Volume 12 Issue 01 Published, February 21, 2008 ISSN 1535-864X DOI: 10.1535/itj.1201

Intel® Technology Journal

Technology with the Environment in Mind

Intel Technology Journal (February 2008) on Intel's Technologies and the Environment discusses a wide range of green technologies and efforts to conserve and sustain our natural environment through greener products, buildings, lifestyle, manufacturing and operations.

Inside you'll find the following articles:

More information, including current and past issues of Intel Technology Journal, can be found at:
http://developer.intel.com/technology/itj/index.htm

Volume 12 Issue 01 Published, February 21, 2008 ISSN 1535-864X DOI: 10.1535/itj.1201

Intel® Technology Journal

Technology with the Environment in Mind

Articles

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

By Lin Chao
Editor and Publisher, *Intel Technology Journal*

Intel Corporation's four-decade history of technology innovations is a well-known fact. Less known but equally important is that for decades Intel has had goals to reduce the environmental impact of its manufacturing, operations, and products. This issue of the *Intel Technology Journal* on Eco-Smart Technologies (Volume 12, Issue 1, 2008) features a podcast by Ted Reichelt, Intel's Principal Environmental Engineer, as he remembers operational meetings with Gordon Moore and Andy Grove held nearly two decades ago on this very subject. Today, Intel's commitment to *design for environment* remains strong: our goal is to reduce the environmental impact of our operations while continuing to meet high-performance requirements for computing. We are designing for energy-efficient performance with platforms that can support usages and applications with dramatically decreased energy consumption, in some cases. We have lead-free and halogen-free materials and packages. We have been reducing our emissions of perfluorocompounds (PFCs) for some time now. We are improving the energy efficiency of our factories and operations. All in all, within the company itself, and with outside partners, we are making environmental sustainability a priority, thereby minimizing the negative effects of our operations and products on the ecosystem.

The seven papers in this issue of the *Intel Technology Journal* delve into the challenges and their solutions in our factories, operations, and products. The first two papers look at green micro-electronic packaging and a more energy-efficient universal serial bus. The first paper, *Materials Technology for Environmentally Green Micro-electronic Packaging* looks at how Intel is working to make substrate materials for micro-electronic packages lead-free and halogen-free. This paper discusses the challenges overcome by Intel to deliver on both aspects of this green packaging technology.

The second paper, *Making USB a More Energy-Efficient Interconnect* looks at the power consumption of the Universal Serial Bus (USB) and its impact on mobile platforms. The USB's consumption of power is due to its architecture; devices are constantly polled. Although this creates a simple and low-cost device model, it is fundamentally inefficient—especially when the device is idle or has little data to transfer. We offer techniques that transform USB into a much more energy-efficient interconnect, primarily by optimizing the idle behavior of USB host controllers and devices.

The next two papers explore the technology behind Intel's first built green building, and look at how green consumerism affects the purchasing decisions and technology usage of environmentally conscious people. The third paper, Intel's *First Designed and Built Green Building,* reviews specifically how a building project in Haifa, Israel was initiated to address the need for a "smart" building; and reveals how the project turned out to be a driver for triggering broader innovations in both the corporation and a largely inexperienced (in the sense of green) local building industry.

The fourth paper, *Green Homeowners as Lead Adopters: Sustainable Living and Green Computing,* reviews how green consumerism affects the purchasing decisions and technology usage of people who prioritize environmental values in their lives and homes. We draw on our ethnographic study of 35 green households in the United States and look at the role technology plays in the lives of these people. How do their definitions of green consumerism affect their purchasing decisions and their relationship to technology in general? We answer these questions and conclude with a framework of strategies for green computing at the intersection of technology and sustainable living.

The last three papers review greener manufacturing and operations. The fifth paper, *Novel Wastewater Conserving Technology Meets Environmental and Business Challenges,* describes how semiconductor wastewater was treated to achieve a high level of wastewater effluent quality meeting the water reuse and conservation priorities of the local authorities. The sixth paper, *Dynamic Data Center Power Management: Trends, Issues, and Solutions,* reviews Intel environmental, health, and safety early screening and materials management programs that ensure supply-chain stability through the early identification and management of the environment-, health-, and safety-associated risks. Future semiconductor devices are critically dependent on the ability of stable and reliable materials to support new generations of semiconductor devices. In this paper, we provide an overview of what drives Intel's procedures, how the procedures are used, and we review a case study example.

In the final paper, *Evaluation Process for Semiconductor Fabrication Materials that are Better for the Environment,* we examine the challenges of increasing data center power consumption and higher energy costs in the face of ever-increasing computing needs. We identify requirements that server platforms must address to solve data center power problems. We offer a solution that includes a platform resident Policy Manager that monitors power and thermal sensors and enforces platform power and thermal policies.

Foreword

By Theodore (Ted) D. Reichelt

Principal Environmental Engineer, Global Environmental Group, Intel Corporation

Intel has a long history of technology innovation and leadership. For more than four decades, Intel and our industry have consistently fulfilled the vision and challenge of Moore's Law: to double the transistor density on integrated circuits about every two years. Intel co-founder Gordon Moore's remarkable prediction for technology advancement is complemented by his personal commitment to environmental and conservation causes for many years. Dr. Moore has helped instill a *design for environment* philosophy throughout Intel and the semiconductor industry that spans microarchitecture, process technology development, product development, manufacturing, and supporting operations.

Today, environmental concerns are at the forefront of our daily lives. Issues such as global warming, clean water access and supply, and the recycling and reuse of waste products and materials are routinely in the news and the subject of public policy deliberations. While there is lively debate about causes and cures for these issues, Intel has taken a responsible approach to continuously reduce our environmental impact while finding ways to help make our technology a part of environmental solutions. To be more environmentally sustainable, an enterprise must look for opportunities both in its products and its operations—manufacturing, facilities, and supporting activities.

High computing performance with minimum energy use is an important goal for our integrated circuits and other products. Lower energy consumption reduces computing costs, extends battery life, and reduces the heat produced by the device. The U.S. EPA estimates that the Energy Star efficiency standard has helped reduce annual national energy consumption by 170 billion kWh (37 million metric tons of carbon equivalent green house gas emissions), the equivalent of taking 25 million automobiles off the road. Despite all the progress to date, there are many remaining opportunities for further efficiencies through novel microarchitecture and design, software, and optimization of both components as well as the total electronic product.

Another important product environmental goal is the use of materials that are safe for the environment and for public health. Safe materials also help to make recycling and reuse of products easier and more financially viable. Global regulations already require that certain metals, some halogenated flame retardants, and a few other materials be phased out.

Through early involvement in the process technology development lifecycle, the carbon footprint of future manufacturing processes can be more effectively minimized. Establishing goals and business processes for energy and water conservation, waste reduction, and the use of environmentally friendly materials ensures the proliferation of more efficient manufacturing processes.

Design and construction of greener buildings is also expanding rapidly. Standards such as the U.S. Green Building Council's LEED (Leadership in Energy and Environmental Design) are aimed at reducing resource use, identifying more sustainable building sites, and creating a better indoor working environment. Intel is working with a coalition of companies through the International Semiconductor Manufacturing Initiative (ISMI) to develop LEED standards specifically for high-tech manufacturing.

Here are some recent highlights from Intel's record of recent environmental accomplishments resulting from our environmental philosophy:

- We have been named to the Dow Jones Sustainability Index for the ninth straight year and have been the Technology Super Sector leader in the index for the seventh straight year.
- We have reduced our global warming emissions to the equivalent of taking 50,000 automobiles off the road.
- Our innovative 45 nm technology is lead-free and has significantly reduced halogen content; moreover, our newest microprocessors are up to 40% more energy efficient than their previous technology generation counterparts.
- In the last three years our operations have conserved over 9 billion gallons of water.
- We recycle 70% of our wastes.
- We are a major purchaser of renewable energy and are working to certify several buildings to the LEED green standard.

This Q1/08 edition of the *Intel Technology Journal* covers a wide application of technologies that enable high-performance computing while providing solutions to environmental concerns. I hope you will find this edition both informative and encouraging.

Technical Reviewers for Q1 2008 ITJ

Mostafa, Aghazadeh, Technology and Manufacturing Group
Jim Blakley, Digital Enterprise Group
Eric Brewer, Corporate Technology Group
Mike Calyer, Mobility Group
Derek Collier, Digital Enterprise Group
Tom Cooper, Technology and Manufacturing Group
Herman D'Hooge, Corporate Technology Group
Michael Goldstein, Technology and Manufacturing Group
Neil Gordon, Technology and Manufacturing Group
John Harland, Technology and Manufacturing Group
John Howard, Corporate Technology Group
Steve Kinsler, Technology and Manufacturing Group
Ravi Mahajan, Technology and Manufacturing Group
Bill McAuliffe, Mobility Group
Jay Melican, Digital Home Group
John Morgan, Technology and Manufacturing Group
Ted Reicheld, Technology and Manufacturing Group
Amir Roggel, Technology and Manufacturing Group
Bob Sankman, Technology and Manufacturing Group
Alan Schijf, Technology and Manufacturing Group
Karthi Vadivelu, Mobility Group

THIS PAGE INTENTIONALLY LEFT BLANK

Materials Technology for Environmentally Green Micro-electronic Packaging

Mukul P. Renavikar, Technology and Manufacturing Group, Intel Corporation
Neha Patel, Technology and Manufacturing Group, Intel Corporation
Ashay Dani, Technology and Manufacturing Group, Intel Corporation
Vijay Wakharkar, Technology and Manufacturing Group, Intel Corporation
George Arrigotti, Technology and Manufacturing Group, Intel Corporation
Vasu Vasudevan, Technology and Manufacturing Group, Intel Corporation
Omar Bchir, Technology and Manufacturing Group, Intel Corporation
Amruthavalli P. Alur, Technology and Manufacturing Group, Intel Corporation
Charan K. Gurumurthy, Technology and Manufacturing Group, Intel Corporation
Roger W. Stage, Technology and Manufacturing Group, Intel Corporation

Index words: Pb-free, halogen-free, solder materials, thermal interface materials, package and board reliability, BGA, substrates, micro-electronic packaging, EU RoHS, C4

ABSTRACT

Intel has been continuously striving to provide environmentally green micro-electronic packaging solutions for high-density interconnect (HDI) product applications. The environmentally green initiative consisted of providing lead-free (Pb-free) packaging materials solutions as well the enabling of halogen-free (HF) substrates technology to eliminate the use of brominated flame retardants. This paper discusses the challenges overcome by Intel to deliver on both aspects of environmentally green packaging. Although Intel's efforts to enable Pb-free and HF-compliant packaging solutions have been wide-ranging, the scope of this paper is limited to discussing the key technology development challenges faced in transitioning to Pb-free materials in first-level interconnects (FLI), second-level interconnects (2LI), solder thermal interface materials (STIM) applications, and halogen-free (HF) substrate materials. The transition to Pb-free micro-electronic packaging materials and HF substrate technology required a paradigm shift in the industry, needing extensive benchmarking initiatives and sharing cross-technology learnings across the industry and academia. The delivery of Pb-free packaging solutions across FLI, 2LI, and STIM applications as well as HF substrate technology has strongly reinforced Intel's One Generation Ahead (OGA) philosophy in micro-electronic packaging.

INTRODUCTION

Intel's drive to "get the lead out of the package" began over five years ago when we produced a Pb-free tin-silver-copper (SAC) solder for 2LI applications that complied with European Union Restriction of Hazardous Substances (EU RoHS) requirements. Continuing on this path to deliver "Pb-free" packaging, Intel recently reached a critical milestone by eliminating Lead (Pb) from the FLI solders in its next-generation 45nm Silicon technology roadmap products. Intel is among the first semiconductor companies to deliver Pb-free FLI solutions in high-volume manufacturing. In order to meet the stringent integration challenges of transitioning to Pb-free-compliant packages, Intel has also successfully developed substrate, FLI flux, and underfill (UF) materials technologies that are compliant with higher Pb-free processing temperatures. Intel has been working with suppliers, customers, and several industry consortia to develop and provide EU RoHS-compliant products. Intel has completed certification of EU RoHS-compliant materials and processes and is manufacturing and shipping many EU RoHS-compliant products today. Additionally, in anticipation of the RoHS regulations, Intel pro-actively worked to develop pioneering Pb-free STIM materials to

meet the challenging needs of heat dissipation from the silicon die.

Intel's drive to enable halogen-free (HF) substrate technology entailed a careful evaluation of HF material properties to identify robust materials sets in order to meet manufacturing, assembly, performance, and use condition-based reliability criteria. The selected HF substrate materials sets were fungible with existing manufacturing and assembly processes used previously with nHF (not halogen-free) cores. Several technical challenges were overcome in enabling the HF substrate technology. Mechanical drilling of HF cores, as well as substrate warpage, was evaluated to verify that their performance was on a par with nHF cores, from a manufacturing and assembly perspective. The electrical properties of the HF core material were also evaluated and it was determined that the impact on performance compared to that of nHF cores was negligible. From a reliability perspective, the key concern with HF substrates was delamination, which can occur due to moisture release at Pb-free reflow temperatures from the HF core material. To verify that adequate reflow delamination margins exist for HF substrates, relative to use condition requirements, a component reflow accelerated test was developed by Intel and used to assess Intel's HF product lineup. The drive to enable HF materials has continued with the development and successful introduction of HF-compliant packaging materials such as molding compounds, underfill materials, and Polymer TIMs.

In the first half of this paper, we discuss Intel's qualification of Pb-free packaging solutions in the first-level interconnect (FLI), second-level interconnect (2LI) and solder thermal interface materials (STIM) applications. In the second half, we address the enabling of halogen-free (HF) substrates and accompanying reliability challenges.

A schematic of a lidded ball grid array (BGA) Intel package is shown in Figure 1 which depicts the FLI, 2LI, TIM, and substrate materials technologies that are undergoing the environmentally green transition in Intel's micro-electronic packaging.

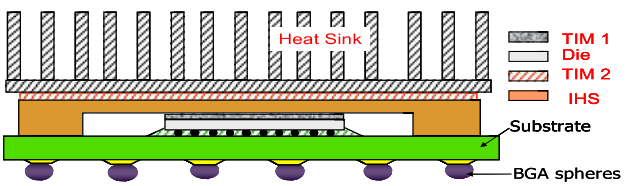

Figure 1: Schematic of Intel's lidded BGA package

PB-FREE INITIATIVE: FIRST LEVEL INTERCONNECT MATERIALS (FLI)

Intel recently announced the achievement of a significant milestone in the quest to deliver RoHS-compliant Pb-free FLI solutions in micro-electronic packaging. The transition to Pb-free FLI interconnects that connect the silicon die to the substrate eliminated the last 5% of Pb remaining in the package. Traditionally, tin-lead solder alloy has been used for FLI chip-to-substrate and 2LI substrate-to-board attachment interconnect materials. The presence of lead in tin-based solder alloys, mostly with the composition of eutectic 63Sn-37Pb, lends the solder superior thermal and mechanical characteristics for microelectronic assembly and reliability. However, the inherent toxicity of lead has raised serious environmental and public health concerns. Developing lead-free alternative solder alloys for micro-electronic substrates is of paramount importance. Intel selected tin-silver-copper (SAC) solder metallurgy as the lead-free chip attachment material for its 45nm CPU products. Compared to their tin-lead counterpart, high tin content, lead-free C4 solders possess physical properties less desirable for assembly and reliability: higher surface tension, increased mechanical stiffness, and a higher melting point. A number of technical challenges have been encountered and solved during Intel's lead-free C4 interconnect development. For example, reduced wettability of lead-free solder with die copper bumps can pose challenges to the downstream underfill process. Moreover, optimization of substrate solder metallurgy has also shown to be very effective in improving the mechanical robustness of the C4 interconnect, and in minimizing the occurrence of C4 brittle solder joints. The change to Pb-free SAC solder alloy necessitated the development of alternate flux materials to clean off the more tenacious tin oxides from the solder surface and form a robust FLI solder joint. The new flux material needed to be stable at high process temperatures as well as be cleanable following the chip attach process to allow strong adhesion between the underfill, the bump metallurgy, and the die passivation. The formation of a robust FLI, Pb-free joint significantly increased the current carrying capability of the joints. The transition to Pb-free FLI solder materials also necessitated the development of underfill materials technology designed to mitigate additional thermo-mechanical stresses imposed on the die due to stiffer FLI joints.

Changing from a Pb-based C4 substrate solder to a Pb-free metallurgy drove an increase in the peak reflow temperature during die attach in assembly requiring Pb-free-reflow-compliant substrate materials technology. The selection of the dielectric materials set (core material, buildup layer, and solder resist) in the Pb-free substrate was therefore critical to ensure robust reliability

performance of the package at higher reflow temperatures. In particular, the glass transition temperature (Tg) and the coefficient of thermal expansion (CTE) values of the dielectric materials set were carefully selected to minimize the risk of substrate warpage, substrate dielectric material cracking, and die damage.

Thus, along with changing the solder alloy material to one that is Pb-free, a judicious choice of the associated materials sets allowed Intel to solve the challenge of removing the remaining 5% of lead, thereby achieving EU RoHS-compliant FLI packaging technology.

PB-FREE INITIATIVE: SECOND LEVEL INTERCONNECT MATERIALS (2LI)

Another critical challenge to meet EU RoHS requirements in Intel's packaging technology was the transition to Pb-free solder technology for 2LI applications. 2LI refers to the interconnect between the substrate and the printed wiring board (PWB). 2LI is accomplished with solder sphere, flux and/or paste, and it involves two reflow processes: ball attachment (BA) and surface-mount reflows for board attach. In the BA process, paste is printed onto the metal pad on the BGA side of the substrate, typically by using screen printing. Then solder spheres are picked and placed onto the paste printed pads. Finally, substrates with solder balls undergo a reflow process, typically in a multizone convective oven. In the surface process, solder paste is applied onto the metal pad on the PWB, typically by using stencil printing. Solder ball-attached packages are then picked and placed onto the fluxed PWB. Finally, the entire package and PWB undergo a reflow process typically in a multizone convective oven. Traditionally, eutectic tin-lead alloy was used for 2LI solder metallurgy applications. Relatively low melting temperatures (T_m = 183°C) and excellent shock resistance of the eutectic Sn-Pb alloy made this alloy highly suitable for Pb-ed BGA applications. Several Pb-free solder alloys for BGA applications were evaluated for this purpose, and SAC405 (tin-4% Ag- 0.5% Cu) was downselected based upon extensive materials characterization and reliability evaluations. The SAC405 solder alloy has a higher melting temperature (217-221°C) than eutectic SnPb solders, as well as a higher elastic stiffness and yield strength. These differences in the physical and mechanical properties of SAC405 solder alloy posed several challenges to packaging processes and reliability performance, especially due to the need to reflow the solder alloys at much higher temperatures (230-260°C peak reflow temperatures) than those used for eutectic SnPb. Figure 2 shows a typical SAC solder alloy reflow profile (BA and SMT) used for SnAgCu solder.

