



Intel[®] Technology Journal

Intel[®] Centrino[®] Duo Mobile Technology

Intel[®] 945GMS Express Chipset for Small Form Factor Platform Based on Intel[®] Centrino[®] Duo Mobile Technology

Intel[®] 945GMS Express Chipset for Small Form Factor Platform Based on Intel[®] Centrino[®] Duo Mobile Technology

Suresh Subramanyam, Mobility Group, Intel Corporation
Taninder Sijher, Mobility Group, Intel Corporation
Sidharth Krishnama, Mobility Group, Intel Corporation
Parthasarathy Ramaswamy, Mobility Group, Intel Corporation
Deepa Mohan, Mobility Group, Intel Corporation
Vikas Shilimkar, Mobility Group, Intel Corporation
Satish Prathaban, Mobility Group, Intel Corporation
Eric Samson, Mobility Group, Intel Corporation
Michael Derr, Mobility Group, Intel Corporation
Samir Gundawar, Mobility Group, Intel Corporation

Index words: Intel 945GMS Chipset, package, reduction, manufacturing, motherboard, routing, power management, DDR2 memory down

ABSTRACT

Small form factor (SFF) platforms aim to position Intel firmly to address the fast growing mini- and sub-notebook market, expected to increase fourfold over the next three years. This trend is proven by the increasing number of design wins for Intel in this market segment. Intel has made significant advances in mobile technology with their SFF platforms based on the new Intel[®] Centrino[®] Duo mobile technology: these include the Intel[®] Core[™] Duo processor Low Voltage/Ultra Low Voltage (LV/ULV), Intel[®] Core[™] Solo processor ULV, and the Intel[®] 945GMS Express Chipset. To support the mobility vectors¹ of battery life and performance, several power/performance features were designed into the processor, chipset, and platform architectures. This paper describes the power management features of the Intel 945GMS Express Chipset Graphics and Memory Controller Hub (GMCH), the Intel[®] 82801GBM/GHM I/O Controller Hub (ICH), and the overall platform, all of which have reduced their power consumption while allowing for greater performance. Intel has designed the Intel 945GMS Express chipset to address the mini- and sub-notebook market segments. Given that the package

size had to be smaller or equal to that of the Intel[®] 915GMS Express Chipset (27mm²) [2], the package design had additional challenges to meet the electrical requirements of the higher frequency Front Side Bus (FSB)/Dual Data rate (DDR) memory and the additional 27 signals. In this paper we discuss the enhancements at the platform level, such as the processor steeper load line and the onboard memory decoupling, which help reduce the Bill of Materials (BOM) cost and save real estate. We also discuss in detail the platform signal and power integrity that enable the DDR2-533 MHz Small Outline Dual In line Memory Module (SODIMM) with onboard memory and decoupling methodology.

This paper serves as a guideline for system designers to understand the benefits of the features described and design better mini and sub notebooks. It also provides an insight into the challenges as package sizes shrink and interface speeds grow. With this SFF platform, Intel reaffirms its strong focus on the four vectors of mobility: breakthrough mobile performance, integrated wireless LAN capability, great battery life, and thinner/lighter design.

INTRODUCTION

The Mobile Intel 945GMS Express Chipset, a member of the Mobile Intel 945 Express Chipset family, is a specially designed Graphics and Memory Controller Hub (GMCH)

¹ Intel Centrino mobile technology is based on four vectors of mobility: breakthrough mobile performance, integrated wireless LAN capability, great battery life, and thinner/lighter design.

targeted for use with small form factor (SFF) platforms. The package and features of the Mobile Intel 945GMS Express Chipset are optimized for size, power, and performance to best suit SFF applications.

The feature set of the Mobile Intel 945GMS Express Chipset makes it a compelling product in mini- and sub-notebook segments for the following features.

- The Intel Core Duo and Core Solo processor Low Voltage (LV)/Ultra Low Voltage (ULV) support with a Front Side Bus (FSB) speed of 667 MHz and 533 MHz, respectively.
- DDR2 533 MHz single channel memory interface. The maximum memory supported is 2 GB.
- Integrated Graphics with a 166 MHz pixel clock. Integrated graphics supports CRT, TV, dual channel Low Voltage Differential Signaling (LVDS), and Serial Digital Video Output (SDVO) interfaces.

The Intel Core Duo LV and Core Solo low-voltage ULV processors are high-performance, low-power, mobile processors with several micro-architectural enhancements over existing Intel® mobile processors. The ULV Intel® Celeron® M processor is a low-power mobile processor with good performance targeted towards value products.

The Mobile Intel 945GM Express Chipset family optimized for SFF-based designs contains two core chipsets: the Mobile Intel 945GMS Express Chipset and the ICH7M (I/O Controller Hub). The Mobile Intel 945GMS Express Chipset provides the host interface controller, System Memory Interface (SDRAM), Direct Media Interface (DMI), and an integrated graphics engine with display output ports.

The ICH7M integrates a number of I/O device controllers and interfaces for legacy and high-speed devices to provide system design flexibility. The DMI, the chipset component interconnect, is designed into the chipset to provide an efficient high-bandwidth communication channel between the GMCH and the ICH7M. The ICH7M also supports PCI Express* x1 I/O ports, next-generation Serial ATA (AT Attachment), and Hi-Speed USB* 2.0 connectivity. An Advanced Configuration and Power Interface (ACPI)-compliant Mobile Intel 945GMS chipset platform can support the Full-On (S0), Suspend to RAM, Suspend to Disk (S4), and Soft-Off (S5) power management states, processor C states, and PCIe-based power management. Through the use of an appropriate LAN device, the chipset also supports wake-on LAN for remote administration and troubleshooting.

