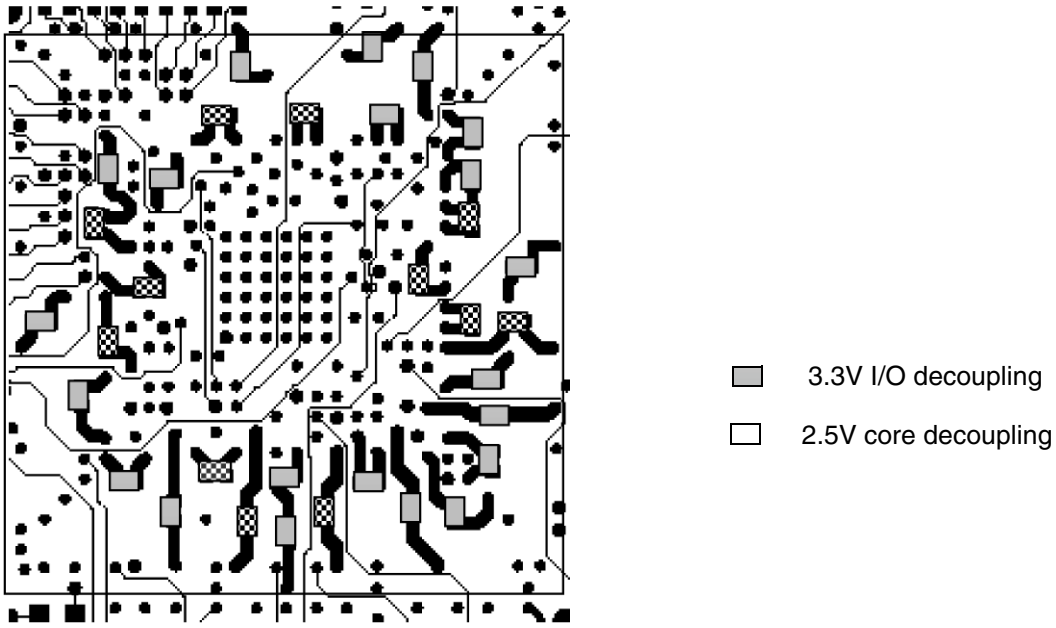


Figure 6. Decoupling Capacitors


Figure 7. Decoupling capacitor pad layout

4. Debug connections

BGA packages have provided a boost in density, but they do bury the pins under the package. Careful consideration must be given before a pad is left unconnected, or an input is tied directly to the power or ground plane. Unless it compromises the ground and power planes it is recommended to make inputs and control pins for interfaces available as a top side pad or as a via. The connections from the processor debug interfaces to the headers, used by the hardware probes, for JTAG and Trace are shown in Appendix A of the RISCWatch ™ User's Guide; which can be found on the web at: www.chips.ibm.com/products/powerpc/tools/rw.html.

5. PLL Filtering

The 405GP contains a PLL for internal clocking. The voltage supply to this PLL must be filtered to minimize jitter. The low pass filter shown in figure 8 is a circuit example.

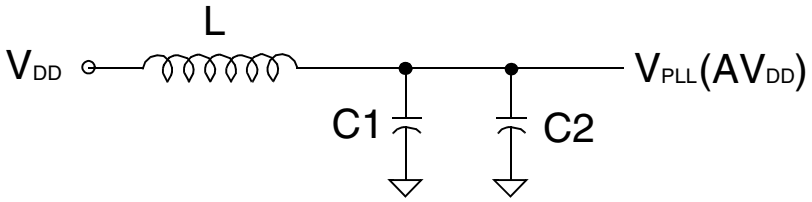


Note: L is a 2.2uH SMT inductor equivalent to MuRata LQH3C2R2M34
C1 is a 0.1uF SMT monolithic ceramic cap with X7R dielectric or equivalent
C2 is a 10uF SMT Tantalum capacitor

Figure 8. PLL voltage filter

The filter reduces any supply voltage noise or ripple by –3.8 db at 40KHz and continues to roll off at –20 db/decade. Initially the filtering is provided by the tantalum capacitor, C1, but as frequencies increase the parasitics make it look less like a capacitor. That is why the high frequency monolithic ceramic capacitor, C2, was added. Minimization of parasitic effects is also why a higher quality dielectric is recommended. The inductor, L1, was also chosen to have a low DC resistance in the range of 0.1 ohms. It is important to minimize the DC drop across the filter network to provide the maximum voltage available to the PLL. It is also important to locate the filter close to the $AV_{DD}$ pin to minimize noise coupling into the filtered voltage. Adding guard traces around $AV_{DD}$ tied to ground provides further isolation.