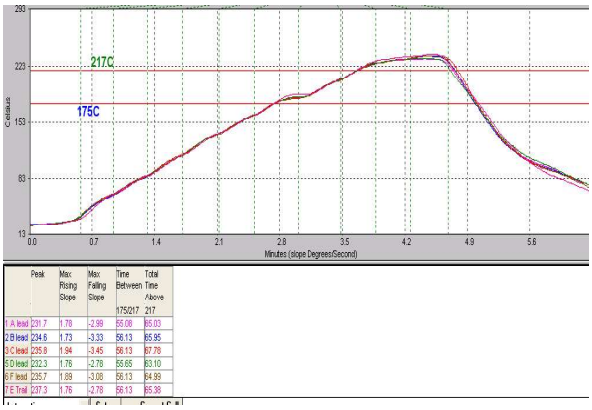

Figure 2: Typical SAC solder alloy reflow profile (BA and BGA)

The transition to Pb-free solders also necessitated a change in the composition of the paste from eutectic SnPb to SAC405 metal powder paste for BA and surface-mount applications. This change in paste composition forced the need to optimize the fluxability of the flux in the paste to ensure robust solder joint formation in conjunction with the SAC405 solder BGA spheres. Higher processing temperatures necessitated the use of appropriate solder paste/flux formulations that could withstand the higher thermal exposure and facilitate excellent joint formation with a myriad of surface finishes such as Electroless Nickel Immersion Gold (ENIG), Cu OSP (Organic Surface Protection), and Immersion Silver (ImAg). The use of SAC solder alloy for BGA applications has the potential to impact both the package- and board-level manufacturing and reliability. Additionally, Intel's packages also utilize paste for attaching BGA spheres to the packages. As mentioned earlier, the need to reflow SAC solders at higher temperatures impacts the package and board materials manufacturing to enable high temperature reflow. In addition to the metallurgical challenges involved in the high temperature reflow of 2LI solder joints, the impact of high temperature reflow on the viscoelastic behavior of the package and board-level materials and accompanying reliability concerns needed to be understood. In the next section, we discuss the 2LI solder joint reliability challenges and board-level surface-mount challenges due to the transition to Pb-free materials.

Pb-free Initiative: 2LI Reliability Challenges

A failure-mechanism-based approach was used for Lead-free (LF) 2LI reliability assessment. An LF material property comparison was made with SnPb solder to understand the failure-mechanism-based reliability risks. The typical LF reliability concerns based on the failure

mechanism and stress testing approach are listed in Table 1.

Table 1: LF failure concerns and recommended stress testing

LF Failures Concern	Stress Test	Comments
Solder fatigue	Temp Cycle	Electrical open/solder crack
Overstressing	Shock Test	Electrical open/solder crack
Overstressing	Vib Test	Electrical open/solder crack
PCB trace, via corrosion	Temp/Humid 85/85°C	Electrical open due to via, trace corrosion
IMC growth, Diffusion and Solder creep	Bake Test	Electrical open IMC growth, diffusion, & shorts due to solder creep

LF 2LI reliability challenges include delamination with increased reflow temperature (35° C higher than Sn/Pb), lower mechanical margin in high strain rate shock testing due to increased stiffness (1.5 times stiffer than SnPb solder), creep performance difference compared to SnPb, and Sn whiskers. The JEDEC specification for max reflow spec of 260° C was established as part of the solution, and Intel components were qualified to meet the JEDEC specification for peak reflow requirements. Use-condition-based reliability performance was used for Pb-free alloy selection. The LF mechanical margin was lower than that for SnPb solder, but it passed intensive board-level shock and vibe testing. The mechanical margin testing showed lower mechanical performance compared to SnPb solder, and the results are shown in Figure 3. The NCTF design rule was implemented at corner joints for shock margins.

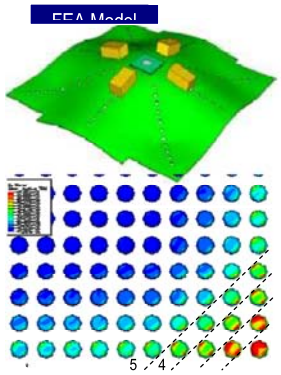

Solder	Crack (%)Row				
	Non CTF		CTF		
	1	2	3	4	5
Sn/Pb	50	0	0	0	0
LF	100	20	0	0	0
Sn/Pb	50	50	0	0	0
LF	100	100	40	0	0
Sn/Pb	100	75	0	0	0
LF	100	100	100	100	100

Figure 3: Mechanical margin assessment for LF solder

The failure mode during shock is manifested in cracking along the solder joint (either on the package or board side) as shown in Figure 4 [1].

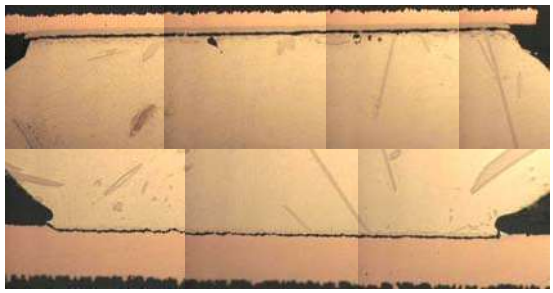

Figure 4: SAC405 solder joint failure in shock conditions

A detailed failure analysis reveals that the shock crack in SAC is along the interface between the IMC and the solder (package side) and through the IMC bulk (board side). In either case, shock failure is characterized by a lack of solder deformation and an absence of solder bulk cracking. This is quite contrary to thermal fatigue failure where the solder joint exhibits extensive inelastic deformation that is often time-dependent.

This difference is partly due to the strain-rate sensitivity of metallic materials. Metallic materials including solders typically become stronger with increasing strain rates. In other words, the flow stress increases with increasing strain rates. The strain rate sensitivity is a strong function of the homologous temperature (T_{hom}), which is considerably higher for solders, due to their low melting temperatures.

As a result of the high strain rate sensitivity, the yield strength of SAC solders increases rapidly with strain rate. This increased yield strength suppresses any plastic deformation and prevents the shock energy from dissipating through the solder joint, thereby transferring more stress to the interface which causes interfacial fails. The yield strength of eutectic SnPb solder is relatively low compared to the SAC405 solder alloy. This means SnPb solders can dissipate more high strain rate energy through deformation and hence can perform better in shock than a SAC405 solder alloy. The higher yield strength of the SAC405 solder alloy is derived primarily from the precipitation hardening of the tin matrix by the Ag_3Sn precipitates/platelets. In addition to the increased bulk strength of the SAC alloy, the higher reflow temperatures can also cause an increased thickness in the IMC layers and thereby degrade the shock performance of the SAC405 solder alloy. However, the increased strength of the SAC405 solder alloy is beneficial for thermal cycle fatigue resistance as it reduces creep damage in each

thermal cycle. Thus the transition from eutectic SnPb to SAC405 solder alloy poses more challenges in high strain rate shock applications, but provides more margin in thermal cycle reliability.

The LF fatigue performance was about 20-30% higher than SnPb for Flip Chip Ball Grid Array 9FCBGA (FCBGAs). Temperature cycle results are compared with the SnPb solder as shown in Figure 5. LF (SAC405) showed improved fatigue performance in both 15-min. and 30-min. dwell time testing.

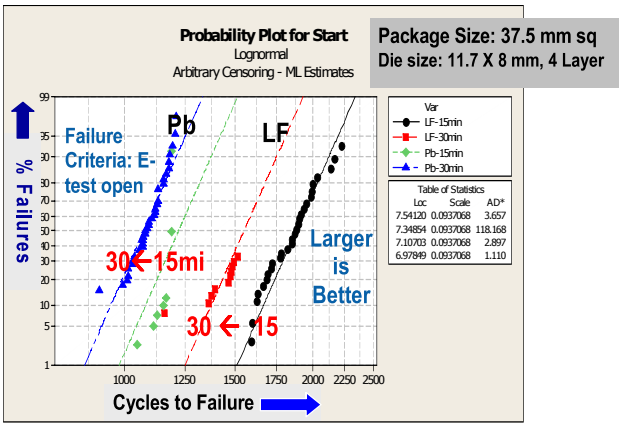

Figure 5: LF temp cycle performance comparison with Sn/Pb solder

LF creep performance concerns were addressed through long dwell time testing. Long dwell time temperature cycle testing (three cycles/day) were completed for SAC405, and the results showed better LF performance than SnPb solder as shown in Figure 6 [2].

Figure 6: LF Temp cycle performance in long dwell time testing

The industry concern over long dwell time was satisfactorily addressed based on these results, and the LF reliability model correlation showed a similar type of fit compared to that of Sn/Pb solder. The corrosion and diffusion concerns were addressed through temperature humidity testing (85/85 test) and bake test (125°C) for 1000 hrs. Bake testing did not show interface-related failure mechanisms. The results are shown in Figure 7.

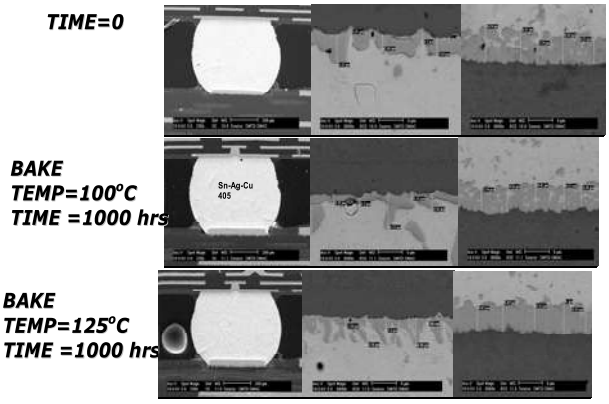

Figure 7: Bake test results for SAC405 with Im/Ag SF

The other LF reliability concern is the formation of tin (Sn) whiskers. Sn whisker formation is not an issue for Pb-free solders, but concerns the Sn surface finish on components. The Sn whisker failure mechanism is an electrical short caused due to the growth of the whisker. An example of Sn whiskers is shown in Figure 8 [3].

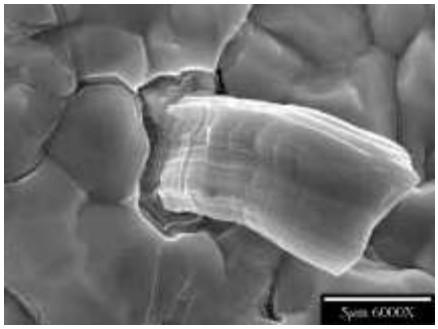

Figure 8: Sn whiskers

The concern about Sn whiskers for Sn-coated components is not related to LF issues, but it is the same as that for Sn/Pb solder. JSTD 201 was established for Sn whisker mitigation that includes the use of matte Sn and anneal at higher temperatures for stress relief prior to SMT.

PCB surface finish quality characteristics have an impact on reliability: micro void associated with Im/Ag PCB SF quality characteristics negatively impacted the solder fatigue margin (40-50% reduction in temperature cycle performance). The impact of micro void on temperature cycle performance is shown in Figure 9. The other concern is a Kirkendall-type void in bake testing for OSP

SF. Kirkendall-type voids were attributed to Cu purity and OSP chemistry and are related to the coating process. Proper control of PCB plating quality mitigated the micro void risk for Im/Ag and Kirkendall-type void for OSP SF.

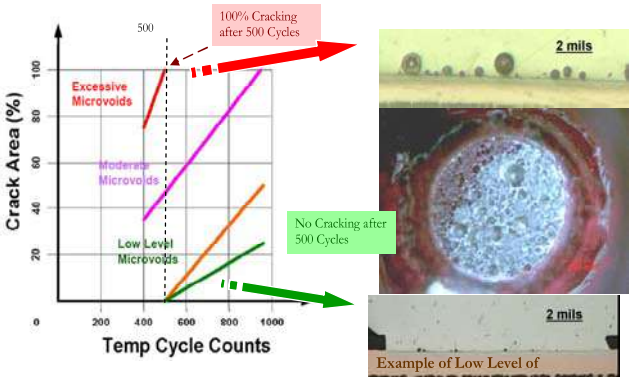

Figure 9: Impact of micro void on temperature cycle performance

Pb-free Initiative: Surface-Mount Challenges

The transition to Pb-free solders for 2LI posed severe challenges to surface-mount technology, and an equivalent effort in adapting board manufacturing and SMT processes was required. The transition affected not just the SMT pastes used in board assembly, but also wave solders and board materials. Process steps significantly affected by the conversion are shown shaded in pink, with their impact described in Figure 10.

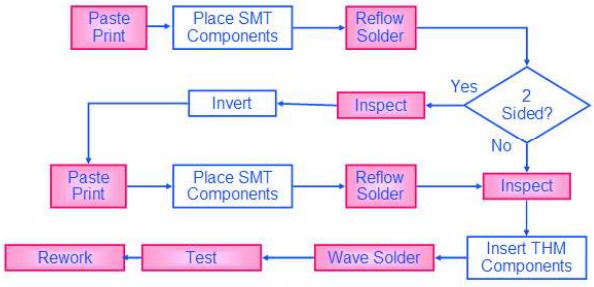

Figure 10: SMT process flow diagram indicating impact due to Pb-free transition

Solder Paste Printing

The same stencil printing equipment, stencil fabrication technology, and stencil thickness used for SnPb can be used for Pb-free solder pastes. If reduced stencil aperture openings (when compared to the land area) are being used for SnPb pastes, these apertures will need to be expanded to cover the entire land area, since Pb-free solders do not wet out to the same extent as Pb-ed. The most popular Pb-free solder paste composition is SAC305.

Component Placement

Component placement is only marginally affected by the change to Pb-free solder. SAC solders are more grainy and less shiny than SnPb solders. There can also be a wider part-to-part variation in solder appearance. When front-side lighting is used for component lead or ball recognition, such as for BGAs, these differences may require adjustments to the vision system hardware and/or algorithms, to properly recognize Pb-free solder balls and surfaces.

Reflow Soldering

Since SAC305 has a higher initial melting point (217°C) than Sn-Pb (183°C), higher reflow soldering temperatures are required as shown in Figure 11. Wetting characteristics of SAC305 are not as superior as that of SnPb. This requires a slightly longer reflow soldering window, typically 230 to 250°C for Peak Reflow Temperature with 40 to 120 seconds above 217°C. Since this window is narrower than that used for eutectic SnPb, a greater level of control is required for Pb-free reflow soldering, and this necessitates better reflow soldering ovens. Ten zone ovens are typical for Pb-free reflow, instead of the seven zone ones sufficient for SnPb. Higher temperatures also lead to deleterious effects on the boards. Metallic surfaces not covered with solder paste during reflow soldering will get more oxidized than for SnPb, degrading the wetting of these surfaces during subsequent soldering processes, such as wave soldering. Using a nitrogen atmosphere in the reflow oven will mitigate this oxidation to a large extent. Higher reflow temperatures also increase PCB warpage, increase risk of delamination and blistering, and weaken the plated through-hole (PTH) copper. PCB laminates with higher glass transition temperatures, higher thermal degradation temperatures, longer time to delamination at temperatures at and above 260°C, as well as lower z axis expansion coefficients, will all help alleviate these issues.

Wave Soldering

SAC305 and eutectic SnCu are two alloys used for Pb-free wave soldering. Both have higher melting points (217°C and 227°C, respectively) than SnPb. Thus, an increased solder pot temperature is required, in the 260-275°C range, rather than the 250-260°C range. Increased temperatures may require a different wave solder flux, one that reaches an optimum activity level at a higher temperature. A new wave solder pot and machine are required to avoid cross-contamination between Pb-free and SnPb solders.

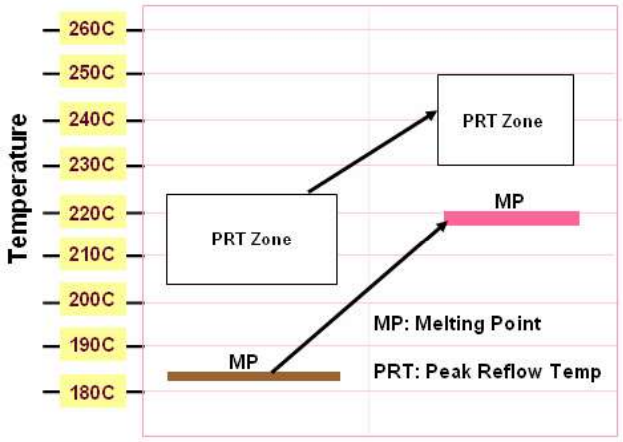

Figure 11: Comparison of peak reflow temperature ranges between eutectic SnPb (L) and SAC (R) alloys

Since the Pb-free solders have diminished wettability under the same flux activity levels, filling PTHs to achieve strong reliable through-hole solder joints and adequate coverage of test point pads becomes a challenge, especially when using surface finishes such as OSP, which is prone to oxidation at higher temperatures. Examples of acceptable and unacceptable wave solder filling are shown in Figure 12. Some ways to mitigate these challenges are increasing flux application density, achieving adequate penetration of the flux in the through holes, and using a turbulent chip wave option. The latter, however, increases the dissolution and erosion of the copper lands and traces, and it also increases copper content in the solder pot. Since the melting points of Pb-free alloys increase quite markedly with an increase in the copper content, an effect that also reduces the fluidity of the solder wave, the copper content needs constant monitoring to avoid certain defects.

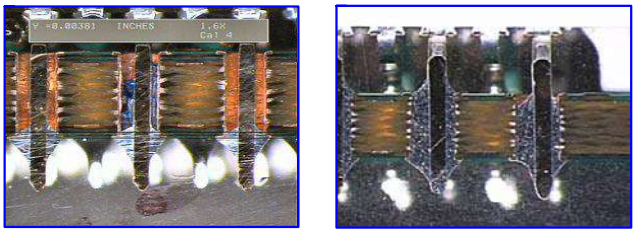

Figure 12: Unacceptable wave solder hole fill (L) and acceptable hole fill (R)

In-Circuit Test (ICT)

Higher reflow and wave soldering temperatures cause two issues for In-Circuit Test (ICT). One, board test points collect extensive polymerized flux residue. This results in a larger build-up of this flux residue on test probe tips and barrels. Consequently, the test probes have to be cleaned more frequently in order to control retest rate increases caused by probe contamination. Two, there is increased oxidation of test points on the board. This results in higher contact resistance, which increases false failures, raising the retest rate and reducing the capacity of the test area. Probes with higher forces, sharper points, or rotating action can reduce the contact resistance issues with Pb-free boards, but can also cause board damage due to board flex and damage at test points. Using a nitrogen atmosphere during reflow and/or wave soldering can significantly reduce surface oxidation of the test points, but this will increase process costs. Also, because Pb-free solders are harder than SnPb, probe tips wear out faster, requiring more frequent replacement.