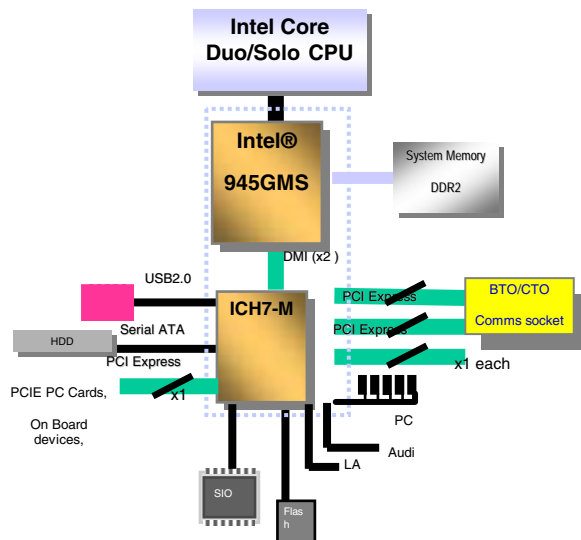


Figure 1: Intel Centrino Duo mobile technology-based SFF platform

In this paper, we discuss the benefits of the Mobile Intel 945GMS chipset over the Mobile Intel 915GMS chipset [2] in terms of the features supported, the challenges in the Intel 945GMS Ball Grid Array (BGA) breakout, motherboard routing, and manufacturability. In comparison to the Mobile Intel 915GMS, the Mobile Intel 945GMS had an additional 27 signals to be routed on the package. Also, the Intel 945GMCH die size is bigger than that of the Intel 915 GMCH by about 15 mils on one side and 22 mils on the other side. The development efforts to deliver the smallest package, meeting all the electrical, layout, and manufacturing requirements are key to Intel's competitive advantage, while pushing the design envelope for thinner, lighter notebooks.

The platform signal and power integrity analysis to enable the DDR2-533 MHz SODIMM with onboard memory and decoupling methodology are discussed in detail. As real estate and BOM costs are key factors in SFF platforms, we describe the platform power delivery with a steeper load line technique and the motherboard decoupling for the Intel Core Solo ULV processors. A steeper load line implementation results in a 50% reduction in the processor decoupling BOM cost compared to the previous-generation SFF platform. Despite the significant improvement in performance in terms of FSB-667 and DDR2-533, the advanced platform power-management features enable lower platform power, compared to previous platforms. The power-management innovation in the Mobile Intel 945GMS Chipset and ICH7M are also discussed in greater detail. Lastly, we discuss the rising challenges foreseen in power management, packaging, design, routing, and electricals.

OVERVIEW OF THE MOBILE INTEL® 945GMS PACKAGE

The Mobile Intel 945GMS Chipset package is 27 mm x 27 mm in size (998 pins). A regular Mobile Intel 945GM Chipset die is used for the Mobile Intel 945GMS package. The biggest challenge was reducing the size of the package since this chipset has an additional pin count (VCCAUX power rail) and more features (LVDS second channel) than the previous-generation Intel 915GMS Chipset (840 pins) [2]. A detailed study of various footprints was done to overcome these challenges. Accommodating additional pins when compared to the Intel 915GMS Chipset in the same package size was possible only by reducing the pin pitch and BGA pad size. Hence, the Mobile Intel 945GMS Chipset was designed with a 0.8 mm uniform pitch and a 14 mils pad size. This smaller pin pitch also introduced a lot of routing complexities. A thorough analysis of motherboard breakout and power-delivery shapes was done to overcome these complexities. This footprint uses the parquet package technique to meet the 27 sq. mm using an 0.8 mm ball pitch. This gives a parquet region of 1.2 mm, utilized as routing channels. Figure 2 shows the uniform 0.8 mm pitch, 27 mm x 27 mm 945GMS package.




Figure 2: 0.8 mm uniform pitch package

The Challenges of an 0.8 mm Pitch in Motherboard Routing

As discussed earlier, the 0.8 mm (32 mils) pitch poses a significant challenge in motherboard routing. Due to the smaller pitch, it is possible to route only one trace between the two vias, assuming that a 10/20/28 mil via is used, as shown in Figure 3.

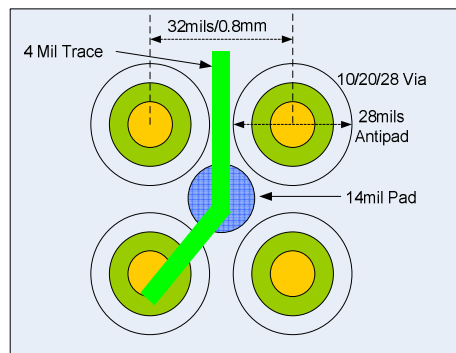


Figure 3: Breakout of one Trace

Therefore, it is difficult to route a 1-byte lane of DDR or FSB signals on one layer. However, this is achieved by routing channels in a special way in the package. Reduced pin pitch provides 70 additional pins in the 27 mm x 27 mm package, which are distributed as No Connect (NC) pins, and additional power pins provide the routing channels for the byte lane in the inner layer as shown in Figure 4. These NC pins or power pins do not require a PTH via, resulting in routing channels that enable to route five traces of 4 mils width through these channels. If those pins are power pins then those corridors are used as power entry corridors on the Top or Bottom layers to feed the power shapes on the motherboard from the Voltage regulator to the GMCH pins, as shown in Figure 4. The Top layer power corridors are used to feed the power to the pins in the parquet region, and Bottom layer copper shapes can be used to feed the power to VCC Core pins located in the center array region.




Figure 4: Routing channels in 0.8 mm pitch

Intel® 945GMS Differential Breakout Challenges and Mitigation

The 0.8 mm ball pitch for the Intel 945GMS requires differential signals to break out of the BGA pads on the motherboard as single ended signals for Serial Digital Video Out (SDVO) and Low Voltage Differential Signal (LVDS) interfaces. This single ended breakout implies that the differential signals go through a higher impedance than the target differential 100 ohms till the breakout ends.

Figure 5 shows a snapshot of the motherboard breakout. The impact has been analyzed in simulations and the timing margins and signal quality impact assessed. The design for robustness requires that the length of the single ended breakout for differential signals be restricted to 10 mm. The impact is particularly significant when the motherboard trace lengths are smaller. The MCH driving signals with lower motherboard trace lengths is the worst case condition. The jitter and noise margin numbers for all worst-case conditions show a degradation of 5-7% at the receiver, which is within acceptable limits.

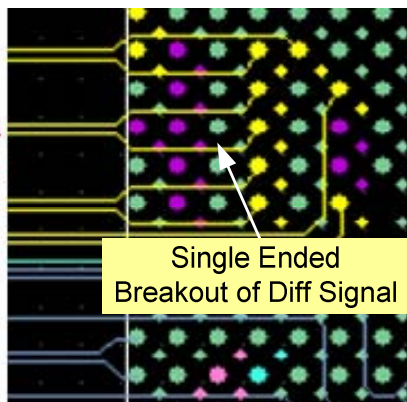


Figure 5: Single-ended routing of differential signals

Challenges in Manufacturing and Solder Joint Reliability (SJR)

Due to smaller pin pitch, the Via to BGA pad spacing is reduced to below 5 mils and hence it adds to the risk of solder bridging during the assembly process. Therefore, efforts have been made to reduce the BGA pad size to 14 mils, which increases the BGA pad to via spacing to 5.6 mils.




Figure 6: Via to BGA pad spacing

It is also suggested to either encroach the via pad with solder mask in the BGA region completely, or at least partially, to reduce any risk of solder bridging. As the pin pitch and BGA pad size for the Intel 945GMS is lower

than the same parameters for Intel 945GM and Intel 915GMS [2], the solder joint reliability for shock and thermal stresses was analyzed by the Advanced Technology group. It was recommended that the five balls in each corner and the center array (die shadow) balls should be sacrificial balls. These sacrificial balls, if used for non-critical to function (NCTF) power pins, should be used as metal defined (MD) pads (50% solder mask opening). Hence, all the center array pins are assigned as NCTF power pins, which serve as additional power pins.

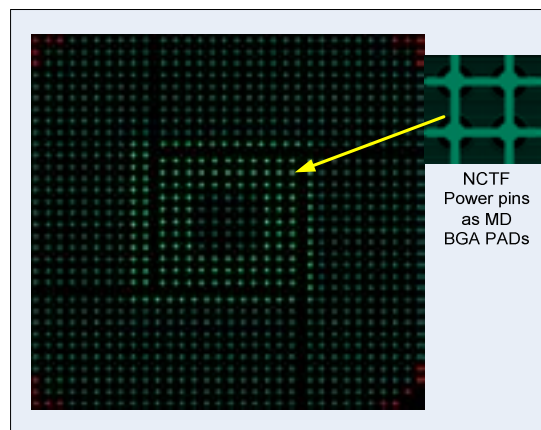


Figure 7: NCTF MD pads for solder joint reliability

INTEL® 945GMS SUBSTRATE LAYOUT AND TECHNOLOGY CHALLENGES

Keep Out Zones

There are certain areas on the surface of the package that need to be clear of any kind of components such as die, capacitors, or materials, such as epoxy. These areas are used for the handling of parts in package assembly and testing and are called Keep out Zones (KOZs). These zones are primarily driven by position of die and the package edge. On the Intel 945GMS Express Chipset with an Intel 945GM Express Chipset die and a reduced package size, KOZs become a considerable challenge as they start overlapping with each other and leave no space for critical components such as package capacitors. In order to resolve some of these concerns, the reduced KOZs that were used in the Intel 915GMS chipset, after a considerable number of experiments, were adopted.

Another challenge arising from KOZs was overlap of the underfill tongue with the package KOZ. Underfill is an epoxy material that is filled under the die for reliability, and it leaves an extended tongue from the side of the application which is east of the die on the 945GMS chipset. This underfill flow was reduced by cutting a trench on the solder mask layer of the package that reduced the capillary effect. This is similar to what was implemented on the 915GMS chipset, but the overlap is

more, due to the larger die size on the 945GMS. DDR and FSB stripline routing in that region was another key factor in enabling implementation of this trench. Figure 8 shows the KOZs and trench.




Figure 8: 945GMS KOZ and trench

Layout

Compared to the previous generation Intel 915GMS chipset [2], the Intel 945GMS chipset, in addition to enhanced features like faster DDR and FSB interfaces, also has additional features like dual channel LVDS which resulted in increased signal count on the same package size as the Intel 915GMS. The package capacitor count has increased to 15 compared to 9 on the previous generation Intel 915GMS. This poses additional routing challenges given the same routing area as 915GMS. Die area on the Intel 945GMS has increased by about 10% over the Intel 915GMS. With all these challenges and requirements the package size was maintained at 27 mm sq.

Intel 945GMS Package Layer Stack-up

The biggest challenge the package development team had was to deliver the smallest possible package size that still met all the electrical requirements for high performance and at the same time stayed within the envelope of manufacturing technology capabilities. The primary drivers enabling smaller size was an 8-layer package with stripline routing for single channel DDR and FSB

interfaces which forms the bulk of the IO count. Compared to the Intel 945GM package, with dual channel DDR, which had a 44 μm trace width and a 54 μm spacing in microstrip resulting in 55 ohms impedance, equivalent stripline routing in the Intel 945GMS resulted in a 26 μm trace width and a 30 μm spacing, resulting in 50 ohms single ended impedance. With 26 μm being the minimum trace width allowed in manufacturing, a 5 ohms mismatch between package and the motherboard was treated as low risk from a Signal Integrity (SI) perspective. Reduced spacing to 30 μm is compensated by reduced trace-to-trace coupling in stripline. Figure 9 and Figure 10 show the microstrip and stripline structures on the package.