Solder Joint Inspection

SAC solder joints are typically less shiny and more grainy than eutectic SnPb joints as shown in Figure 13. SAC solder joints also spread less on lands and pads than SnPb joints. For these reasons, both manual and automated visual inspection criteria need to be adjusted for Pb-free board assemblies. Due to the lack of Pb, a heavy element, SAC solder has lower stopping power for x-rays. This causes x-ray images to appear a lighter shade of gray in transmission X-ray images used for inspection of solder joints that are not visible, such as BGA and QFN joints. Hence, Automated X-ray Inspection (AXI) criteria may also require adjustment.

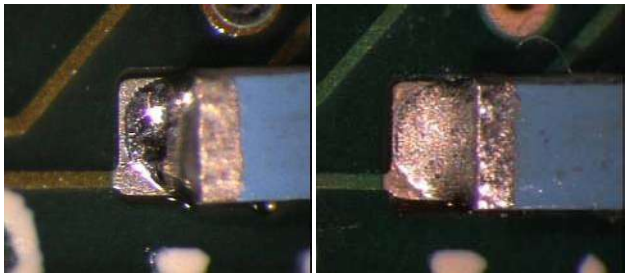

Figure 13: Visual inspection difference between SnPb solder joint (L) and SAC solder joint (R)

Rework

Repair and Rework become increasingly more difficult for Pb-free board assemblies, mainly due to the higher temperatures required for removing and replacing defective components. A greater amount of rework temperature control is required for Pb-free assemblies and this increases the total rework time in most cases. Due to the different solders, soldering irons and other equipment need to be changed from those used for SnPb rework. The risk of damage to the PCB lands, pads, and laminate is also increased, especially when there is direct contact between hotter soldering irons and the pads, such as during the Site Redressing process step for BGA component rework. Thinning of the copper thickness in the PTH barrel is significantly increased during the mini-pot rework process for connectors and other through-hole

components as shown in Figure 14. The higher tin content of the SAC solders causes more copper dissolution and erosion, increasing barrel thinning. Use of newer Pb-free alloys can alleviate this copper dissolution/erosion to a large extent.

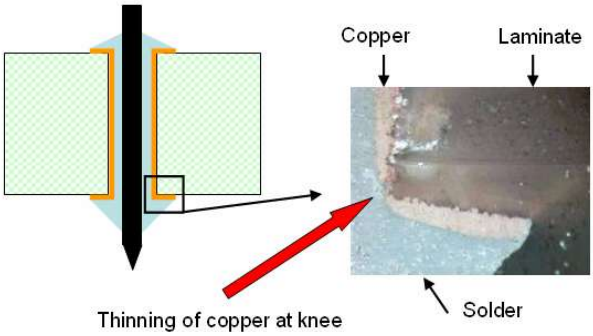

Figure 14: Copper erosion from minipot rework

Pb-free Initiative: Customer Manufacturing Enabling Challenges

Given the significant challenges posed in the SMT materials and processes to enable Pb-free transition, an equivalent effort was needed to influence and enable the ODM and OEM customer base to adopt the Pb-free technologies and processes. The efforts undertaken by Intel to influence the industry are outlined in this section. Customer enabling for Pb-free board assembly consisted of multiple steps over four years prior to the July 2006 RoHS date.

1. A customer-ready document, informally called the Pb-free Manufacturing Advantage Service (MAS), was prepared to capture technical learnings from the Intel Lead Free Board Assembly Team (LFBAT). It included tutorial information on all Pb-free board assembly modules in a typical production line, and also the Intel Reference Process for each module, with detailed process parameters and process material selections. It served as a starting point for customer process development. The Pb-free MAS, containing a substantial wealth of information, was developed in order to meet customer's requirements for more detailed information.
2. Customer Manufacturing Enabling (CME), a new group at the time, engaged major ODM and OEM customers, delivered the Pb-free MAS in person (primarily in APAC), and invited certain customers to participate in further enabling activities.
3. One challenge with developing a board-level process was the limited availability of Pb-free boards and components. The CME team developed Pb-free test boards, with Pb-free Bills of Material (BOMs), representing designs from desktop, mobile, and server market segments. Selected customers in each segment were provided with board designs, or physical boards, and BOMs, or physical parts, depending on their preference.
4. Customers used the board kits (typically 75 boards) to develop a board assembly process on their own. Both Intel and customers then performed reliability tests on samples of the boards and shared results.
5. After results from the development builds were incorporated into customer processes, customers completed another round of builds with Intel representatives present, using additional Pb-free board kits. These were called validation builds, intended to confirm that the customer process could consistently produce good boards over a larger quantity (up to 200) in a manufacturing environment (rather than a lab), without tweaking the process mid-build. Again, Intel and customers separately performed reliability testing, Failure Analysis (FA), and Materials Analysis (MA) and again, shared results.
6. Intel and customers held Manufacturing Readiness Assessment meetings to review all results and customer status regarding their own further development work and builds of prototypes.
7. After the launch of Pb-free products, Intel monitored customer manufacturing performance during launch and ramp, providing assistance as needed.

By this process, Intel ensured a smooth launch and ramp for the initial Pb-free platforms in each segment. In addition, it helped drive industry convergence on a narrow set of Pb-free materials that allowed Intel to produce Pb-free components with only one ball alloy for all desktop, mobile, and server products.

PB-FREE INITIATIVE: SOLDER THERMAL INTERFACE MATERIALS (STIM)

Intel's foray into Pb-free materials in electronic packaging began more than five years ago through the introduction of STIM in the 90nm technology node products. The relentless progress of Moore's law, leading to a doubling of transistor density in silicon chips every generation, drove the need to develop thermal solutions to dissipate additional heat generated in the silicon die. Consequently, Intel's packages have evolved from a bare die solution catering to mobile market segments to Integrated Heat Spreader (IHS) lidded products in desktop and server market segments as shown in Figure 15 [4].

Figure 15: Evolution of Intel's Package technology for meeting thermal performance requirements

There are several technical and cost drivers to enable lidded thermal architecture such as minimizing the impact of local hot spots by improving heat spreading, increasing the power-dissipation capability of the thermal solutions, expanding the thermal envelopes of systems, developing thermal solutions that meet business-related cost constraints, as well as developing solutions that fit within form-factor considerations of the chassis.

The primary role of the IHS is to spread the heat out evenly from the die and to provide a better bondline control of the interface material. This can be achieved by increasing the area of the IHS and by using a high thermal conductivity thermal interface material with low interfacial resistances. In order to meet thermal dissipation targets, Intel introduced polymer thermal interface materials (PTIM) initially with 3-4 W/m^{o}K bulk thermal conductivity and then successfully transitioned to Pb-free solder-based thermal interface material to meet the ever increasing demand for thermal cooling capability as shown in Figure 16 [5].

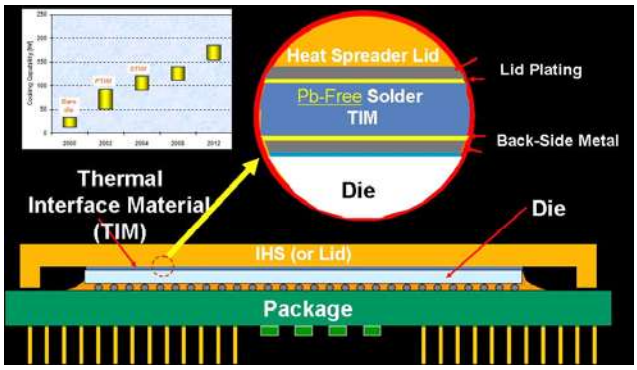

Figure 16: Improvement in thermal cooling capability with TIM materials (Polymer vs. Solder)

The introduction of Pb-free solder-based TIM materials posed significant integration challenges. The STIM needed to relieve the mechanical stress caused by CTE mismatch of the integrated heat spreader lid and the silicon die and to minimize stress transfer to the silicon die during thermal cycling [6]. The thermal conductivity and the mechanical compliance requirements resulted in the development and qualification of low melting temperatures (157^{o}C T_m), low mechanical yield strength (4-6 MPa), and relatively high thermal conductivity (~87 W/m^{o}K) pure Indium (In) metal for STIM applications. In order to use In for STIM applications, appropriate flux vehicles had to be developed to a) effectively reduce the thermodynamically stable native Indium oxide on In performs; b) to control solder joint voiding post joint formation; c) to control interfacial reactions with surface finishes on the IHS lid and the back side metallization (BSM) on the silicon die; and d) to deal with reliability issues faced in small and large die products, such as thermal fatigue cracking of the Indium during thermal cycling. The assembly process, including the reflow of the Indium STIM to form uniform intermetallic compounds (IMCs) post assembly, is a key challenge. A schematic of the STIM microstructural development as a function of packaging assembly steps is shown in Figure 17. The Indium oxide on the surface of the Indium needs to be effectively reduced in order to form uniform and defect-free intermetallic layers at both the die/Indium and the IHS lid plating (Ni/Au) and the In. Indium oxide is an extremely tenacious and thermodynamically stable oxide as shown in Figure 18 [7]. The presence of voiding in the joint can potentially lead to an increase in local thermal resistance and consequently lead to the degradation of the thermal performance of the joint. Additionally, excessive spallation of the binary Au-In IMCs as well as the formation of excessive Kirkendall voiding due to relatively different diffusion coefficients of In-Au and Ni can result in an increase in the thermal resistance of the joint.

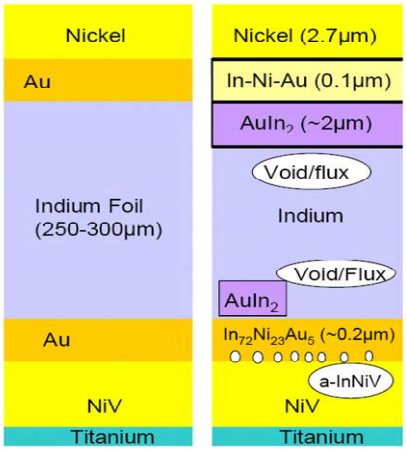

Figure 17: Use of Indium metal as STIM and interactions with surface finish on IHS and BSM pre and post assembly

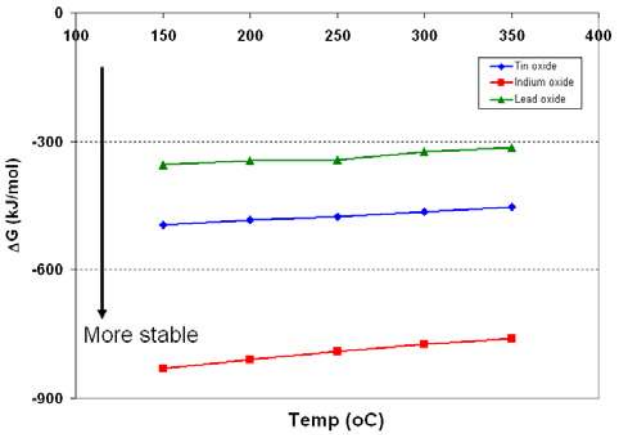

Figure 18: Thermodynamic stability of Indium Oxide as a function of temperature

In reliability testing as in thermal cycling, tensile and shear stresses are imposed on the STIM joint due to the mechanical coupling of the die to the IHS lid and the package as shown in Figure 19.

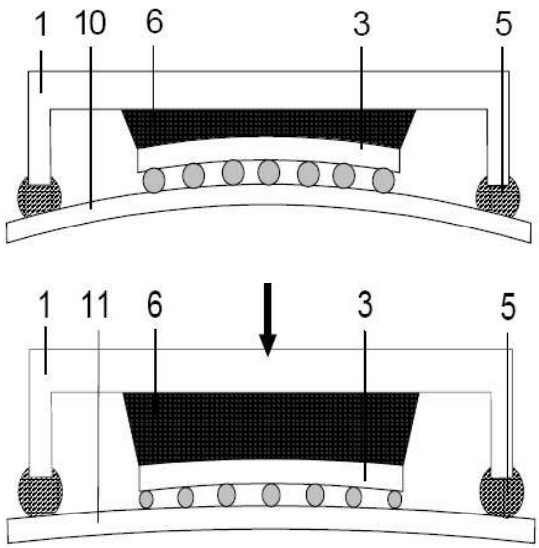

Figure 19: Warpage induced stresses on STIM joint at low temperature and high temp thermal cycling (Numbers in figure indicate different locations along package)

Typical failure modes encountered in STIM joints relate to thermal fatigue cracking of Indium close to the IHS/Indium interface which is manifested in the form of a white signature in CSAM imaging as shown in Figure 20 [6].

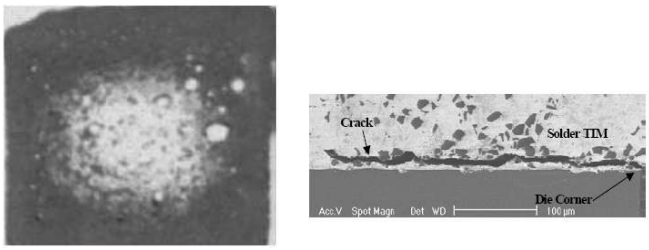

Figure 20: White CSAM image showing delamination at the Indium/IHS interface and corresponding cross-section SEM image showing cracking post reliability testing

The reliability performance of the STIM joints has been found to be modulated by the relative thickness and morphology of the binary and ternary IMCs as influenced by the fluxing ability of the flux used and the reflow profile used, as well as several mechanical design attributes of the IHS dimensions/die size, package stiffness, and preform dimensions. In addition to the technical challenges faced in enabling Indium, a significant effort was made to establish a strong supply chain for IHS lid manufacturing and plating, sealant materials technology for attaching the IHS lid to the substrate, as well as the development of appropriate back side metallization on the die to enable interfacial reactions with STIM to ensure robust joint formation.

In summary, Intel's transition to Pb-free packaging materials technology was attained through a judicious choice of materials across all functional areas such as FLI, 2LI, and STIM. The Pb-free materials solutions met all the integration assembly and surface-mount challenges as well as component and board-level reliability requirements. Intel worked closely with industry partners including suppliers and the ODM and OEM customer base to achieve a smooth launch and ramp of the Pb-free packaging materials technologies.

We now discuss Intel's stewardship in enabling HF-compliant packaging materials initiatives specifically the enabling of HF substrates technology.

HALOGEN-FREE PACKAGING MATERIALS INITIATIVE

As part of Intel's broad strategy to support an environmentally sustainable future, Intel is introducing environmentally conscious HF and Pb-free packaging at the 45nm CPU and 65nm chipset technology nodes. HF packaging materials introduced by Intel include several materials such as molding compounds, underfill materials, and substrates. The scope of this section of this paper is limited to HF-compliant substrate technology. Historically, components and printed circuit boards

(PCBs) have used nHF flame retardants, which have been the subject of an environmental impact debate for a number of years. Primary concerns regarding nHF flame retardants include bioaccumulation and toxic dioxin formation during recycling. Intel's drive to meet HF requirements in substrates is to substantially reduce the Br and Cl levels in the substrates to meet HF requirements, which are currently being established by industry consensus.

Table 2: Comparative core material properties

Core	Tg (°C) TMA	CTE (ppm/C)			Modulus (GPa)		Peel Strength (kN/m)	Dielectric Constant	Dissipation Factor	Moisture Absorption (%)
		x1,x2	y1,y2	z1,z2	25°C	260°C		@ 1 GHz	@ 1 GHz	
Halogenated	176	12,11	11,10	22,166	21	11	1.1	4.5	0.009	0.07
HF	172	13,10	12,11	15,130	23	10	1.0	4.7	0.018	0.05

A key component in Intel's HDI package that requires conversion to HF is the substrate, which is the focus of HF enabling in this paper. Typically, nHF substrates contain Br and Cl-based compounds in the core material, the buildup dielectric layers, the solder resist, and the PTH plug material. To enable HF substrates, each of the above mentioned materials was changed over time. A schematic cross-section of a package is shown in Figure 21.

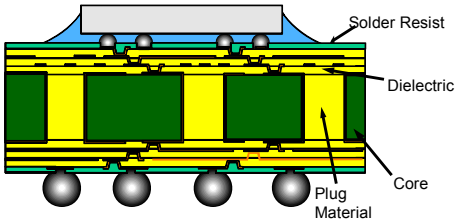

Figure 21: Assembled substrate schematic, labeled with key material items that required a change to HF

Intel has several years of HVM experience with HF dielectric, plug and solder resist materials, and hence the final push to enable HF substrates required a change to the core material, wherein the HF core does not contain brominated flame retardants. HF core material candidates, which meet Intel's assembly and reliability criteria, have been identified, and these are presented in the next sections.

HF Core Material Selection

Challenges

The ideal HF core material is one that can serve as a "drop-in" solution, with material and reliability properties that meet or exceed those for nHF core material. By close matching of material thermo-mechanical and electrical properties, the degree of change to substrate manufacturing, assembly, board-level reliability, and performance can be minimized. Table 2 shows a sample comparison of nHF vs. HF core material properties, where x1, x2 refers to CTE in the x direction below Tg (x1) and above Tg (x2). As indicated in Table 2, the thermo-mechanical and electrical properties of selected nHF and HF cores were similar. This enabled a relatively smooth conversion from the standpoint of assembly, 2LI reliability, and electrical performance.

The mechanism for flame retardency in nHF vs. HF core materials is different, due to differences in the flame retardant used in the core. nHF cores typically use a brominated flame retardant, wherein the Br reacts with combustion reactant species, suppressing reaction propagation and creating a layer of char, both of which help to stop the fire. In contrast, HF core materials typically use a metal hydrate as the flame retardant, wherein the metal hydrate releases water to cool the polymer and simultaneously creates a char passivation layer:

Metal hydrate + heat → char + water

In practice, use of HF core materials in Pb-free packages can be challenging, because the HF core tends to undergo more decomposition/water release at Pb-free reflow temperatures (~260°C). This is due to differences in the flame retardant type and content in HF vs. nHF cores. This poses a challenge specifically for BGA component reliability, as a significant amount of moisture release from the core during repeated Pb-free reflows (for BA, board mount, etc.) can facilitate delamination in the substrate as shown in Figure 22.

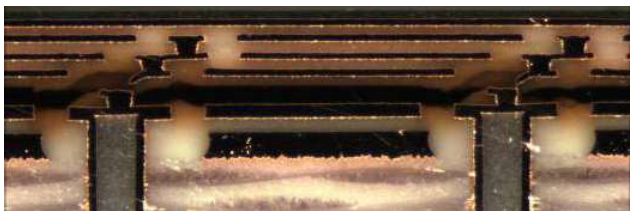

Figure 22: Cross-section of HF substrate, which has suffered a delamination in the buildup layers

To address this concern, a reflow accelerated test best known method (BKM) was implemented and used in the HF core material downselection process. The reflow accelerated test BKM incorporated JEDEC Pb-free Level 3 preconditioning (L3 precon) [8] with stepwise, additional Pb-free reflows (@ 260°C), to check for the delamination margin in the HF product substrates. Up to 15x additional reflows (beyond L3 precon) were run to test the delamination margin. Through careful material selection and screening, robust core materials, that met Intel's reliability requirements, were identified.