Figure 9 : Microstrip signal routing




Figure 10: Stripline signal routing

Intel 945GMS Substrate Power Delivery Challenges

While taking advantage of the technology developments made on the 915GMS package (KOZ reduction, shift to 0201 FF package decoupling capacitors, 8L stripline), the Power Delivery (PD) team continued to push design boundaries in order to deliver next-generation performance features. Additional package decoupling capacitors (see Figure 11) and improved power and return path routing on the internal substrate layers helped control AC loop inductance to acceptable levels. The reduced impedance versus frequency response of the power delivery network allowed the PD team to meet stated design targets, thus helping to enable FSB 667 MT/s and DDR 533 MT/s on the Mobile Intel 945GMS Express Chipset.




Figure 11: Substrate top layer with package decaps

Another significant technology enabler was the shift to a uniform 0.8 mm (32 mil) ball pitch. The PD team was able to take advantage of the resulting increased pin count to opportunistically assign VCC “power corridors” and VCC islands in the package (see Figure 12). These power corridors allow a wide uninterrupted copper flood on the surface and bottom layers of the motherboard and serve a dual purpose. They can accommodate higher currents, thus improving DC and AC response of the power delivery network.

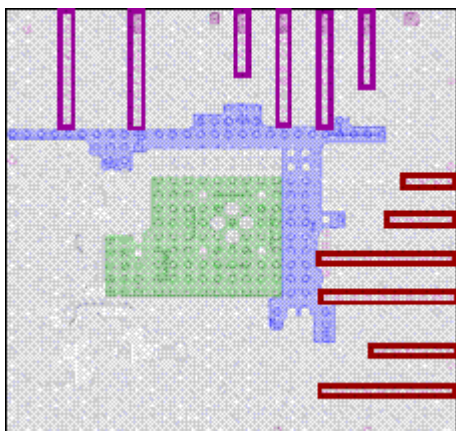


Figure 12: Substrate power corridors and VCC islands

Another benefit of the corridors is that they allow higher signal density breakout on the inner stripline layers of the motherboard due to that region being absent of vias. The

absence of a via field on internal signal layers creates additional routing channels, and this allows for greater signal density within a smaller package size. Increased ballout flexibility with the 32 mil ball pitch also enabled the PD team to uniformly distribute “edge” decoupling capacitor pins along the periphery of the package. These capacitors utilize the much smaller substrate routing for current return paths, thereby increasing their effectiveness and improving the robustness of the overall power delivery solution.

PLATFORM POWER DELIVERY

SFF platforms are designed for the mini- and sub-notebook market segment wherein power (battery life), real estate, and BOM costs are the key factors. This section explains how these aspects are taken into account while designing a platform power-delivery solution. Power-delivery architecture for the platform and different power-delivery optimization techniques are discussed.

Platform Power-Delivery Architecture

All the products are battery based and use an adapter with a voltage that varies from 16-28 V. For such a large variation in the input voltage of all the voltage regulators, it is hard to predict the efficiency of the power conversion. In case of higher input voltages and lower output voltages, the efficiency is very low. To resolve the issue, the platform implements a narrow VDC technique in which input to the VR controllers is regulated to give 12.6 V. With this technique, the efficiency of the VRs can be closely controlled, and higher thermal efficiency is obtained as there are fewer thermal losses. This significantly improves battery life. Another important aspect of power-delivery architectures is explained below.

Processor Core Power Delivery

Among the multiple ways to optimize the power-delivery network, is to design the network for higher DC and AC impedance targets to achieve greater savings on real estate and cost.

All the power-delivery network optimization techniques are aimed at a lower droop in order to meet V_{min} specifications at the die of the device. This droop basically depends on DC and AC impedance of power delivery networks over the operating range of frequency. Figure 13 shows the relationship of core voltage with maximum frequency.

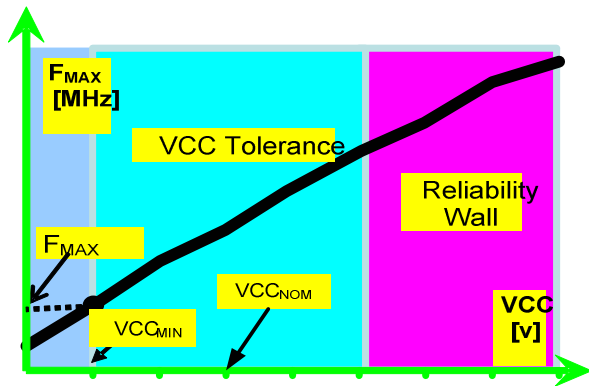


Figure 13: Fmax dependency on Vcore-min

As is clear from the figure, to maintain the performance, i.e., high speed, a particular minimum voltage needs to be maintained on the core rail. This gives a maximum frequency at which processor core logic can operate without failure. The maximum voltage or overshoot on the core rail is limited by the metal oxide limit of the particular silicon technology as shown by the reliability wall.

The processor platform power-delivery solution comprises a single-phase voltage regulator with an adaptive voltage positioning technique in which the power-delivery network is designed for a steeper load line in the case of ULV.

The ULV processors maximum core current has been significantly reduced from 36A (for Standard Voltage processors)/20A (for LV processors) to 8A (for ULV processors) and the power-delivery network can be designed for a load line with a steeper slope. The concept is explained in Figure 14.

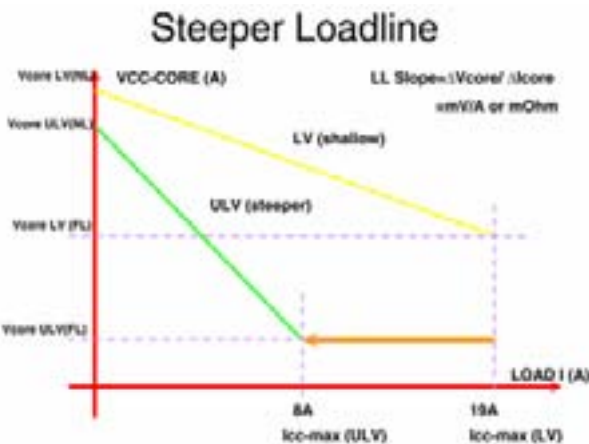


Figure 14: Steeper load line (LL) concept

$$\text{Load line slope} = \Delta V / \Delta I,$$

As $\Delta I \downarrow$, LL can be increased keeping V_{min} the same as shown in Figure 14.