HF Core Material Results

HF core material selection required us to focus on substrate manufacturing, component and board-mount assembly, reliability testing, and performance. The results from these evaluations are presented in the following sections.

Substrate Manufacturing

From the substrate manufacturing perspective, the key challenges were selection of robust core materials, followed by mechanical drilling and flatness assessments of those materials. To select the most robust core materials, substrate suppliers were engaged and enabled with the reflow accelerated test BKM. Thorough evaluations were performed on various short and full loop test vehicles (TVs) with different designs, to understand the impact of core material and design on delamination reliability. Based on the number of reflows before the occurrence of delamination, the core materials with the most robust heat resistance were selected for further evaluation. These core materials were confirmed to be HF at Intel through ion chromatography testing, with the Br and Cl contents measuring <4 ppm. Core material drillability, as well as drill bit life parity between nHF and HF core, was established across the substrate supply base through drilling evaluations on the downselected materials. The parity in drill bit life ensures that there is no increase in drilling costs when an HF core is used. By measuring substrate flatness on incoming substrates, it was shown that HF core and nHF core units were equivalent across the substrate supply base.

Intel Assembly

The above trend carried forth through Intel assembly, wherein HF substrate flatness was the key assembly concern. Figure 23 shows comparative HF vs. nHF BGA ball coplanarity data for a 13x14mm TV. The data confirm that HF and nHF cores have similar flatness performance.

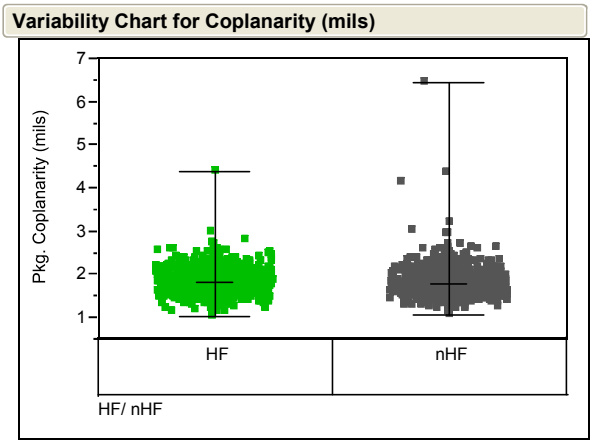

Figure 23: Comparison of HF to nHF package coplanarity for 13x14 mm TV

Board mount of nHF vs. HF components was also studied, with results indicating similar board-mount yields for both nHF and HF units.

HF Substrates Component Reliability

Reliability testing at the component level involved use of Intel's new Pb-free L3 preconditioning plus up to 15x additional reflows at 260°C, to check for delamination margins. A number of HF core materials were dropped from consideration because of delamination margins. The selected HF core materials were robust during reliability testing, and for the given form factors/designs in Table 3, passed more than 10x additional reflows beyond L3 preconditioning before any delamination was observed. Reliability results were similar for substrates across Intel's supply base, indicating sufficient reliability transparency.

HF Substrates Enabled Reliability

Enabled reliability testing (component mounted on the board) also showed performance parity between nHF and HF core TVs as shown in Table 4. In shock testing (a.k.a. dynamic bend testing), neither nHF nor HF substrates showed cracks in any critical to function (CTF) BGA solder joints.

Flammability Rating

Due to concern about flame retardant decomposition during multiple Pb-free reflows for BGA products, and the potential implications of this for flammability rating, a check of the UL-94 flammability rating before and after extended Pb-free reflows was done. Table 5 shows the results, which indicate that both nHF and HF core materials were able to maintain a V-0 flammability rating after 10x Pb-free reflows.

Table 3: Example HF core component reliability performance for different form factors

Substrate Form Factor	Reliability Stress	HF Core Related Fails
13x14 mm	L3 Precon + TCB (-55°C / +125°C)	0/149 post 750 cycles
	L3 Precon + HAST (130°C/85% RH)	0/59 post 100 hrs
	L3 Precon + add'l Pb-free reflows	0/39 post 10x add'l reflows
22x22 mm	L3 Precon + TCB (-55°C / +125°C)	0/95 post 700 cycles
	L3 Precon + HAST (130°C/85% RH)	0/56 post 100 hrs
	L3 Precon + add'l Pb-free reflows	0/21 post 10x add'l reflows

Table 4: Enabled BGA package reliability results for nHF vs. HF packages

Substrate Form Factor	Reliability Stress	HF Fails	nHF Fails
13x14 mm	Shock Test (40 mil test board, 5 drops, 295G input)	0/9 CTF cracking	0/9 CTF cracking
	Temp Cycle test (-25 to +100°C)	0/8 through 2500 cycles	0/8 through 2500 cycles

Table 5: Core material UL-94 flammability rating after extended Pb-free reflows

Core Mat'l	Criteria	Number of Reflows at 260°C		
		0	5x	10x
Halogenated	Passed V-0 criteria?	Yes	Yes	Yes
HF		Yes	Yes	Yes

Electrical Performance

Lastly, from a performance standpoint, nHF and HF components were tested side by side to determine the electrical impact of an HF core. Test results confirmed no impact on maximum operating frequency due to the use of an HF core, and electrical performance parity between nHF and HF was achieved.

In summary, a careful choice of HF core materials enabled Intel to introduce and ramp HF-compliant substrates that met assembly processing requirements as well as use condition component and board-level reliability requirements while maintaining the electrical performance of the package. Additionally, the selection of HF core materials with similar properties to nHF core materials enabled Intel to use existing recipes for component and board-mount assembly. This reduces the impact on the factory, and it potentially minimizes the impact on board-mount processes at customer sites due to the use of HF substrates.

CONCLUSION

The implementation of Pb-free and HF-compliant packaging materials is critical to Intel, and it is vital to Intel's broader strategy to support an environmentally sustainable future for its industry. Enabling of Pb-free materials with robust resistance to Pb-free assembly processing was made possible through prudent downselection of solder TIM materials, FLI solder, and substrate materials, 2LI ball attach and surface-mount paste materials, and BGA solder metallurgy. The transition to HF-compliant substrate materials posed significant assembly and reliability challenges that were addressed and successfully met. The transition to Pb-free packaging and HF-compliant substrate materials was achieved through close collaboration with industry partners, suppliers, and our customer base, and we had to establish a robust materials supply infrastructure to sustain the environmentally green micro-electronic packaging ecosystem. The adoption of these initiatives has maintained Intel's OGA philosophy in micro-electronic packaging.

ACKNOWLEDGMENTS

This work would not have been possible without the support of several people within Intel, specifically in ATD pathfinding and development, and in assembly materials, modules, quality and reliability (Q&R).

For the HF work, the authors specifically thank Derek Hetherington and Yi He in the AMCL chemical lab for substrate halogen content evaluations. The authors also thank Raiyo Aspandiar for his contribution to the surface-mount challenges section for 2LI.

REFERENCES

[1] Daewoong Suh et al., "Effects of Ag content on fracture resistance of Sn-Ag-Cu lead-free solders under high strain rate conditions," *Materials Science and Engineering A*, 46-461, 2007, pps. 595–603.

[2] Vasu Vasudevan et al., "Slow Cycle Fatigue Creep Performance of Pb-free (LF) Solders," *ECTC,* 2007.

[3] Irina Boguslavsky et al., "Recrystallization Principles Applied to Whisker Growth in Tin", APEX Conference, March 31, 2003, Anaheim, CA, pps. 1–14.

[4] R. Mahajan et al., "Advances and Challenges in Flip-Chip Packaging," *IEEE 2006 Custom Integrated Circuits Conference,* (CICC), 2006.

[5] V. Wakharkar et al., "Microelectronic Packaging Materials Development & Integration Challenges for Lead Free," *TMS Conference*, 2006.

[6] Carl Deppisch et al., "Material Optimization and Reliability Characterization of Indium Solder Thermal Interface Materials for 1262 Package technology," *IATTJ Journal, Vol. 5*, 2002, pps. 227–240.

[7] *Smithell's Metals Reference Book*, E. A. Brandes and G.B. Brook (editors), Seventh edition, 1992, Butterworth Heineman, London, England.

[8] JEDEC J-STD-020.

AUTHORS' BIOGRAPHIES

Mukul Renavikar is a Senior Packaging Engineer in the ATTD-Materials group. He has been with Intel for over four years and leads the Metallurgy/Materials Core Competency group in Assembly and Test Technology Development (ATTD). Mukul received his Bachelor's degree in Metallurgy from the College of Engineering, Pune, and M.S. and PhD degrees in Materials Science from the University of Pittsburgh, PA. His e-mail is mukul.p.renavikar at intel.com.

Neha Patel is a Senior Materials Engineer in the ATTD group. She has been with Intel for three years and has worked extensively in the area of first-level interconnect polymer materials and thermal interface materials. She obtained her bachelors degree in Physics from St. Xaviers college, Bombay and M.S. and PhD degrees in Physics from Case Western Reserve University Cleveland, OH. Her e-mail is neha.m.patel at intel.com.

Ashay Dani obtained his B.S. degree in Chemical Engineering from UDCT, Bombay, India and a PhD degree in Chemical Engineering from Clemson University. Ashay joined Intel in 2000 and has managed several packaging materials and suppliers' technology development, certification and ramp. His e-mail is ashay.a.dani at intel.com..

Vijay Wakharkar is a Principal Engineer in ATTD and has worked at Intel for 17 years. Vijay got his Bachelors degree from the College of Engineering, Poona in 1982 and his PhD degree in Materials Science and Engineering in 1988. Prior to joining Intel, he spent two years as a Post Doctoral Fellow at IBM Almaden Research Center in the Polymer Materials Division. Currently he is responsible for managing the development and path-finding of materials technologies that end up in ATTD's packaging solutions. Vijay was one of the founding editors of the Intel Assembly Technology and Test Journal (IATTJ) and also serves on Intel's Thermo-mechanical Patent Committee. His e-mail is vijay.s.wakharkar at intel.com.

Omar Bchir is a Senior Packaging Engineer in the substrates pathfinding group with Intel-ATTD, and he has been with Intel for over three years. His focus areas have been metallization, integration, and HF enabling. Omar obtained his Bachelor's degree in Chemical Engineering from the Georgia Institute of Technology in 1997 and his PhD degree in Chemical Engineering from the University of Florida in 2004. His e-mail is omar.j.bchir at intel.com..

Amruthavalli Pallavi Alur is a Senior Packaging Engineer in the substrates pathfinding group with Intel-ATTD, and she has been with Intel for two years. Her focus area includes Mechanics of Materials and HF enabling. Pallavi obtained her Bachelor's degree in Metallurgy and Materials Sciences from the Institute of Technology (India) in 2000 and her PhD degree in Materials Science and Engineering from Brown University in 2005. Her e-mail is amruthavalli.p.alur at intel.com.

Charan Gurumurthy manages ATTD's substrate pathfinding group and has been with Intel for over five years. Before joining Intel, he worked at IBMs Microelectronics Division in Endicott, NY. Charan obtained his Bachelor's degree in Mechanical Engineering from the Indian Institute of Technology in 1994 and his PhD in Materials Science and Engineering from Cornell University in 2000. He has 30 + patents awarded/filed and 10+ journal publications. He currently serves as "Associate Editor" of IEEE Transactions on Advanced Packaging and is a senior member of IEEE. His e-mail is charan.k.gurumurthy at intel.com.

Roger Stage has a PhD degree in Mechanical Engineering from Arizona State University. Dr. Stage manages package quality and reliability engineering in Intel's Assembly Technology Development Quality & Reliability organization in Chandler, Arizona. His e-mail is roger.w.stage at intel.com.

George Arrigotti is the compiler and editor of the Pb-free Board Assembly MAS, which is a materials and process reference document for customers. He works in ATTD Customer Technology and Manufacturing. He is currently the Initiative Development Team leader for High-Bandwidth Digital Content Protection (HDCP), helping customers prepare their designs and manufacturing for Blu-ray and HD DVD. George joined Intel in 1979 and has held a variety of board-related positions, including board design, product engineering, test engineering, manufacturing technology development, and customer manufacturing enabling. He earned his B.S. degree in Electrical Engineering from Gonzaga University in Spokane, Washington, and his M.S. degree in Computer and Electrical Engineering from Oregon State University in Corvallis, Oregon. George is a member of IEEE, has served as a Segment Roadmap Champion for NEMI, and holds ten Intel patents. His e-mail is george.arrigotti at intel.com.

Vasu S. Vasudevan graduated from Purdue University with a PhD degree in Materials Engineering. He is a quality and reliability engineer for board-level technologies and he has been working in Intel for the last seven years. He is primarily involved in board-level reliability assessment including package-to-board interaction for new technologies. He is also responsible for proliferating Pb-free reliability standards, techniques, and new learnings to industry as needed. He has been working in the area of electronic manufacturing and technology development for more than 14 years. He has more than 30 technical publications and holds two patents. He is an active member in IPC and iNEMI. His e-mail is vasu.s.vasudevan at intel.com.

THIS PAGE INTENTIONALLY LEFT BLANK

Making USB a More Energy-Efficient Interconnect

Barnes Cooper, Mobile Platforms Group, Intel Corporation
Paul Diefenbaugh, Corporate Technology Group, Intel Corporation
Jim Kardach, Mobile Platforms Group, Intel Corporation

Index words: USB, low power, power friendly devices, platform power management

ABSTRACT

The ’64 Mustang is a classic: a car that people still talk and reminisce about 44 years on. Do you mess with success and change a winning formula? No, but you do update a design to fix weaknesses in the original (better audio, air-conditioning, reliability, etc.) and to address new consumer desires as the market changes (efficient engines to address fuel economy, catalytic converters to address the need for a cleaner environment, etc.).

The Universal Serial Bus (USB) is a classic of the computer world. It was introduced in 1996 and is now a ubiquitous computer interface. When it was developed in the mid ‘90s it was targeted for mainstream computers of the time, optimized primarily for consumer ease of use and low device cost. Around 2002, USB 2.0 was introduced offering a performance bump to 480 Mbps; again optimized to meet similar criteria.

Although many characteristics of the USB are top-notch, its impact on *platform* power consumption has been downright abysmal. While power consumption was not an important criterion of its original design, the USB has become a defacto feature for battery-powered platforms where low power is key. In addition, global concerns over energy consumption and carbon emissions have made energy efficiency an important market requirement even for desktop and server systems [1]. Therefore, like the classic Mustang, it’s time to overhaul the USB in a manner that preserves the goodness which has helped make it such a successful interconnect.

In this paper we first outline the USB power issues and look at their impact on mobile platforms. We then discuss ways of resolving these issues. Although the focus here is clearly on notebook systems, most issues and solutions apply to other systems as well.

INTRODUCTION

To comprehend USB’s power problems you first need to have a basic understanding of how it works. We won’t try and make you an expert on USB architecture; rather, we will just provide enough detail so you can understand the fundamental problems and how the proposed fixes address these.

The root of most of the power issues is the fact that the USB is based on an architecture that constantly polls devices. Although this creates a simple and low-cost device model, it is fundamentally inefficient—especially when the device is idle or has little data to transfer. Specifically, a USB device is incapable of transferring data or generating an interrupt without being polled by the host. The best it can do is indicate the rate at which it wants to be polled in the event that activity occurs. This rate is typically assigned statically when the device is first configured and tuned for highly active phases (e.g., to maximize throughput).

We will go into a little more detail about how a USB device is designed to work in this polled environment and then discuss why polling creates power problems.

Figure 1 illustrates the behavior of normal (non-polled) data transfers for PCI devices. In this bus model, devices are generally implemented as fully capable bus masters. When a PCI device needs to transfer data it simply requests control of the bus and initiates one or more cycles to main memory (green line #1), which also results in a snoop cycle to the CPU (green line #2) to ensure data consistency in case the memory contents reside in the CPU cache.

Contrast this to the USB model where the device must wait until the next time it is polled by the host to transfer data, or more importantly, the host must continually poll a device just to see if it has data to transfer. The USB

provides two general models for data transfers: synchronous and asynchronous. Synchronous transfers are polled at a guaranteed periodic rate with a maximum frequency of once every microframe (125 microseconds). This corresponds to the Isochronous and Interrupt endpoint types. Conversely, asynchronous transfers are not polled at a guaranteed rate, but for most implementations this occurs quite frequently (many times per microframe) to achieve high data throughput when needed. Bulk and Control endpoints belong to this transfer type.

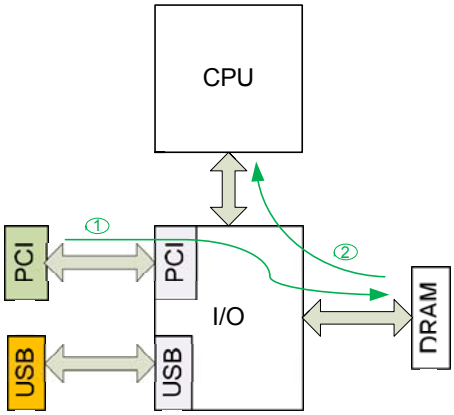

Figure 1: PCI data transfer (non-polled)

In addition, many USB host controllers rely on main memory for their schedule information. Data structures within the USB schedules inform the host controller of the (active) synchronous and asynchronous endpoints that need to be serviced, the polling frequency for synchronous endpoints, memory locations for data transfers, etc. The host controller must access these structures frequently, both to understand when endpoints need to be serviced (polled) and to initiate each transfer request—regardless of whether data are actually transferred.

Figure 2 illustrates the behavior for a typical USB Bulk IN transfer (read from device, write to main memory). The host controller first reads the transfer descriptor information from its schedule in main memory (red line #1), which in turn causes a snoop cycle to the CPU (red line #2) to maintain cache coherency. Once read, the host controller initiates the transfer to poll the targeted device (red line #3). If the device has no data to transfer it returns a NAK response (tan line #1). Otherwise, an ACK is returned along with whatever data the device needs to transfer (tan line #1), which the host controller then writes to main memory (tan line #2), and again causes a snoop cycle to the CPU (tan line #3).

USB transfers are inherently less efficient than equivalent PCI transfers, requiring a total of six cycles (three being snoops) versus two cycles (one snoop) on PCI. But the bigger issue is that USB endpoints that have no data to move (constantly NAK) continue to be polled by the host resulting in a fairly active USB subsystem that generates frequent memory accesses, snoop cycles, and USB transfers. This behavior does not occur on PCI or other non-polled interconnects. Thus, USB works quite hard at doing nothing, which translates into poor energy efficiency.