As the core maximum current is reduced, the load line gets steeper maintaining the V_{min} (refer to Figure 14). Due to the steeper load line, the capacitance required to meet the higher AC impedance is reduced. Therefore the bulk and Multi Layer Ceramic Chip (MLCC) capacitor recommendations for motherboard decoupling are reduced by more than 50%, saving real estate and BOM costs, which are very significant factors in the SFF platform designs.

Power Delivery for Onboard Memory

The Intel 945GMS Express Chipset supports only a single channel 4-rank DDR2 interface, with a maximum memory support up to 2 GB (1 GB per two ranks) at 400 MHz and 533 MHz, respectively. In this section, we discuss the decoupling methodology for the onboard memory configurations that use a total of eight components where four components are on the Top and four are on the Bottom.

Onboard Memory Decoupling Methodology

For the eight DDR2 SDRAMs directly assembled on the motherboard with the four devices on Top and four devices on the Bottom, extensive simulations were done to obtain the optimum decoupling solution. The motherboard is an 8-layer board with Layer 1 (L1), Layer 3 (L3), and Layer 8 (L8) providing the power delivery for the DDR2 core and the I/O. L1 and L8 are the primary layers providing the power to the device pins. The internal layers, L4 and L5, cannot be used for power planes as the memory interface signals in L3 need to have ground referencing on L2 and L4, and signals in L6 need to have ground referencing on L5 and L7 layers. A frequency domain analysis was done on the layout to achieve constant impedance up to the DDR2 operating frequency. A time domain analysis, including single pulse, even mode switching, and worst-case patterns, was performed assuming ideal transmission lines, to verify if the voltage requirements of the devices were met. The DDR2 voltage regulator supplies 1.8 V with a load line of -8.75 mE for a max of 8 A. The VR is modeled as a simple Resistor-Inductor-Capacitor (RLC) network using an extended adaptive voltage positioning concept. The RLC parameters are extracted from the layout to obtain the impedance profile near the pins of the device. The impedance profile is observed at the farthest device VCC pin from the VR output. The simulations were carried out to meet the impedance profile of around 6.75 mohm (3% of 1.8V/8A) up to the operating frequency of the device so that the voltage droop is within the specification of the device. Figure 15 shows the layout implementation for the 1.8 V plane.

We analyzed the generic layout with VR placements both on the East and North of the onboard memory devices

giving the customer maximum VR placement flexibility. The decoupling solution reduced the BOM cost compared to previous platforms providing optimum power-delivery solution.




Figure 15: Power layout for memory down

INTEL® 945GMS DDR2 SDRAM SYSTEM MEMORY INTERFACE

The Intel 945GMS Express Chipset supports only a single channel 4-rank DDR2 interface, with a maximum memory support up to 2 GB (1 GB per two ranks) at 400 MHz and 533 MHz, respectively. The DDR2 memory interface consists of four signal groups: clock, command, control, and data. Figure 16 shows the block diagram of the memory interface. Intel gives design recommendations and routing guidelines for various SODIMM and onboard memory configurations for this chipset. In SFF platforms, implementing onboard memory along with the SODIMM provides a significant reduction in the space utilized and the z height. In this section we cover the signal integrity challenges inherent in enabling the 2rank-2rank DDR2 533 memory interface, the key parameters that need to be considered, and some of the best-known routing methods.

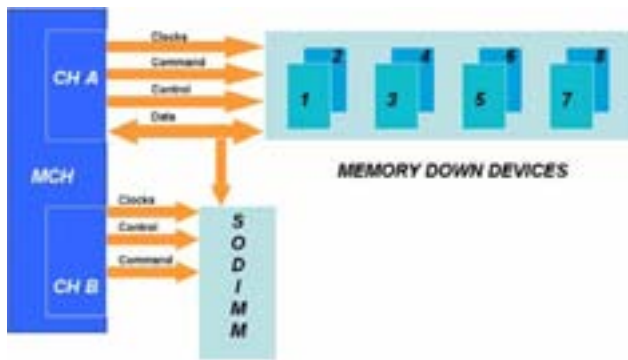


Figure 16: DDR2 memory interface overview

Signal Integrity and Design Recommendations

For the data signal group, being source synchronous, minimizing the flight time skew within the same group of signals is a key parameter that allows their high-frequency

operation. As the interface speeds and memory density and loading get higher, the increased skew reduces the timing margins between the data and the strobe signals. In order to achieve good timing margins and also allow maximum routing flexibility in terms of trace lengths, the loading on the data lines needs to be restricted. For example, with a dual rank SODIMM and onboard memory configuration, a user could go with 1-rank on an onboard memory instead of two ranks to get the same memory size. This reduces the loading on the data bus from four SDRAMs to three. The 3-load scenario reduces skew between data and its strobe significantly. However, there is a tradeoff on the power dissipated, as all devices in the 3-load case need to be ON. We switch OFF all devices not being accessed in the 4-load case. However, by restricting the trace lengths between the SODIMM and the onboard memory to less than one inch, good timing margins can still be obtained with four loads on the data lines, but this requires the placement of the onboard memory underneath the SODIMM. Customers have demonstrated such designs by having an optimum cooling solution. Crosstalk is another important parameter to be considered, and it is observed that crosstalk could impact timing significantly with a worst-case timing hit of the order of 200 ps. Increasing the spacing between the data lines reduces crosstalk. Since real state on mobile platforms comes at a premium, designers should choose the right spacing to satisfy all aspects including signal integrity. Using smaller values of the on-die termination (ODT) and having a tighter tolerance on the ODT values yields better timing margins, but the tradeoff is increased power dissipation.