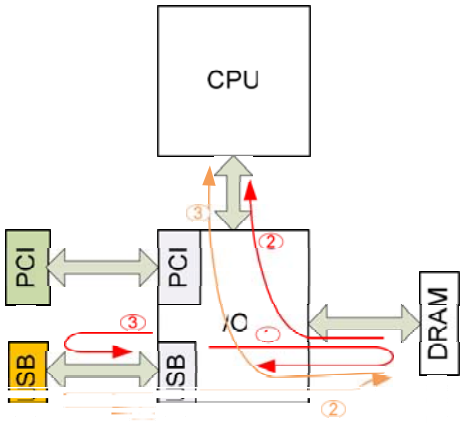

Figure 2: USB data transfers (polled)

It is also important to notice the majority of power increases occurs upstream of the USB host controller. For example, a host controller polling a single Bulk IN endpoint can generate bursts of activity every 8-16 microseconds (us), which prevents most of the core logic (CPU, memory, backbone busses, clocking, etc.) from entering a low power state. This in turn can have a huge impact on platform idle power and drastically decrease battery life.

Background on USB 2.0

The Universal Serial Bus 2.0 specification [2] is defined by the Universal Serial Bus Implementer's Forum, Inc. (www.usb.org). It supersedes and is backwards-compatible with the USB 1.1 specification. USB 2.0 encompasses three distinct data rates: low-speed at 1.5 Mbps, full-speed at 12 Mbps, and high-speed at 480 Mbps. USB 2.0 uses a 4-pin bus with two differential signaling lines (D+/D-). Fundamentally, the USB 2.0 bus is a polled bus in that data and control transactions are initiated by the host, not the device. Because polling directly translates to increased power consumption across the platform, device design techniques are especially important. The USB 2.0 bus standard has a low power state known as Suspend, but today the latencies associated with entry and exit make it problematic to use as a dynamic flow control and link power management mechanism.

The USB 2.0 specification defines four distinct traffic classes (control, bulk, periodic, isochronous) and three data rates (low, full, high). This is typically managed on

Intel® Architecture (IA) platforms using two different host controller types: Enhanced Host Controller Interface (EHCI) for high-speed devices and Universal Host Controller Interface (UHCI) for low- and full-speed devices.

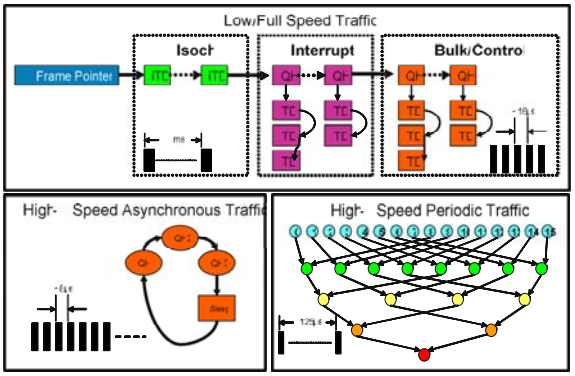

Figure 3: Host controller schedules

Figure 3 illustrates the various schedules, traffic classes, and data patterns for low and full-speed transactions associated with low-, full-, and high-speed devices. For low- and full-speed devices serviced by the UHCI controller, the host controller maintains a frame list pointer that references a physical address in main memory. The host controller parses this schedule every frame (1ms interval) to fetch memory structures (descriptors) that tell the host controller how to poll devices. The operating system (OS) software is responsible for populating the schedule. This specifies which transactions the host controller will attempt during each frame. In the Windows* OS, periodic transfers are layered first starting with isochronous endpoints that are allocated a fixed bandwidth. After this, the OS places interrupt endpoints that are generally polled at some derived periodicity, typically using a binary tree (poll rates of 1ms, 2ms, 4ms, 8ms, 16ms, 32ms, etc.). Bulk and control endpoints are added next and typically arranged as a linked list. The host controller typically parses the periodic elements once per frame, spending the rest of its time (until the next frame) processing bulk and control endpoints.

Table 1: General UHCI/EHCI power implications

	UHCI (Low/Full-Speed)		EHCI (High-Speed)	
	Periodic	**Asynchronous**	**Periodic**	**Asynchronous**
Bus Cycles	Every 1ms	Every ~16us	Every 125us	Every ~8us
Power Impact	Low	Very high	High	Very high

The EHCI controller services USB 2.0 high-speed devices and contains two distinct schedules. The asynchronous schedule consists of bulk and control endpoints that are typically arranged as a linked list. The periodic schedule contains isochronous and interrupt endpoints that are linked at a specific periodicity. The EHCI controller is capable of processing periodic transactions at an accelerated rate referred to as a microframe (125us)—eight times more frequently than UHCI. Thus, periodic transfers may be scheduled at a maximum rate of once every microframe (125us).

Table 1 summarizes the platform power implications when servicing low-, full-, and high-speed endpoints using traditional UHCI and EHCI host controller designs.

Effect of USB Activity on System Power

When bus master traffic is generated by a USB host controller on an otherwise idle system, the platform will immediately transition out of a low power state to process this traffic. This flow is represented in Figure 4 which loosely depicts an Intel® Core™ 2 Duo mobile processor-based system.

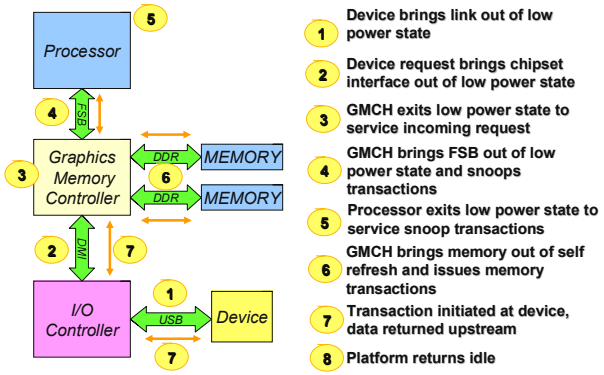

Figure 4: System power impact of USB activity

Because this activity is a platform-wide event, the resulting power impact can be large. Figure 5 illustrates a bus master transfer from a WLAN device fielding a keep-alive packet from an 802.11g access point. Although the actual transfer is short-lived, the component and platform power scales up dramatically to process this activity.

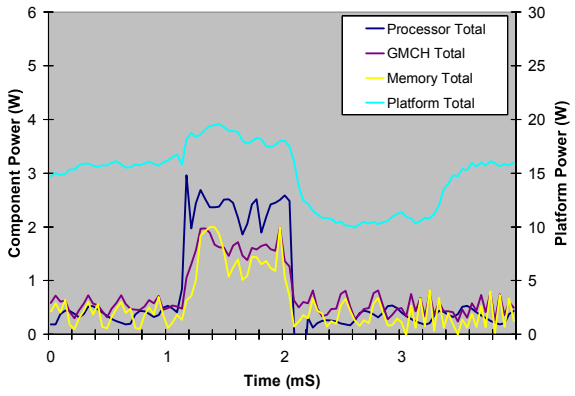

Figure 5: Platform power impact for WLAN activity

Thus, the general solution for addressing USB's power issues requires that we significantly reduce the amount of activity the host controller generates, especially when USB devices are otherwise idle (no data to transfer).

ARCHITECTURE

Addressing the power issues associated with USB is a challenging one. Figure 6 depicts the high-level vision for a truly energy-efficient model for USB. The general idea is to keep the entire path from main memory through the host controller and down to the device completely quiescent until meaningful data needs to be transferred, thereby transforming today's continuously polled architecture to one where devices are only polled when needed.

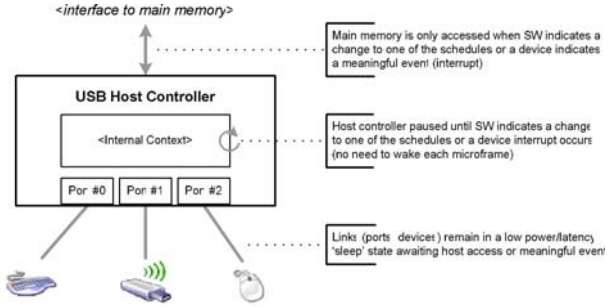

Figure 6: Energy-efficient USB vision

But in order to maintain compatibility with mainstream OSs it was important to avoid changes to the upper levels of the USB software stack. This focused the scope of our solution on the lower levels (miniport driver and hardware), as illustrated in Figure 7.

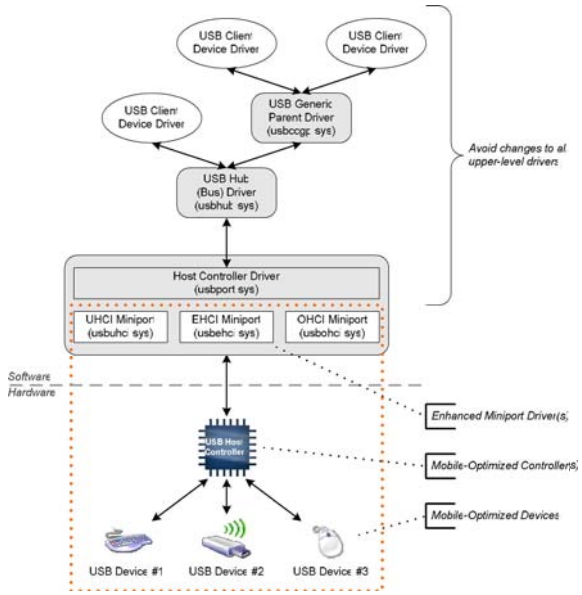

Figure 7: Design constraints

Making USB Power Friendly

In the next few sections, we discuss various energy-efficiency optimizations based upon the following criteria:

1. If no devices are connected or no work is scheduled then USB hardware should remain in a low-power state.
2. Suppress host-side activity (upstream of the host controller, e.g., to main memory) when there is no meaningful work to do.
3. Suppress device-side activity (downstream of the host controller, on the USB bus) when there are no data to send to/receive from devices.

Miniport Drivers

Because of the polled architecture, the host controller's interaction with devices is very important, and if it is not done properly it can adversely affect platform power. In the architecture overview we talked about host controller schedules and how these are used to poll devices and to perform data transfers. Proper management of these schedules is absolutely necessary for producing a power-friendly USB subsystem.

For example, suppose no devices are attached to the system. Obviously a power-friendly USB software stack should schedule no work when there are no devices, and it should immediately remove all associated work from the schedules when a device is removed. If this sort of basic "schedule" and "controller" management is not performed well, any additional power-efficient enhancements will have limited impact. Thus, it is critically important to ensure the miniport drivers do effective work scheduling, turn controllers off when not used, and remove all associated descriptors from host controller schedules when devices are unplugged, disabled, or the work has completed.

Several critical changes were identified in Windows XP* SP2 that have resulted in tremendous power savings. Intel worked with Microsoft engineers to develop these changes and make them available for both Windows XP and Windows Vista*. This includes support for the UHCI run/stop bit, EHCI run/stop, and asynchronous/periodic schedule enable bits, as well as aggressive schedule idle detection. These software optimizations have in turn enabled other hardware optimizations, which we discuss in the next section.

Host Controllers

New features for Intel's mobile USB host controllers were identified to allow for power management opportunities when one or more schedules are enabled (endpoints

present and active). The enhancements include the following key concepts.

Caching

The Caching technique allows the host controller to store schedule information (descriptors) in controller-local memory in order to significantly reduce accesses to main memory, particularly in the case where devices are relatively idle. These data are typically stored in an abbreviated format where just enough information is provided to generate a transfer request (poll). Figure 8 illustrates this technique.

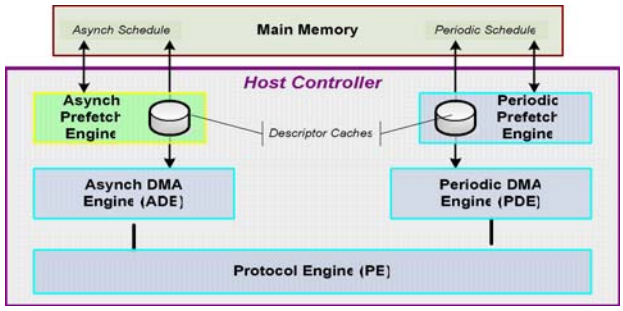

Figure 8: Caching technique

If the device NAKs the transaction, the host controller remains completely idle (since the information needed to generate the transaction was stored locally). If the device ACKs the transaction, the host controller typically must open the path to memory to move the actual data (to/from the device).

The caching feature is especially helpful for endpoints that observe a high NAK rate (for example, streaming or networking devices with bulk asynchronous endpoints open all of the time).

Deferring and Link Power Management

Although caching is fairly good at quiescing host-side activity it does nothing to address downstream (device-side) activity. The USB 2.0 Suspend state was originally intended for this purpose, but it is very difficult to use because of entry and exit latencies and other limitations. It takes considerable time to enter and exit this state (3ms + OS overhead for entry, 30ms + OS overhead for exit), and devices are severely limited in the amount of power they can consume while residing in this state. Additionally, Suspend is coupled with the device D3 state where the OS assumes hardware context is lost (and thus the device needs to be re-initialized and context restored upon exit), which adds significant latency and often interrupts device functionality.

The L1 state is a new Link Power Management (LPM) state that addresses the key deficiencies of the existing Suspend state (herein referred to as L2) by reducing state latencies and decoupling the link state from the device state (allowing the device to remain in D0). The L1 state is intended to be used dynamically when the device is operational (D0), but otherwise idle, and able to quickly enter and exit this low power state without disrupting normal operation. Host controllers can safely negotiate L1 entry with idle devices, progressively decreasing downstream (device-side) activity until all devices reside in L1 (or L2), at which point no downstream activity will occur until either the host or device wakes the link to an active (L0) state.

The L1 transitions have significantly lower entry and exit latencies (10s of µs) than those of L2 (10s of ms). As with L2, both device- and host-initiated wake events are supported from the L1 state, noting that L1 device-initiated wake events play a prominent role in another key technique known as Deferring.

Supporting the L1 state requires modifications to both USB host controllers and devices. The L1 state is a new feature that augments USB 2.0 power management; it does not replace the existing L2 (suspend/resume) mechanism. The proposed L1 definition is backward compatible in that a new host can determine whether a device supports L1. A new device will continue to work properly with legacy hosts (obviously without L1 transitions), and old devices will continue to work on new host controllers. The only time L1 will be used is when a device acknowledges support for this feature on a new host controller.

The policy for using the L1 state is platform and implementation specific and will likely depend on the type of endpoint being served by the host controller. For periodic (interrupt or isochronous) transactions, the host controller would likely implement a policy whereby the device is immediately placed into the L1 state as shown in Figure 9.

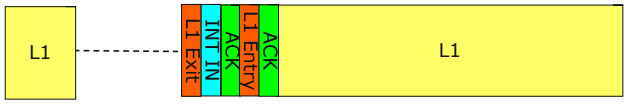

Figure 9: Example L1 policy for periodic devices

For asynchronous (bulk or control) transactions, the host controller would likely implement a policy whereby the device is polled some number of microframes or frames at the nominal asynchronous poll rate before attempting to transition the device to L1, as shown in Figure 10. This is done in order to reduce the overhead for devices that stall for short periods between subsequent data phases.

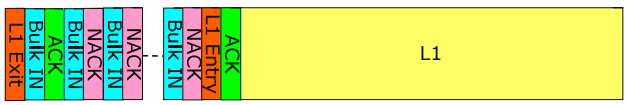

Figure 10: Example L1 policy Asynch devices

The L1 state benefits all types of devices and traffic patterns, and when coupled with the associated host controller enhancements, it can aggressively save power across the entire platform by allowing the entire USB subsystem to enter and remain in a low-power state until some meaningful event occurs.

Devices

When we analyzed the behavior and power impact of many USB 2.0 peripherals currently in the market it became evident that a clear set of device recommendations was required to promote energy-efficient designs [3], both for present-day systems and forward looking to future optimizations. We summarize these recommendations in this next section.

Periodic-Triggered Asynchronous Transfers

In general, it has been observed that there is a multitude of devices that generate traffic in a continuous stream using bulk (asynch) endpoints, with a high NAK rate (>90%). While the design is simplistic, it has a key downfall: bandwidth, and hence device buffering/throughput is highly variable and hard to quantify. A principal recommendation is to use an interrupt (periodic) endpoint to indicate that a device requires service and to use bulk endpoints dynamically for moving data to or from the device. This concept is termed "periodic-triggered asynch" and is illustrated in Figure 11.

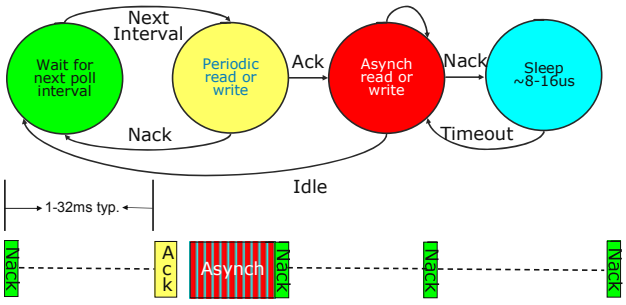

Figure 11: Periodic-triggered asynch transfers

By using this scheme, the response time is well defined (namely, the polling interval requested), and streaming bandwidth is more carefully managed for data movement. This is also a more platform-friendly approach in that it preserves bus bandwidth (a shared resource for USB) for use by other devices. The key virtue of course is that this scheme is much more power friendly as illustrated by the idle time between USB poll events.

Minimize Polling Rate for Periodic Endpoints

Using the aforementioned periodic-triggered asynch scheme or when periodic interrupt/isochronous endpoints are used for other purposes, it is important to maximize device buffering such that the poll rate of the device can be as slow as possible. A power-friendly device should employ poll rates of at least 1ms, and preferably 2-4ms or longer. It may be also possible to support endpoints with different periodic rates that are used selectively based on the bandwidth needs of the device. In such a case, if the device has a high-speed connection, device buffering may mandate a 1ms poll rate interval: when the device has a slower connection, the device may have sufficient buffering to tolerate a much longer (e.g., 4-8ms or higher) poll rate.

Use Isochronous Transfers for Streaming Devices

One common characteristic observed for streaming devices is the use of bulk endpoints for data transfers. There are several problems with this approach. First, asynchronous bandwidth is shared across all ports on a given controller, and thus, realized bandwidth may vary dramatically depending on whether other devices are actively consuming bus bandwidth. This can be readily observed with two devices that use asynchronous transfers for streaming content: in many cases the streams become unstable whenever both devices are active on the same host controller at the same time. This is because bandwidth is shared across a single host controller instance, highlighting the fact that USB is fundamentally a broadcast bus where multiple streams are time-sliced rather than served concurrently.

On the contrary, the isochronous traffic class is time scheduled, and bandwidth is properly allocated by host software. As such, a device can receive a dedicated amount of bandwidth to service its endpoint where this traffic effectively runs at a higher priority level than asynchronous transfers. Moreover, since isochronous transfers reside on the periodic schedule, the effectiveness of power management techniques are generally better (versus the asynchronous schedule)—at least when the periodicity of these transfers approaches 1-2ms or more.