For the onboard memory clock signal group, we recommend a routing topology that yields good differential slew rates of more than 2.0 V/ns with monotonic rising and falling edges and positive noise margins. The signal integrity of the clocks has an impact on all the other signal groups. As the trace lengths of certain routing segments increase, CLK rings back into the DC thresholds causing slew rate degradation. This is overcome by restricting the lengths of certain segments, keeping in mind the layout feasibility, and choosing optimum values for the differential capacitance.

All the data lines that belong to the same group should be routed on the same internal layer for the entire length of the bus. This practice results in a significant reduction of the flight time skew, since the dielectric thickness, line width, and velocity of the signals will be uniform across a single layer of the stack-up. DQ Bit swapping within a byte lane is allowed whether the routing is to the SODIMM or onboard SDRAM components. Data Bit swapping works best for the memory down configurations that use a total of eight components where four components are on the Top and four are on the Bottom. With bit swapping the lengths to the Top and Bottom

SDRAM components can be routed to exact lengths for matching and also ease of routing. Figure 17 illustrates the bit swapping for data lines for the memory devices placed back to back.




Figure 17: Bit swapping in data lines

GMCH POWER MANAGEMENT

The challenge in every new generation of chipsets is how to increase the platform performance while also increasing the platform battery life. Platform battery life is measured using several benchmarks. A typical benchmark is the Mobile Mark 2005 (MM05). Figure 18 shows system memory bandwidth versus time for this particular benchmark. It is apparent that the GMCH workload primarily consists of a relatively constant demand for a small percentage of read bandwidth, and a much smaller write bandwidth.

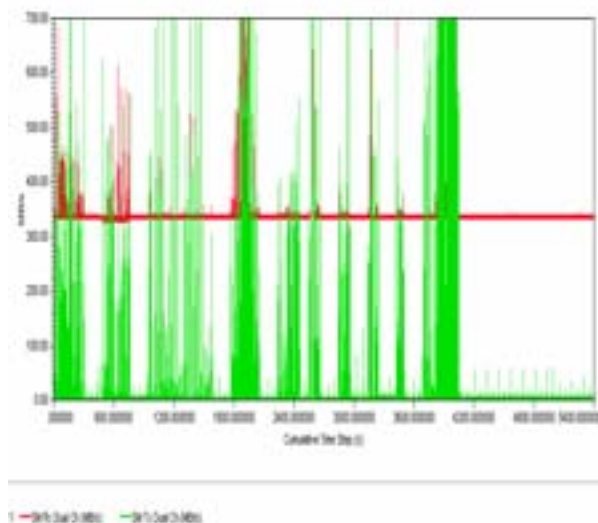


Figure 18: Memory traffic vs. time for MM05 OP

It turns out that while there is negligible usage of the graphics engine during the benchmark (therefore that power component can be ignored as long as there is good dynamic clock gating of this engine), the monitor/flatpanel needs to keep being updated at the required display refresh rate even if nothing is visually changing on the monitor/flat panel. Display refresh requires relatively power hungry accesses to the external memory, since the display frames are too large to store internally to the GMCH, and this represents the majority of the read bandwidth.

As the FSB and memory speeds increase, GMCH internal clock frequencies also increase, and hence their possible contribution to average power also increases, especially if a voltage increase is required to maintain critical timings.

Finally, another potentially large contributor to platform power is bus mastering from attached devices traffic (USB, PCIE, FIR, Audio).

These types of workloads are a small fraction of the MM05 benchmark, but a badly configured platform can easily suffer a significant power impact. This is because these types of workloads (especially if they are making unnecessary memory accesses when they are idle) have the potential to force the platform into a higher platform power state that is required to provide the short memory request latency and CPU snoops they need.

In summary, the solution to lowering the GMCH average power contribution is to do the following.

- Keep core voltages as low as possible.
- Minimize the impact of increasing I/O frequencies.
- Keep the overhead of display refresh power low.
- Minimize the impact of bus master traffic.

GMCH Power Management Features

As in previous generations of the Intel[®] mobile chipset [1], a large factor in achieving these goals is via main memory power management during normal operation and in low-power ACPI Cx states. This includes such power-management features as Dynamic Memory Row Power Management (DRPM) and Conditional Memory Self-Refresh (CMSR). CMSR mode includes dynamic gating of most GMCH clocks, since a MHC has no useful work to do whenever memory can be placed in a self-refresh state. In an analogous fashion, I/O circuits are also placed in lower power states.

One of the most efficient ways to control power is to minimize voltage. The 82945PM (targeted at larger form-factor platforms requiring up to two channels of higher speed memory), for example, needs to run some of its core logic at a higher voltage than the rest of the chipset core in

order to operate at higher FSB and DDR frequencies. The Intel 945GMS Express Chipset runs its core at the lowest voltage that can meet its performance requirements.

Given that display refresh must be kept serviced and is the predominant memory traffic, the GMCH attempts to maximize the amount of time spent in the CMSR state, and minimize the amount of time spent exiting the CMSR state. One way the GMCH does this is by using a large First In First Out (FIFO) Random Access Memory (RAM) to store very quick but large read bursts from memory so that the FIFO contents can be used to feed the slower display pipeline while the GMCH and memory are in the CMSR state. The self-refresh efficiency (= time spent in self refresh/total time) improves as a function of this FIFO size, as shown in Figure 19, although this is modulated as a function of the exit time from self refresh mode.

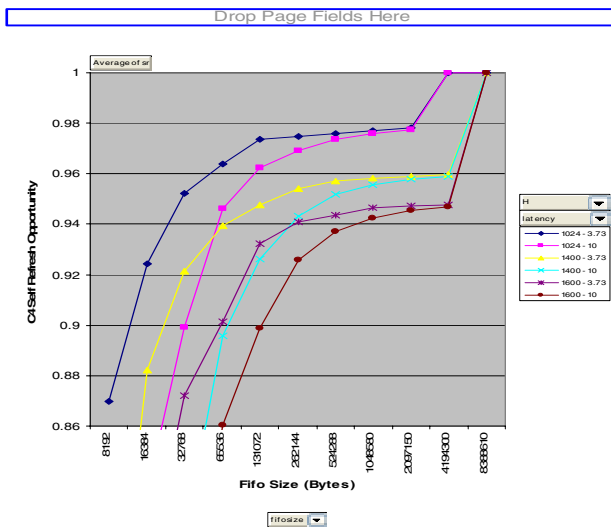


Figure 19: Self-refresh efficiency

Previous mobile chipset generations automatically jumped and stayed in the high-power CPU C2 state during bus mastering traffic. The only way back to a lower power C3 state had also required a jump back to the C0 (highest power of the C states) where the decision to go back to the C3 state could be made by the ACPI software.