Use LPM L2 Dynamically (Selective Suspend)

Devices should support and use Suspend (L2) whenever the device is idle and use of this state is possible, occasionally waking to look for activity, incoming connections, or other device state changes. This is important as a device should not continuously post periodic (and certainly not asynchronous) transfers when it is not active or actively connected. For example, in the case where a USB network device is scanning for network connectivity, it should take care to do this very infrequently or provide hardware capabilities in the device to do this without requiring continuous transfers from its function driver. For other classes of devices, inactivity can be easily determined by whether the device is in use or not (for streaming devices such as audio/video, occasional use devices such as fingerprint sensors and GPS). The most difficult class of device to make use of Suspend is typically human interface devices (HID) such as mice and

keyboards, where the end-user may perceive the increased latency associated with L2 entry/exit (e.g., choppy mouse movement) when using these devices.

Use LPM L1 Dynamically
As discussed previously, the long-term path to fully addressing the power efficiency limitation of USB 2.0 requires that the device and platform implement a new low-power link state known as LPM L1. For device implementations, it is important to note that entry into the L1 state should not result in any loss of functionality, as it is intended to be used while the system and device may be idle between bursts of activity. It is also important that the device pay attention to the Host Initiated Response Duration (HIRD) field in the host command sent to the device to request entry into the L1 state. This parameter is indicative of the depth of lower power state the platform is expecting to enter. If the platform is semi-active, the field may indicate a light response duration (e.g., <200us), whereas if the platform and devices are more deeply idle, the field may indicate a bigger number (~1ms). The device should use this parameter to control the depth of power management in use by the device to save power, for example, by shutting off PLLs only when a "long" (~1ms) L1 entry transaction is identified.

Design True Composite Devices
The use of integrated hubs within multifunction devices has been a common practice to streamline and simplify hardware implementations. Although convenient, this approach has a number of power management pitfalls and is therefore strongly discouraged. For example, many Deferring scenarios are not feasible for devices that are attached to a downstream hub rather than directly to one of the host controller's root ports.

The most energy-efficient designs involve true composite devices. Here multiple logical functions (devices) reside behind a single USB 2.0 physical device interface where each independent function is exposed as sets of one or more endpoints.

Application/Driver Synchronization
Many devices such as streaming (media playback, cameras) or occasional use (fingerprint sensor, GPS) are bundled with application software. It is critical that when the application stream is shut down, care must be taken in the device function driver to ensure that the application properly cleans up driver requests on exit or inactivity (pause, mute, etc.) to avoid dangling transactions pending on the device; otherwise, these transactions remain un-serviced or are continually retried.

Avoid Polling Integrated Buttons
Many devices such as integrated cameras support a so-called "Instant On Feature," whereby the device has local buttons that are typically serviced by a periodic interrupt endpoint. The buttons require a continuously running periodic interrupt endpoint to poll the button, and this wastes power. It is recommended that devices purposefully designed for mobile platforms do not support buttons (better to enable through applications or traditional keyboard hotkeys), or if they do support buttons that must be functional, you should work with the platform designer to provide platform-level notifications mechanisms through sideband signals and Advanced Configuration and Power Interface (ACPI) BIOS modifications. By using such a scheme, the notifications may be delivered on demand, and the function driver can be the target of these notifications providing the same net effect for Instant On Features without having to continuously run the periodic schedule.

If the button can't be avoided, then architect a very long poll interval for the button (10s to 100s of milliseconds) to reduce the inevitable platform power impact. Such a long polling interval will give other hardware optimizations a chance to kick-in (Caching, Deferring, L1, etc.).

Challenges

Clearly there were and are numerous challenges associated with making USB 2.0 an energy-efficient interconnect. We are quite pleased with the progress thus far, but note the biggest remaining challenge is the broad and timely adoption of these devices, OS, and platform features by the ecosystem.

RESULTS

There are two main reasons why the USB needs to be overhauled: to reduce the power directly consumed by USB devices and host controllers, and (more importantly) to eliminate the drastic increase in power consumption that current USB behavior has on other platform components. The techniques described herein fully address both.

As an example, Figure 12 illustrates the total platform power savings opportunity for the *Caching* and *Deferring* techniques on an Intel® Core™2 Duo mobile processor-based system. The results were derived using measured data and best-known practices. Note that platform power increases by a whopping 5.7W when a high-speed bulk endpoint is active but constantly NAKing, as is the case for most wireless network devices. Here the *Caching* technique achieves a ~4.7W improvement by localizing activity to the USB host controller (EHCI), link (Port), and downstream device (WLAN), enabling the processor, GMCH, memory, and DMI interconnect to enter and remain in a low-power state. The *Deferring (LPM L1)* technique addresses most of the residual power by

allowing the controller, port, and device to only become active when necessary (no longer polled).

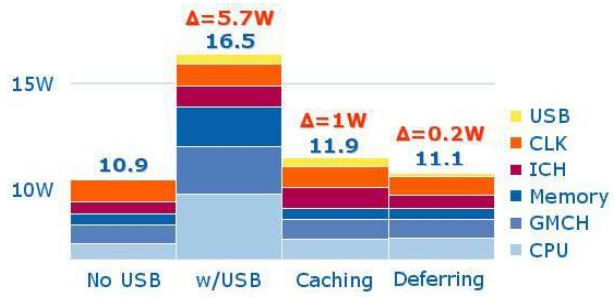

Figure 12: Resulting platform power savings

DISCUSSION

We discussed a number of techniques to transform USB into an energy-efficient interconnect in this paper. These techniques are based on several basic principals for power-friendly design:

1. An efficient transfer rate (bandwidth per Watt) is important but should not be the only focus when designing an interconnect. Specifically, robust and low-power idle states are an absolute necessity.
2. Power management states should be defined such that these states can be used effectively across a variety of idle to pseudo-active scenarios.
3. It's all about platform power. Optimizing for low subsystem power while ignoring the subsystem's impact on the rest of the platform is a recipe for failure. Developers need to analyze both component and platform power consumption in order to catch unexpected behavior or other power-related artifacts. A poorly designed device or host controller may only consume tens of milliwatts but this can result in a multi-watt increase throughout the platform.
4. Good idle behavior is key. An idle device should burn (nearly) zero power; the same applies to buses and host controllers. "Do nothing efficiently!"

Although USB has always had a relatively efficient data transfer rate, this provided little advantage to the platform designer because of the interconnect's significant idle penalties.

Concerning a low-power idle state, the original Suspend state was intended to address a variety of usage cases, but high latencies and other characteristics have prevented its widespread use. And although certain flow control mechanisms do exist, these were designed to address platform performance (vs. power) concerns and failed to address the fundamental issues of constant polling and associated upstream and downstream activity. The new LPM L1 state fills this void.

The Caching technique addresses upstream (host-side) activity that has prevented much of the platform from residing in a low-power state even when all USB devices (and the rest of the system) are pervasively idle.

The Deferring technique addresses downstream (device-side) activity, where entry into L1 state is used as a means by which host controllers can safely defer polling when all endpoints for a device become idle. Here the host can resume the device (and thus polling of its endpoints) when it has meaningful data to transfer and vice versa.

The combination of techniques has transformed USB from a constantly polled architecture with frequent activity to one where activity occurs only when there are meaningful data to transfer, approaching the energy efficiency of other non-polled interconnects such as PCI.

CONCLUSION

Like the original '64 Mustang, USB is a classic. Although it has been widely successful, the time has come for an "energy efficiency" overhaul. The key attributes that contributed to USB's success (simple, low cost, decent bandwidth) needed to be preserved while at the same time modernizing and enhancing this interconnect for today's environment where "green-ness" and "power efficiency" have become equally important.

We have demonstrated techniques and offered suggestions that transform USB into a much more energy-efficient interconnect, primarily by optimizing the idle behavior of USB host controllers and devices. This complements USB's relatively good bandwidth per watt characteristics to produce a robust and power-friendly solution.

ACKNOWLEDGMENTS

We thank John Howard, Corporate Technology Group, Intel Corporation; Karthi Vadivelu, Client Components Group, Intel Corporation; and Michael Derr, Client Components Group, Intel Corporation.

REFERENCES

[1] "Energy Star* System Implementation," April 2007, at http://www.intel.com/design/core2duo/316478.pdf

[2] "Universal Serial Bus Revision 2.0 specification," April 27, 2000, USB SIG, at http://www.usb.org/developers/docs/

[3] "Designing Power-Friendly Devices," May 8, 2007, WinHec 2007, at http://download.intel.com/technology/EEP/designing_power_friendly_devices.pdf

AUTHORS' BIOGRAPHIES

Barnes Cooper is a Principal Engineer within the Intel Mobile Platforms Group (MPG) and has been working on mobile platform architecture and power/thermal management since 1992. He was a key contributor to numerous power/thermal innovations such as ACPI 1.0, Enhanced Intel SpeedStep® technology, and multi-core/thread power management. He works extensively with Microsoft engineers on defining OS power management innovations. His e-mail is barnes.cooper at intel.com.

Paul Diefenbaugh is an architect within the Intel Corporate Technology Group (CTG) leading efforts there on Platform Power Management. He holds an M.S. degree in Computer Science from Oregon Graduate Institute. Paul joined Intel in 1992 and has focused on mobile platforms and energy efficiency during the last ten years. He was a key contributor to the ACPI 2.0 specification, the Linux implementation of ACPI and OSPM, and various Intel power-management innovations addressing processors, interconnects, displays, storage, and platforms. His e-mail is paul.s.diefenbaugh at intel.com.

Jim Kardach is a Senior Principal Engineer within the Intel Mobile Product's Group (MPG). He joined Intel in 1986 where for the last 20 years he has worked on notebook processors and chipsets. Jim worked on developing technologies for the notebook platform (including the original ACPI specification). Jim has about 50 issued patents, was runner-up in the 1996 Discover Magazine's Innovation Award, and was one of the original inductees into the Bluetooth Hall of Fame for his work on that standard. His e-mail is jim.kardach at intel.com.

THIS PAGE INTENTIONALLY LEFT BLANK

Intel's First Designed and Built Green Building

Dori Hershgal, Technology and Manufacturing Group, Intel Corporation
Marc Denner, Technology and Manufacturing Group, Intel Corporation
Ayelet Harush, Technology and Manufacturing Group, Intel Corporation
Revital Bitan, Legal Corporate Affairs, Intel Corporation

Index words: Green Building, sustainable construction, Intel's first, LEED, R&D building, energy, productivity, maintenance cost

ABSTRACT

The goal of the "Green Building" project is to reduce the impact of construction on the environment. By doing this Intel joins the world's efforts in sustainable construction. The Green Building project differs from conventional building projects by assigning equal priorities to economical, social, and environmental goals. The new Intel R&D building, located in Haifa, Israel, is designed per the Leadership in Energy & Environmental Design (LEED*) rating system. These were the objectives of the design process:

- To achieve energy and water savings.
- To incorporate environmental friendly materials into the building.
- To implement waste and recycling policies.
- To provide a high-quality internal environment.
- To position and plan the site to reduce the impact on the environment.

The new Intel R&D building is now under construction and is undergoing a certification process. Once awarded, it will become Intel's first certified Green Building worldwide!

This project was initiated to address the need for a "smart" building. The Green Building concept was found to be the most suitable for this purpose, particularly for an R&D building. Research shows that such a building improves tenants' satisfaction and health, enabling higher employee productivity; and it reduces energy and maintenance costs for the owners. The estimated ROI due to operational costs is approximately five years.

In this paper we show the innovation, logic, drivers, and triggers that helped us overcome significant challenges in dealing with an inexperienced local building industry, and in driving innovation in a large organization.

Our engagement with the local construction industry adds a new dimension to the concept of Intel social responsibility.

Our key message is that it is possible to execute a local and innovative initiative successfully, even in global corporations, once the initiative supports the corporate values.

INTRODUCTION

R&D Environment Need for "Smart Buildings"

Design Center building is an example of an ever changing construction environment in a leading corporation. The energy consumed by our existing R&D buildings has increased since 2003 from 2MW to 8MW. That's an increase of 400% in less than four years. Such growth requires many fast conversions of office spaces and energy infrastructure, while at the same time keeping the quality high and the cost low. There are no construction rules for this kind of turnaround; our business continuation policy requires creative engineering solutions to provide adequate power and space, within the same building structure, and without or with minimum interruption to the business. Innovative and creative solutions already resulted from our R&D operational challenges (i.e., installing an ice storage reservoir as an Uninterruptible Chilled Water Supply to a Data Center or water-cooled server racks to optimize Data Center space).

Back in 2004, Intel Corporation decided to build a new development center in Israel. At that time, the practical decision makers in the local construction industry were totally unaware of the concept of Green Buildings (although this concept was known to environmental groups and to a few in the academy). The public at large (excluding environmental groups) had as yet not been exposed much to this concept. In light of this fact, the decision to build and formally certify a first Green

Building, new to Intel Corporation and new to the local industry, is considered a breakthrough.

The experience gained by the local engineers helped establish the "Site Engineering Management Mission," documented in the project scope: *"…New IDC Building will be built smartly, will be flexible to modifications and designed with energy savings in mind…" (Marc Denner, 2004).* This mission defined the scope for an efficient energy facility, with modular and low-cost conversions of offices and spaces, suitable as working environments for researchers, and with effective Total Cost of Ownership. Actually, this common strategy affected the project team's mindset, thus broadening their goal: reduce the impact the building has on the environment, decrease the contribution to global warming, and create a healthier place to work. This "Sustainable Development" approach created a new standard for Intel R&D buildings.

Sustainable Development

The concept of environmental management and education for nature conservation has changed dramatically over the last few decades. A romantic approach was dominant in the early 60s, where conservation of the wild areas outside of the cities was considered to be the issue, and the public was encouraged to go out from the urban environment to nature. During the 70s a more apocalyptic approach was adopted, mainly due to the large number of environmental catastrophes reported during those years: In 1967, the Torrey Canyon oil tanker went aground off the southwest coast of England, and in 1969 oil spilled from an offshore well in California's Santa Barbara Channel. In 1971, the ruling of a lawsuit in Japan drew international attention to the effects of decades of mercury poisoning on the people of Minamata, and later in 1978, in Love Canal, New York, the discovery of toxic waste buried beneath an entire neighborhood attracted the world's attention.

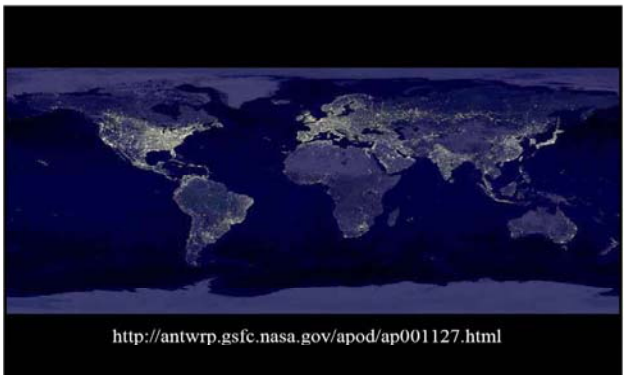

Figure 1: Photo of Earth from space emphasizing the interrelation between development and ecology impacts

Meanwhile, photos of Earth from outer space emphasized the deleterious consequences of our technological "progress," and showed us how truly small a place the Earth is in the universe (see Figure 1) [1].

The 1980s saw a rise in public concern about the role of business in society following a series of notorious environmental and social incidents, and by the early 1990s, the Brundtland Report (1987) and the Earth Summit (Rio de Janeiro, 1992) had inspired a shift in the understanding of the concept of development. These landmark events pushed the idea of development beyond economic growth and the continued generation of wealth defined narrowly in financial terms, to include the concepts of intra- and inter-generational equity and quality of life [2].

The concept of "Sustainable Development" came into general usage following the publication of the 1987 report of the Brundtland Commission, the World Commission on Environment and Development (WCED). This commission coined what was to become the most often-quoted definition of Sustainable Development: development that "meets the needs of the present generation without compromising the ability of future generations to meet their own needs" [3].

Sustainable Development does not focus solely on environmental issues. More broadly, Sustainable Development policies encompass three general policy areas: economic, environmental, and social (Figure 2). Only at a point where all three areas merge, is real sustainable action taken [4].

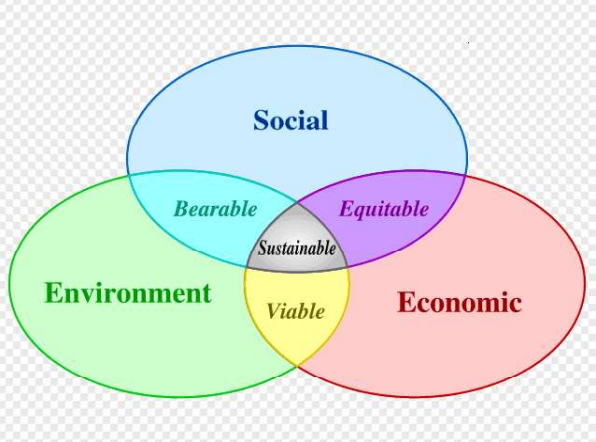

Figure 2: Scheme of Sustainable Development: the confluence of three policy areas

The World Business Council for Sustainable Development, founded in 1995, has formulated the business case for Sustainable Development and argues that "Sustainable Development is good for business and business is good for Sustainable Development" [5].

This concept of Sustainable Development spawned the emergence of the "triple bottom line" concept, which lies

at the heart of corporate responsibility and corporate citizenship. Central to this is the need to measure and report on a company's performance with respect to economic prosperity, social justice, and environmental quality.

The roles and responsibilities of business in society, in particular global businesses, are being defined more broadly. Stakeholder demands are increasingly going beyond the obligation to "do no harm" to the responsibility for being "a positive force" in contributing to worldwide social development goals [6].

Green Building and LEED Certification Benefits

The built environment has a profound impact on our natural environment, economy, health, and productivity [7]. In the United States alone, buildings account for

- 65% of electricity consumption
- 36% of energy use
- 30% of greenhouse gas emissions
- 30% of raw materials use
- 30% of waste output (136 million tons annually)
- 12% of potable water consumption

In Israel, buildings account for [8]

- 30% of the total of electricity consumption, 12% of which is for heating, cooling, and lighting
- 30% of the total energy consumption in the commercial sector, largely used for heating, cooling and lighting
- 7.5 million tons of building waste, generated every year
- 30% of the fresh water consumption

Building green, if it is done right, reduces the impact on the global environment, supports community improvement, increases the useful life of building systems and equipment, increases property values, and contributes to user satisfaction and productivity resulting in reduced absenteeism [9]. It also enhances the corporate profile, supports the local economy, and reduces the building life-cycle cost.

In order to promote environmentally-responsible building in Israel, the Israel Ministry of Environmental Protection published guidelines on the subject in 1997 [10]. The guidelines are based on the environmentally-responsible building guidelines prepared for New York City in 1996 by Dr. Miriam Haran and Professor Victor Goldsmith, on behalf of the Center for Applied Studies of the Environment of the City University of New York.

The guidelines deal with all stages of construction including the following:

- Project and program planning, including definition of the project and consideration of non-construction options, site selection, and building for the long term.
- Design processes including integrated design processes and decision-making models for evaluating trade-offs.
- Building energy use including life-cycle cost.
- Indoor environment.
- Material and product selection including life-cycle analysis.
- Water reduction/management use, storm water runoff, use of non-potable water and landscaping.
- Operation and maintenance considerations.
- Construction considerations including health and safety issues.

There is no characteristic look of a green building. While natural and resource-efficient features can be highlighted in a building, they can also be invisible within any architectural design aesthetics [11]. However, meeting the above guidelines, together with complying with a rating system, provides a relative measure of how "re" the building is. By applying a measurable standard and third-party audit, following these guidelines can ensure recognition and validation to the level of commitment. In recent years, there has been a proliferation of labels for building products said to be produced in an environmentally and socially responsible manner.

The term "eco-label" is any label, seal, or logo used to give purchasers, including architects, designers, and builders, an assurance that the environmental characteristics or production methods of the product used improve the performance of the building envelope [12]. Not all Green Buildings' eco-labels are the same, but most will cover these five broad areas: sustainable site planning, safeguarding water, energy efficiency, conservation of material and resources, and indoor air quality. LEED was chosen as the project's rating system.

LEED is a rating system developed by the U.S. Green Building Council (USGBC), and it is the most acceptable one in the U.S. LEED has four levels, each characterized by earned points: "Certified" (minimum 26 points), "Silver" (minimum 33 points), "Gold" (minimum 39 points) and "Platinum" (minimum 52 points). A level qualifies and quantifies a building's environmental and energy characteristics compared to known common standards and to other certified Green Buildings. This process encourages and accelerates global adoption of sustainable green building and development practices

through the creation and implementation of universally understood and accepted tools and performance criteria. This instrument is used to evaluate environmental performance from a "whole building" perspective over a building life-cycle, providing a definitive standard for what constitutes a green building. LEED for new construction rating system is designed to guide and distinguish high-performance commercial and institutional projects, including office buildings, high-rise residential buildings, government buildings, recreational facilities, manufacturing plants, and laboratories.

The USGBC has asserted that a LEED-Certified level or LEED-Silver level should not cost more than a conventional building [13]. Studies show that LEED buildings fall within the typical cost range of conventional similar building types. A comprehensive study, done by Gregory H. Katz [14], shows that the green cost premium for LEED Silver Certification level, averaged over 18 office and school projects, is around 2% (see Figure 3).

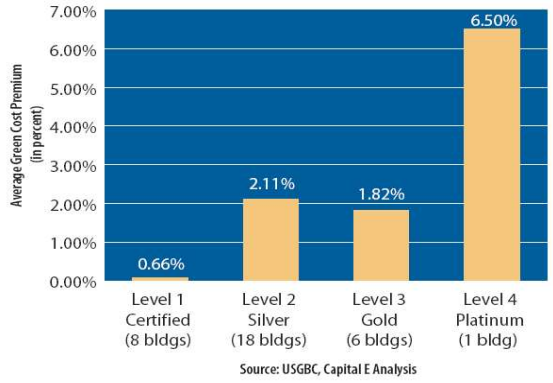

Figure 3: Average green cost premium vs. level of green certification for offices and schools

Davis Langdon [15] updated in 2006 his previous study and showed essentially the same results he published in 2004: there is no significant difference in average costs for green buildings as compared to non-green buildings. Many green buildings' projects end with little or no added cost, and with budgets well within the cost range of non-green buildings with similar programs.

Source: Carnegie-Mellon University

Source: National Research Council, 1998

Figure 4: Personnel and operation costs contribution to building life-cycle costs

As important as they are, real costs are not the only costs to be considered. In many cases, opportunity costs hinge on how a building affects employees as they are by far the largest expense for most companies (Figure 4). Steven Morton from HOK [16] wrote that for offices, salaries are 72 times higher than energy costs, and they account for 92% of the life-cycle cost of a building. On average, Americans spend more than 90% of their time indoors, while case studies show that healthy indoor environments can increase employee productivity, due to less absenteeism, and can cause employees to stay in their jobs longer.

Israeli "Green" Construction Market Capabilities

Many revolutionary "green" developments and solutions carry the stamp of an Israeli company and are constantly being invented by Israeli scientists, developed by Israeli companies, and funded by local investors. Nonetheless, when it comes to implementing these high-technology environmental solutions, it seems that Israel lags behind the rest of the world [17]. Construction management and engineering firms still have limited experience with sustainable design and construction. Trade contractors have never used sustainable construction specifications or equipment and material delivery. Local material suppliers are still not aware of the developing green business opportunities in Israel. Only limited available "eco-labeled" construction materials are available to the local market. In many cases, vendors and suppliers do not know if their product carries a "sustainability" label, such as recyclable content data, and they lack the required supportive documentation.

Difficulties have been evaluated and were acknowledged during the project's planning phase. As a result, the team assumed responsibility to share the experience gained in this project and pave the way for the local construction industry.

DRIVERS

Management Commitment to Social and Environmental Responsibility

Intel co-founder Gordon Moore, a long-time champion of the environment, instilled a legacy of environmental consciousness at Intel that continues today. Intel strongly encourages its employees to apply the same level of knowledge and creativity to solving the environmental challenges of design and production as they do to creating and developing the next breakthrough technology.

"We strive to minimize our environmental footprint and achieve the highest standards of environmental consciousness in everything we do, from how we design

and manufacture our products to how we build and operate our facilities, manage resources, and handle waste materials" [18]. In order to adhere to this environmental philosophy, a "design for environment" approach is being used.

Building and designing the world's most sophisticated products in cutting-edge factories involves addressing many environmental challenges, such as energy efficiency, air quality, water and materials recycling, and more. These challenges grow in complexity as we push the technological boundaries and explore the limits of materials science at Intel.

The Rational for the Adoption of a Sustainability Mindset

It goes without saying that energy-saving designs will cut down on maintenance costs in the long run. This simple example shows how three factors: social, environmental, and economical, merge to form a sustainability mindset. To change someone's mindset to that of Sustainable Development, a socially responsible organization will constantly evaluate the weight of these three main elements during the decision-making process. However, the corporation's true commitment to environmental sustainability is shown when one of the elements is not fully met, and that element is usually the economic one. When the ROI is questionable, only an organization that sincerely adopts and believes in its social business responsibility will decide in favor of a Green Building.

Processes and Procedures for High-Quality Results

Intel operates a Corporate Services (CS) entity to self-manage real estate assets, construction, operation, maintenance, and general services. The associated teams and groups are well educated to ensure that proper codes, standards, and procedures are followed at a defined level of quality. Being a global company, these business procedures are managed practically the same all over the corporation.

The Intel Construction Management Procedures (CMP) database incorporates standards and specifications that are well defined to ensure a building's trademark profile uniqueness, quality, flawless data communication; and that it is based on industrial cost benchmarks. Corporate policies, specifications, and standards are updated through a global consensus change process, after verifying the justifications and the outcome benefits. This process requires that every stakeholder group, affected by the change, will review and confirm all changes before final approval and deployment.

Innovations by nature are being executed based on the same processes. Intuitively, innovations may be implemented and deployed more effectively if they are based on clearly stated corporate values and on the use of existing business processes. In this context, Green Building innovation falls under this category. The major challenge was to justify throughout the approval chain that the changes did indeed comply with Intel's basic rules.

CHALLENGES AND SOLUTIONS

Strategic Core Team

The strategic core team consisted of Project Engineering, Environmental, Health and Safety (EHS), and Public Affairs (PA). These stakeholders were invested in building an environmentally sustainable building, where EHS took the role of environmental responsibility, PA took the role of social responsibility and communication, and engineering provided the practical green building planning, design, and execution.

Corporate Approval for an Innovation

The project's stated mission and the building's type created a clear message that the approach to build a sustainable building required a new mindset. An Internet search revealed the "Green Building" advantages and activities around the world. The "Smart Building" strategy for the new Intel R&D center in Israel has been converted to a "Sustainable Development, Green Building." LEED was chosen as the project rating system. The new Israeli standard SI-5281 was in its development stage at that time, and it was later adopted by the team as a second rating system.

The concept of a green building was first presented to high-level corporate managers to avoid or eliminate potential roadblocks allowing for a flawless review throughout the regular approval process. By getting their approval, midline managers were able to focus on standard construction procedures, without the need to make decisions on innovative strategies. The approval to add capital cost and ROI was left to upper management, where the advantages to the business as a whole, and to the corporate profile, could be better evaluated. As a result, the scope of the project and the budget for a Green Building have been approved. The building's imaging picture and its facilities are shown in Figure 5.

Figure 5: The building's imaging picture and its facilities

Integrated Design Process

All Intel building designs are integrated, which means that Contracted Architects and Engineering (A&E) firms mobilize all the required personnel from the start. Teams, made up of owner's engineers, A&E, and construction managers create owner/contractor groups called Discipline System Teams. By leveraging this process and implementing the same practice in the green building project, we have created a "Green Building" team, made up of both the design managers and the LEED consultant. Owning the management charter and authority, this core team enabled a flawless and directly controlled process, and it coordinated and led single and multi-discipline discussions where ideas, solutions, and cost evaluations were hammered out. For instance, this was the most effective way to bring the Architect, HVAC, and the electrical consultants together at the early stages of the design so that they could discuss natural daylight strategies and consequences, when building envelope options had just been developed. In addition, it was the best method to start the commissioning process. System teams created an effective structure to define and execute the enhanced planned commissioning processes during project design, execution, and delivery.

Preliminary Cost Evaluation and ROI

Initially, during programming (conceptual design), the project team assessed credits options and opportunities that were defined by the site conditions and parameters. A preliminary techno-economic Return of Investment (ROI) study encouraged the team to proceed (see Figure 6).

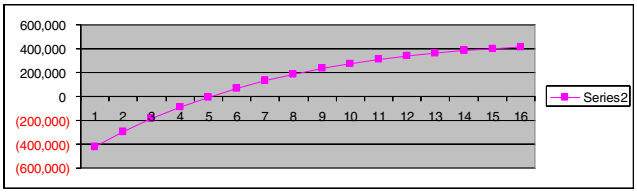

Figure 6: Model used to estimate ROI

This assessment determined LEED "Silver" as a feasible target. The processed preliminary study showed potential operational savings that could cover the estimated premium of less than one percent (<1%) in less than five years. These figures are in line with the works of Davis Langdon [15] and Gregory H. Katz [14]. Actual premium construction costs are still being tracked and will be published after occupancy. However, a few details and remarks can already be made:

1. The design cost did not increase due to the green building scope, relative to the company's regular benchmarked cost. The Green Building contract paragraph contained a clear expectation that the design to efficient systems and structures was a basic expected professional approach, while no unique system consultancy was required (i.e., renewable energy was not implemented).
2. The new role of the LEED Accredited Professional consultant added a reasonable additional cost required for documentation, meetings, and communication with the design consultants and the third-party audits.

3. Certification fees of $12,500 were based on the building size.
4. Premium construction costs are referenced to the standard new Intel constructed R&D building. In this project, we have estimated less than $600,000 in construction premiums. Cost distribution (see Figure 7) shows that most of the premium costs are due to energy improvements (40%) and indoor environmental quality improvements (35%). These improvements costs were estimated to save more than $200,000 a year.
5. The estimated savings did not include immeasurable impacts like productivity, absenteeism, public affairs gains, or "Green" marketing results.
6. ROI was based on total premium costs to total operational cost savings assuming that the LEED-Silver target is achieved. Four sub-projects above the Silver requirement were kept as a contingency to cover for potential missed targets.

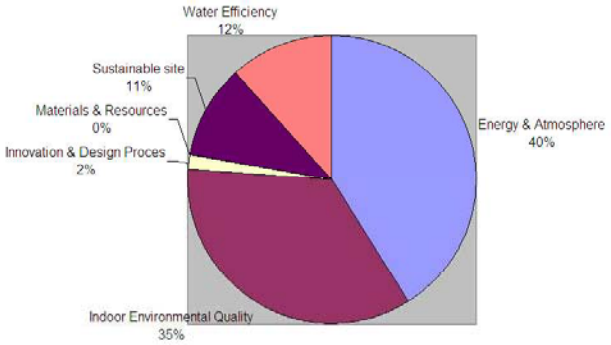

Figure 7: Project premium's cost distribution

Based on the feasibility study's positive results followed by corporate approval, the Public Affairs group got the green light to release information on the initiative internally and also externally to the public.

Intel's Social Responsibility Results

Intel's commitment to corporate responsibility, together with its positive media exposure and support from public officials, lead to numerous invitations from organizers of the foremost corporate and social responsibility conferences in Israel.

The "Green Building" was a great platform to promote Intel's relations with key government officials and leaders. Israeli Prime Minister Ehud Olmert, who laid the cornerstone for the "Green Building" during the ground-breaking ceremony, expressed deep appreciation for Intel and called on other companies to follow in Intel's footsteps by cultivating and advancing environmental awareness.

Extensive newspaper coverage of the establishment of the "Green Building," including features in Israel's largest daily newspaper and leading financial publications, have described Intel as a leader in the field of corporate responsibility and care for the environment. At a leading and well communicated Social Responsibility conference in Israel (Maala), which highlighted local corporations' activities, Intel Israel's CEO delivered the keynote address to the leaders in this field and described Intel's commitment to the global environment. These publications created a baseline that triggered Intel's supportive involvement in the local sustainable construction industry.

Greening Israel Construction Industry

As a way to better sense the local industry capabilities, the project team started to participate in local conferences, starting as members of the audience. Driven by the vast exposure to the public of the Green Building, the requests to participate as presenters were only a matter of time. The project has been presented at several events to architects, engineers, enterprises, and building owners. The project engineer participated as an Intel member in the local Green Building Association, bringing with him a unique experience in leading an active Green Building project. The project was volunteered as a pilot to establish the new Israeli Standard SI-5281 just issued at that time. Today Intel is mostly appreciated for its representatives' activities within the association's framework.

Green materials suppliers were exposed for the first time to new business opportunities. The project procurement strategy to create competitive bids required more than one supplier for each tender. In the case of paints, only one local qualified supplier was valid. A second well-known manufacturer was disqualified because he did not have the required formal document to prove his product's sustainability. Driven by business considerations, the manufacturer sent his samples to a certified known test lab out of the country, certified his product, participated in the bidding with the appropriate documentation, and finally was awarded the contract on account of his lower cost. In other words, a win-win situation.

The green building will have a modest visitor center to educate the public about the building's sustainability. The building is still under construction, but visitors and groups are already being hosted in the adjacent existing campus and attend lectures on the green features, ideas, and opportunities. Among the hosted groups are construction managers, students, journalists, hi-tech operation managers, and more. The stated team's mission to support and green the local industry begun months before the project was even completed and certified, through

community involvement and by exposing the challenges of making the R&D building green.

Intel R&D Green Building Design Features

The following are the implemented strategies to achieve a LEED-Silver rating:

- The building includes a data center, a high-energy consumer. The building's heating system is based on heat recovery from the DC water-cooled condensers. Recovered heat is also used for preheating of the domestic hot water supply to the kitchen and gym showers. There is no conventional fuel-powered boiler heating system.
- Sufficient daylight and enhanced ventilation increase employee satisfaction, productivity and health, key factors in reduced absenteeism.
- The façade includes reflective panels to maximize daylight penetration, providing natural lighting levels to more that 75% of the occupied areas (see Figure 8).

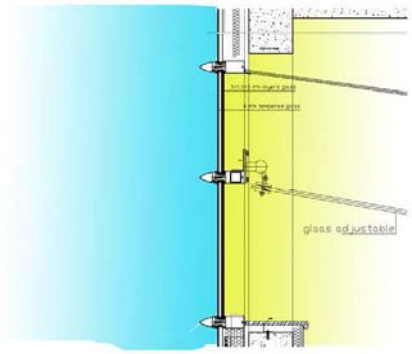

Figure 8: Internal reflective panel used to maximize daylight penetration

- Indigenous plants and controlled computerized irrigation will reduce water usage by 55% at mid-summer.
- A 31% reduction of water usage and related sewage will be achieved by efficient urinals, dual flush water closets, and low-flow faucets and showers.
- The gardens will be watered by an automatic irrigation system using recycled (condensate) water. A HVAC condensate water drain, collected at a rate of 5300 gallon/day during summer, is sufficient as the sole source of irrigation (no potable water is required).
- Lighting fixtures' controlled system enables each fixture to be dimmed separately (DALI com protocol).
- The quantity of fresh air in the offices is controlled by a CO2 monitoring system to save cost and improve ventilation.
- As a prerequisite, a recycling program for non-hazardous waste was implemented during the building's life-cycle.
- Two symmetric patios allow daylight penetration deep into the office at the three upper floors, combined with automatic daylight controlled lighting (Figure 9).

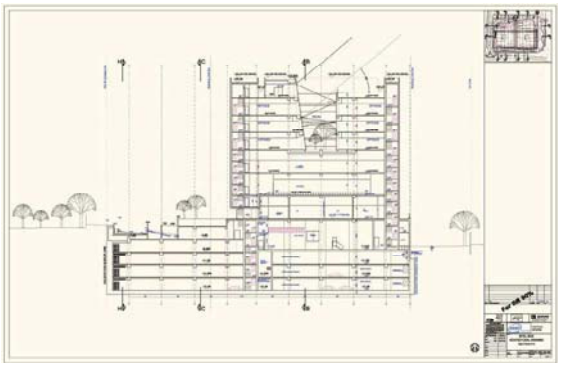

Figure 9: Building section showing patio and north tilted reflecting surface

- Individual light and thermal comfort is controlled via a Building Management System.
- Each worker will be allowed to adjust his ambient air comfort and overhead lighting from his or her personal computer.

Figure 10: Green roof and patio aperture

- Heat islands are reduced by planting a "Green Roof," and by using reflecting materials, shades, and landscaping. The "Green Roof" (Figure 10) also provides thermal insulation to the highest story (Level 5).
- The building site was originally a parking lot so there is no damage to natural assets.
- The site is located at a walking distance from main transportation stations, such as a train station, a central bus station, and a city bus (Figure 11).

Figure 11: The building position (in blue) relative to the original parking lot, walking distance from main transportation

- Motors and chillers were chosen to save on energy consumption (using Variable Speed Drives and high-efficiency motors).
- Bicycle Storage (for 40 units) with showers and changing rooms for the bicyclists' benefit are located in the underground parking area.
- The best environmental measures will be used in the construction to reduce dust, wheel tracks' wash and storm water runoff (Figure 12).

Figure 12: Storm water run-off prevention during construction

- The quantity of storm water runoff post-development has been kept less than that of pre-development.
- The paving of internal roads and pavements was offset by developing the adjacent public area with plants, while 50% of the building footprint is covered by "the Green Roof."
- A glazing surface area (low-e type) has been reduced on East/South/West facades (Figure 13).
- Only low-emitting (low-VOC) paints, adhesives, and sealers are applied.
- A carpool policy is implemented.
- A measurement and verification plan will be implemented during the first year of operation. Lessons learned will be available for other projects.