The Intel 945GMS Express Chipset has increased the amount of time that the platform is in a lower power state by not always requiring the platform to go fully into the C2 state to handle snoops. Instead, only those portions of the platform required to handle snoop traffic are awakened momentarily to handle the snooping.

ICH Power Management

The ICH7M (82801GBM and 82801GHM) builds upon previous generations. It continues to implement the advanced platform power-management control functionality that was present on the previous generation.

This functionality aggressively lowers the processor, GMCH, and clock chip power consumption during brief periods of idleness through control signals to the clock chip, processor voltage regulator chip, GMCH, and the processor. Additionally, the ICH7M integrates other features and new capabilities that are valuable in the SFF segment through power and board area savings.

The ICH7M implements aggressive on-die clock gating. With many independent I/O subsystems integrated on the ICH7M die, the clock-gating is implemented in a distributed manner such that particular branches of a clock tree are shut off to those subsystems that are idle. This reduces idle and average power consumption.

The ICH7M also reduces the core operating voltage from 1.5 V on the previous generation to 1.05 V. This significantly reduces idle, average, and thermal power, all of which are critical to the SFF segment.

The ICH7M also integrates on-die regulators and power switching logic that avoid external components and allow for efficient power distribution. Because the ICH7M (and previous generations) contains system suspend/resume (for S3, S4, and S5 states) control logic, the system real-time clock, system CMOS, and the LAN MAC support multiple isolated power wells that must stay on when the main power supply is off. The ICH7M contains integrated voltage regulators that can supply low-voltage rails from the 3.3 V rail that is available externally in these suspended and off-power states. The platform only needs to supply the 3.3 V rail for each of these power wells (avoiding switches or VRs on the board), while the ICH7M reduces the voltage—and therefore power—on the silicon die. In the S0 state, the ICH shuts off the internal regulators and switches over to the equivalent voltage rails in the main power well. This guarantees that the efficient power supply infrastructure that is in place for the higher power S0 state is utilized, thereby minimizing on-die and platform power loss (idle, average, and thermal) due to power regulation inefficiencies in S0.

The ICH7M also adds features that reduce the system flash footprint on the board. First, the ICH7M adds a Serial Peripheral Interface (SPI) for the system flash as an alternative to the existing Low Pin Count (LPC) interface. This enables the use of high-volume, SFF, low-cost SPI flash devices from a variety of vendors. Secondly, the SPI flash controller in the ICH7M supports flash-sharing with an Intel® LAN NIC such that a single SPI flash component contains the LAN configuration information as well as the system BIOS data. This saves one non-volatile memory component on the board.

Silicon Results

As a result of the chipset power optimization work, the average power of the Mobile Intel 945GMS Express Chipset has gone down from that of the previous generation, all the while providing increased platform performance.

Table 1: Silicon results

	ICH MM05 Power	GMCH MM05 Power
Intel® Centrino® mobile technology platform (Previous Generation)	1.5 W	1.5 W
Intel® Centrino® Duo mobile technology platform	.67 W	1.2 W

CPU AND OTHER POWER MANAGEMENT FEATURES

Other than 945GMS and ICH power-management techniques discussed in the previous section, this platform also supports a few other power-management techniques discussed below.

Intel® Media Storage Manager (IMSM)

This is the power management technique used in the SATA interface of ICH7M. It uses the link power management of the SATA interface by changing the physical SATA link to lower power states. This technique can provide a power savings of up to 400 mW in AHCI mode.

Active State Power Management (ASPM)

An active state power management feature contained in the PCIe* specification is supported by all the PCIe x1 links of ICH7M. The two low-power states, L0s and L1, also called standby states, are supported by ICH7M in ASPM-enabled systems. L0s is a short entry and exit latency state where any transmit link can be placed into a low-power state if there is no traffic on the link. The L1 state is optimized for more power savings than the L0 state, but it comes with the cost of longer exit latencies. ASPM is supported by ICH7M but needs to be enabled by the software.

Processor Power Management

LV/ULV processors support C1/AutoHALT, C1/MWAIT, C2, C3, C4 for optimal power management. Figure 20

explains the processor power management states and signals that transition the state. Enhanced Intel SpeedStep® technology features multiple voltage/frequency operating points for optimal performance. The processor voltage regulator is designed for smooth voltage transition with low latency. Depending on the thermal sensor, the processor can change the core voltage with voltage identifiers. This provides better adaptive thermal management. Voltage regulator efficiency can be optimized with power status indicator pins that assert when the processor is in a lower power state.



Figure 20: CPU power states

System Power Management

SFF platform supports S0 (active), S3 (stand by-suspend to RAM), S4 (hibernation-suspend to disk) and S5 (soft off).

FUTURE CHALLENGES

As the form factors of mobile platforms are getting smaller with faster processor and memory speeds, power management, packaging technology, platform power delivery, signal integrity, and the motherboard implementation have become more challenging than ever before. These challenges are discussed below.

Power Management Challenges

Even though excellent performance and power results were achieved on Intel Centrino Duo mobile technology systems that contain the GMCH and ICH, both performance and average battery life need to keep increasing.

There are three vectors to be considered:

- Graphics features and performance requirements are continuing to follow Moore's Law at a faster than average pace.
- We need to support ever increasing I/O bandwidths that have surpassed what can be done with unterminated I/O.
- Process technology is increasing leakage as a power component.

Packaging Challenges

To support additional functionalities on the chipset, the pins required on the chipset are continuing to increase. This increase in pin count and demand for reduced package sizes impose a big challenge in the substrate design and manufacturability.

- Smaller package ball pitches and reduced motherboard vias/spacing are key drivers for reducing package size.
- Smaller package size means smaller solder ball sizes due to reduced motherboard pad size and substrate pad size. This reduces the solder joint reliability.
- Smaller packages drive more layers on the package, in turn impacting the cost significantly.
- Decoupling capacitors need to be accommodated due to smaller packages.

Platform Challenges

In the sub and mini notebooks, one does not have the luxury of available real estate as the form factors are continually being reduced.

- The need to support additional features to improve performance yet still reduce the BOM cost and real estate.
- Signal integrity challenges in terms of timing margins and signal quality for higher and higher DDR/FSB frequencies need to be overcome. These impose stringent motherboard routing guidelines.
- Due to smaller package sizes and reduced pin pitches, the motherboard breakout becomes even more difficult.

SUMMARY

The Intel 915GMS chipset, which was the first of the smaller form factor chipsets, launched by Intel was very well received by customers, and this motivated Intel to come up with a better next-generation SFF chipset, the Intel 945GMS Express Chipset. Mobile Intel 945GMS

comes in 27 mm x 27 mm package size with additional features (such as additional LVDS channel), higher FSB and DDR speed, optimized motherboard power delivery, platform power management, and innovations in chipset power management. All the above features make this platform ideal for mini and subnote segments of notebook PCs. With this platform Intel has delivered a best in-class SFF platform. Intel's customers will benefit significantly from Intel's strong focus on this market segment.

ACKNOWLEDGMENTS

The authors thank A. Khandekar and S. Pal for providing the power measurements used in this paper. We acknowledge all the teams who worked on this platform: Product Development, Package Development, Power Delivery and Signal Integrity, Boards Development, Technical Marketing, and all the validation teams.

REFERENCES

- [1] "Innovation Brings Low Power Integrated Graphics to the Intel® Centrino® Mobile Technology Platform," *Intel Technical Journal*, Vol. 7, Issue 2, May 2003.
- [2] "Intel 915GMS Chipset: in Mobile Platforms, Smaller is Better," *Intel Technology Journal* Vol. 09, Issue 01, Feb. 2005.

AUTHORS' BIOGRAPHIES

Eric C. Samson joined Intel in 1983 to design communication chipsets. He currently focuses on power management architecture for chipsets. His e-mail is eric.c.samson at intel.com.

Michael N. Derr joined Intel in 1992 and has held various positions in the Client Chipset Development group over the years. Most recently, he has been working as an architect on I/O features and platform power management. His e-mail is michael.n.derr at intel.com.

Sidharth Krishnama is a power delivery design engineer with Intel India, supporting mobile chipsets. He has a B.S.E.E. degree from India and an M.S.E.E degree from the University of Oklahoma, Norman. He started his career with Intel in 1999. His e-mail is sidharth.krishnama at intel.com.

Samir Gundawar is a hardware design engineer in the Mobile Platform Architecture and Development team. He has an M-Tech degree from IIT-Delhi. He started working at Intel India in 2003 after working in Wipro Technologies for four years. His e-mail is samir.v.gundawar at intel.com.

Suresh V. Subramanyam is the package design manager with Intel, India, supporting mobile chipsets. He has a B.S.M.E. degree from India and an M.S. degree in

Materials Science & Engineering from Singapore. He joined Intel in 2003 after spending six years with HP/Agilent, Singapore. He started his career with Bharat Electronics, Bangalore, India, in 1988. His e-mail is suresh.v.subramanyam at intel.com.

Deepa Mohan joined Intel's Mobile Platforms group in 2003. She is a hardware design engineer in the Mobile Platform Architecture and Development team, designing Customer Reference boards for next-generation Intel[®] processors and chipsets. She has also worked on memory signal integrity simulations and technical marketing. Deepa received her B.E. degree in Electronics & Communication from R.V. College of Engineering, Bangalore, India in 2003. Her e-mail is deepa.mohan at intel.com.

Satish Prathaban is a hardware design engineer in the Mobile Platform Architecture and Development team. He has an M-Tech degree from IISc-Bangalore. He started working at Intel India in 2003 in the area of motherboard and package power delivery. His e-mail is satish.prathaban at intel.com.

Vikas Shilimkar is a hardware design engineer with Intel, India. He joined Intel in 2004 and works in the Mobile Platform Architecture and Development group. He works on customer reference motherboard designs for Intel's next-generation platforms. His specializations include platform and package power delivery. He holds a B.E. degree in Electronics and Telecommunication from COE Pune, India. His e-mail is vikas.s.shilimkar at intel.com.

Taninder S. Sijher received his M.S. degree from the University of Mississippi in 2004 and his B. Tech. degree from the Indian Institute of Technology, Bombay in 2001. He has been with Intel since 2004 working as a package electrical engineer with Mobility Group India and is responsible for the electronics of the next-generation Mobile GMCH packages. His e-mail is taninder.s.sijher at intel.com.

Parthasarathy Ramaswamy obtained his Ph.D. degree from IISc., Bangalore, India and is working with Intel Technology India Pvt. Ltd., Bangalore as an SIE manager. He is responsible for signal integrity and power-delivery solutions for the chipsets group. His areas of interests include high-speed PKG/Board designs, power and signal integrity studies, etc. His e-mail is ramaswamy.parthasarathy at intel.com.

* Other names and brands may be claimed as the property of others.

This publication was downloaded from <http://developer.intel.com/>.

Legal notices at <http://www.intel.com/sites/corporate/tradmarx.htm>.

Copyright © Intel Corporation 2006. All rights reserved. Intel, Centrino, Core, and Intel SpeedStep are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

THIS PAGE INTENTIONALLY LEFT BLANK

For further information visit:

developer.intel.com/technology/itj/index.htm