Figure 13: Building illustration, South-West view

- 10%-20% of all supplies will be made up of recycled material (a fraction of all supplied material costs).
- Forest Stewardship Council certified wood will make up 50% of the wood used in the building (doors, finishes, furniture).
- Construction waste (Figure 14) is recycled or reused (parking lot asphalt is reused on site; soil and aggregates were reused at adjacent construction site: steel, wood, cardboard, organic materials, etc. to meet the target of keeping 50% of the construction waste from having to be disposed of).
- Neighborhood density: the building is built within a dense neighborhood where the built area to surface exceeds 60ksf/acre (15ksf/dunam) within a radius of 990 ft.

Figure 14: Rock waste crushed on site and recycled

- The lobby includes a display area ("visitor center") and interactive presentations to educate visitors and employees about sustainable construction and the building's features.

- An enhanced commissioning plan is provided to ensure the building complies with the design intent, includes the strategic features, and is built to full owner satisfaction.

SUMMARY

Key Differences in Green Building Use/Sustain

Greening an R&D building has many functional, economical, environmental, and social advantages and is proven to be a logical construction strategy in today's world. The utilization of used materials diverts much construction waste from landfills. The modular design and use of recycled materials wherever possible, supports fast and low-cost modifications and retrofits down the road. The Green Building indoor environment enhances employee productivity, thereby retaining employees, reducing absenteeism, and enhancing researchers' concentration, which all results in gains for the corporation. Implementing recycling strategies during the life of the building helps to reduce waste and the cost of resources. Further, this reuse, recycle mindset becomes habitual for employees and is carried over into their after-work lives. The building site and planning maximizes natural light penetration thereby reducing operations and sustaining costs, and thus increasing corporate profits. Saving energy reduces CO2 emissions and subsequently global warming. Water-saving strategies reduce depletion of a valuable resource and reduce overloads in municipal waste systems.

Last but not least, Intel Green Building boosted employees pride and morale. The employees are proud to be the leaders in promoting Intel's corporate responsibility initiatives.

Leveraging the Success for the Future

The IDC9 building is not only the first registered Intel project to be certified as a "Green Building" project; it is also a role model for the Israeli construction industry. The challenge now is to ensure that this "green" mindset will continue to grow and will become part of Intel's construction standard as well as being implemented in the Israeli construction industry.

Intel has registered another two of its buildings to be certified by LEED-EB, the rating system for existing buildings. The lessons learned from the R&D project will help Intel with these two projects.

Motivating Innovation Initiatives

It was a challenge to drive innovative strategies in three broad areas: transform a structured corporation (Intel) engineering business process to include an innovative "Green" component and get consensus on this new mindset; succeed in getting the Israeli national construction industry to implement a "Green" mindset in how they conduct business; and to be audited and awarded certification by third-party experts.

In order to encourage continued employee engagement we are planning an internal campaign for employees to strengthen the connection between the "Green Building" and out-of-the-box environmental thinking.

The key message from the authors is to motivate: we now know it is possible to execute a local and innovative initiative successfully, even in global corporations, once it supports clearly-stated corporate values.

ACKNOWLEDGMENTS

The authors thank the people who supported us in transferring our huge collection of material into one readable paper:

Amir Roggel for his masterful help with the writing of this technical paper.

Ted Reichelt for his encouraging mentoring.

Alan Schijf for his focused directions.

Special thanks to Marian Lacey for her guidance and patience.

REFERENCES

[1] Sheldon Kamieniecki, editor, Environmental Politics in the International Arena: Movements, Parties, Organizations, and Policy, Albany: State University of New York Press, 1993.

[2] "Our common future," World Commission on Environment and Development (WCED), 1987.

[3] United Nations, "1987 Report of the World Commission on Environment and Development," *General Assembly Resolution 42/187*, 11 December 1987.

[4] Esty, D. C., Porter, M. E., "Industrial Ecology and Competitiveness: Strategic Implications for the Firm," *Journal of Industrial Ecology,* Winter 1998, Vol. 2, No. 1: 35–43.

[5] J. Elkington, Cannibals with Forks: the triple bottom line of 21st century business, Oxford, Capstone Publishing Limited, 1998.

[6] S. Sagawa and E. Segal, "Common interest, common good: creating value through business and social sector partnerships," *California Management Review 42* (2000) (2), pps. 105–122.

[7] http://www.usgbc.org

[8] "Israel Environment Bulletin," Vol. 31, Oct. 2006.

[9] Alex Wilson, "Making the Case for Green Building," *EBN,* Volume 14, Number 4, April 2005.

[10] "Guidelines for Environmentally-Responsible Building," *Israel Ministry of Environmental Protection*, updated 15/10/2006.

[11] Alexis Karolides, Green Building: Project Planning and Cost Estimating, 2nd Edition 2006, Reed Construction Data Inc., Kingston, MA, p. 3.

[12] Stephens, Jeff, *Ecostructure Magazine*, November/December 2004.

[13] Lisa Fay Matthiessen, Peter Morris, "Costing Green, a comprehensive cost data base and budgeting methodology," Davis Langdon (Financial Consulting Firm), July 2004.

[14] Gregory H. Katz, *Green Building Costs and Financial Benefits,* Massachusetts Technology Collaborative, USA 2003.

[15] Lisa Fay Matthiessen, Peter Morris, "*2007 The Cost of Green Revisited,"* Davis Langdon (Financial Consulting Firm), 2007.

[16] Steven Morton, "Business Case for Green Design," *Building Operating Management*, November 2002.

[17] Shelly Paz, Alissa Gold, *Jerusalem Post Daily Journal*, October 7, 2007.

[18] "Intel Corporate Social Responsibility Report," 2007.

AUTHORS' BIOGRAPHIES

Dori Hershgal is with EMEA Corporate Services, managing Intel's first Green Building engineering and design. He has held Project and Engineering Management positions and currently is developing the sustainable construction engineering concept in EMEA. Dori is a LEED Accredited Professional (LEED-AP) and an active member of the Israel Green Buildings Association, proliferating sustainable construction in the local market through lectures and conferences. Dori has an M.Sc. degree in Fluid Mechanics & Heat Transfer from Tel Aviv University and has a specialty in Wasted Heat Utilization for Refrigeration. His e-mail is dori.hershgal at intel.com.

Marc Denner is with EMEA Corporate Services, the Engineering Department Manager for Israel and Business Enterprise Group (BEG). He was part of the initial team that started the IDC9 "Green Building" concept. Marc holds a B.Sc. degree in Chemical Engineering and also holds an MBA. His e-mail is marc.denner at intel.com.

Ayelet Harush is leading environmental programs at the Intel Israel site. Her main interest is water management and design for environment. Permitting and regulatory activities with local agencies are her main focus right now, supporting the business changes and growth at the Israeli site. Ayelet has an M.Sc. degree in Environmental Science from the Jerusalem Hebrew University. Her e-mail is ayelet.harush at intel.com.

Revital Bitan is an Intel Program Manager. She developed the Corporate Social Responsibility (CSR) strategy and worked with local CSR organizations to spearhead the CSR concept. She has driven numerous local community environmental projects and national education programs to bridge the digital divide. She participated in Harvard University's program "Women and Power—Leadership in the New World." Revital holds a B.A. degree in Social Science from the Open University, Israel. Her e-mail is Revital.Bitan at intel.com.

Green Homeowners as Lead Adopters: Sustainable Living and Green Computing

Jay Hasbrouck, Digital Home Group, Intel Corporation
Allison Woodruff, Intel Research Berkeley, Intel Corporation

Index words: ethnography, green homes, green consumerism, green technology, sustainability

ABSTRACT

In this paper, we explore opportunities for green computing, drawing on our ethnographic study of 35 green households in the United States. We begin by discussing recent shifts in Western environmentalism and the growth of green consumerism. We then explore green social networking processes and how these relate to perceptions of corporations, openness to new technological solutions, and our participants' perceptions of computing devices. We conclude with a framework of strategies for green computing technologies that is based on our analysis.

INTRODUCTION

The dramatic importance of environmental issues has permeated the awareness of consumers, regulatory bodies, OEMs, and others over the past few years. This awareness of environmental issues and global climate change has reached an inflection point with the potential to impact Intel's products, brand, and operations. Intel has already taken significant actions such as reducing emissions from its Fabs, playing a formative role in initiatives such as the Green Grid [1] and Climate Savers [2], and pursuing energy-efficient products. Many additional opportunities lie ahead for "green" (environmentally conscious) computing.

In this paper, we discuss emerging trends in environmental action and how these relate to green computing technologies. Our analysis draws on an ethnographic study of 35 households in the United States, as well as expert interviews and secondary sources. Although green home owners are clearly a distinct and unique population, in many ways they are at the forefront of personal environmental action, and they offer valuable insights into how consumers perceive technologies in terms of sustainability and green values. We explored questions such as our participants' motivations for environmentalism, their openness to new technological solutions from corporations, and their perceptions of computing devices in relation to green values.

METHOD

In addition to surveys, focus groups, and other market research, Intel gains insight into the daily practices and everyday behaviors of consumers through ethnographic research. This typically involves visiting with people in their homes, workplaces, and other locations with the objective of developing a deep understanding of the values, beliefs, and worldviews that inform their perspective of technology in their lives. Methods usually include informal interviews, participatory exercises (such as collaboratively mapping the home and asking how participants interact with various areas of it; or working with an image-association card deck), home tours, and participant observation.

The ethnographic data from this research (video footage, photographs, interview transcripts, artifacts such as annotated maps, and other relevant materials) are then analyzed, and the findings are used to identify opportunities for our computing platforms to enable experiences that consumers value. This process typically involves refining a series of user experience models (or usages) that are designed to best tailor technological capabilities to the lives and priorities of consumers.

For this study, we visited households whose occupants made significant changes to their homes and behaviors in order to be more environmentally responsible. From simple changes that reduce water consumption to the installation of photovoltaic systems or the use of recycled and salvaged materials, the home has become the vehicle for expressing environmental values for many people in the United States and beyond. This study set out to gain a deep understanding of the domestic practices, values, and priorities that shape these decisions and to determine the role technology plays within them.

We chose to locate this phase of the study (October 2006 to February 2007) in the United States, in large part because the rate of energy consumption in proportion to population makes the U.S. a clear candidate for the development of new technologies that help improve energy efficiency. As the world's largest single energy consumer, the United States is also the largest emitter of carbon dioxide, accounting for 23% of energy-related carbon emissions worldwide [3]. Much of this can be attributed to home energy usage, which can be broken down accordingly: heating/cooling (44%), lighting/appliances (33%), water heating (14%), and refrigeration (9%) [4].

As for selecting regions for this study, www.sustainlane.us recently published (2006) a sustainability ranking of cities, in which Portland, OR; San Francisco, CA; and Seattle, WA consistently rank at the top in all categories (renewable energy use, alternative fuel use, food availability, tap water quality, air quality, walkability, park space, and roadway congestion). In terms of progressive energy and climate change policy, Portland, San Francisco, and Seattle were tied for first place, followed by Los Angeles and Albuquerque. Our selection of metro areas for this study roughly follows this ranking:

- Portland, OR (5 households)
- Bend, OR (3 households)
- San Francisco Bay Area, CA (7 households)
- Los Angeles, CA (8 households)
- Santa Barbara, CA (2 households)
- Albuquerque, NM (7 households)
- Taos, NM (3 households)

In these places, we sought participants who expressed commitment to improving the environment by building or retrofitting home functions that minimize negative (and enhance positive) effects on the ecosystem in which the home is situated. In addition, we sought participants who represented a range of diverse living conditions and perspectives on sustainable domestic praxis.

Our analysis is informed by data from these participants, additional interviews we conducted with domain experts, and a critical examination of secondary sources about environmentalism.

SOCIAL CHANGES IN ENVIRONMENTALISM

To set the stage for our research, it is important to consider the broader context of environmentalism. Beginning in the late 1980s and early 1990s [5, 6], environmentalism in the Western world began to change from a counterculture movement with a focus on saving endangered species and preserving the integrity of their ecosystems to a more holistic approach that began to integrate social, economic, and political considerations into ecology. This shift was accompanied by changes in terminology as well and is probably best represented by the increasingly common usage of the term "sustainability."

In keeping with this more holistic perspective, our participants' motivations for taking environmental action were wide ranging: they included everything from religious values to a desire to save on home energy bills to new forms of patriotism that focus on reducing dependence on foreign oil. In a recent report, Euromonitor [7] cites four primary drivers for sustainable behavior that may explain some of the rapid growth of interest and action in this area, especially among those who have not traditionally self-identified as environmentalists. These include motivators such as saving money by saving energy, a desire for a healthy home, freedom from guilt associated with purchasing environmentally-damaging products from ethically compromised companies, and an appreciation for "unspoiled nature." Janet's story below (in her own words) is evidence of a wide range of motivations, and it illustrates how reasons to take environmental action are not necessarily altruistic, or even always "green" (note that participants' names have been changed in order to protect anonymity):

> *Janet*: Motherhood was really what got me on this. There is nothing like having children to plug you into the future and it so happened that when I came to California in '66 I felt as though I was coming home although I had never been here before. And I think one of the reasons that it felt that way is because I was able to reconnect to nature here in a way that I hadn't on the east coast for many years. And so I joined many environmental organizations and their newsletters and magazines provided me with my environmental education. And so in the 70s I was already aware of some of the environmental problems that were kind of lurking in the wings waiting to make their entrance on stage before mainstream media picked it up at all.
>
> And so as I learned more and more about that, I became really concerned about my children's future and the future of all of life on Earth and that is what started me on my quest was, you know, I came to realize that it is our modern urban lifestyle except for the military which does do a lot of environmental junk damage. It is true. But other than the military, one can trace back virtually all of

> our environmental problems to our modern urban lifestyle which is that was me. You know. I was living it. And so it was both empowering and overwhelming. I decided to go with the empowering.
>
> And so I started to just look into, well, how can I change my lifestyle so that it is not so destructive to the environment but is less destructive and maybe even in some ways can help to support three generations of a healthy environment within an urban setting. That was the big challenge of it because the resources that I had for these systems were magazines like *Mother Earth News* and *Organic Gardening* that were targeting really rural folk.
>
> But as I read them and then I read it in my environmental magazines about the problems, I began to put the two together and realize well here I have a front yard and a back yard. There are a lot of these strategies that are described in…*Organic Gardening* magazine and *Mother Earth News* that I can implement right here just having a front and back yard. So that's what started me. I have a perfect little mini cosmos here to use as my lab to see if I can create a lifestyle that doesn't destroy the environment, you know. So that's what started me off. My children and my environment. So I don't have any formal training, although when I started my front yard I did take some courses in horticulture and native plants and landscape design and things like that to give me some skills that I didn't have.

As these more integrated views began to spread, some previously held positions within environmentalism began to be questioned and challenged. In particular, a notable shift away from critiques of technology began to take place, and a new optimism arose that technology and design could provide solutions to environmental challenges. Wikipedia describes this phenomenon (sometimes labeled "bright green environmentalism") accordingly:

> "Bright green environmentalism aims for a society that relies on new technology and improved design to achieve gains in ecological sustainability without reducing (indeed, increasing) the potential for economic growth. Its proponents tend to be particularly enthusiastic about green energy, hybrid automobiles, efficient manufacturing systems and bio and nanotechnologies, [and] are supportive of dense urban settlements. 'One-planet living' is a frequently heard buzz-phrase.
>
> Their ideas can be contrasted with what they consider traditional environmentalism: pessimistic, return-to-primitivism, unattractive, 'dark green' ideas that depend on a reduction in human numbers or a relinquishment of technology to reduce humanity's impact on the Earth's ecosphere."

Together, bright green environmentalism and the increasingly holistic view of the environment have set the stage for an expanded understanding of environmentalism in the Western world. It is here that the media—sparked in large part by rising energy costs, an unstable oil supply associated with conflict in the Middle East, major climate events such as Hurricane Katrina, and Al Gore's clear depiction of the causes and consequences of global climate change in *An Inconvenient Truth*—gripped the public imagination with an appeal far broader than 1960s counterculture-based environmentalism. As a consequence, people began to identify ways in which they personally could take environmental action without necessarily adhering strictly to the traditional strategies of protest and political action. Some called it the birth of green consumerism.

GROWTH OF GREEN CONSUMERISM

Many of the drivers for environmental action seem to extrapolate well to the general population in the developed world, and there are some clear indicators that green consumerism as a new form of environmental action is taking root in substantive ways. In terms of economics, it is hard to ignore the shifts in policy and practice that have taken place in corporate America over the past five years or so. For example, Wal-Mart's "Sustainability 360" program now evaluates suppliers based on progress toward sustainability benchmarks in areas such as energy efficiency, durability, package size, materials that help reduce the use of hazardous substances, the ability to upgrade, and recycling possibilities. Likewise, Home Depot's Eco Options Program highlights "green" products for consumers who are looking for sustainable options in home projects. Further evidence can be found in the fields of architecture and design, where green building and design standards are rapidly becoming the norm. The U.S. Green Building Council has recently extended its Leadership in Energy and Environmental Design (LEED) standards into residential structures, and some local real estate Web sites (see Portland's MLS [8]) now include green features such as solar electricity or water reclamation as value-added features of a home. In design, a quick glimpse at popular blogs such as www.core77.com will illustrate how sustainability has come to the forefront in that profession. Similarly, interest in green entrepreneurship is spreading across the country (and the world), as evidenced by the rapidly growing popularity of green networking efforts such as Green Drinks gatherings

(www.greendrinks.org), and the new flush of venture capital investment in alternative energy in the Silicon Valley [9, 10].

These corporate and entrepreneurial efforts offer potential benefits such as higher product margins, increased market share, and positive consumer perception of a company's corporate social responsibility. According to a 2006 Mintel Research study [11], the green marketplace is now estimated at somewhere between $300-500 billion a year. The same study showed that there are approximately 35 million Americans who regularly buy green products, and that 77% of consumers changed their purchasing habits due to a company's green image. Marketing statistics from many different industries support this. Green homes, for example, are estimated to cost between 2 to 5% more to construct, but are valued at 10 to 15% more in the marketplace [12]. Likewise, organic dairy products are priced typically 15 to 20% more than conventional ones, and organic meats are often priced two to three times more than traditional meat [13]. In terms of technology in particular, we see increasing consumer interest in energy conservation, as well as growing expectations that consumer technologies should enable consumers' personal wellness in ways that are both sustainable and ethical.

While these statistics are certainly promising, it is also important to note that consumer valuations of technology also include a history of expectations that inform perceptions of technology products in unique ways. For example, a recent study by Saphores *et al*. found that most consumers "are willing to pay only a 1% premium for 'greener' CEDs" [14]. Part of this may be attributed to the fact that "consumers expect manufacturers to innovate to make their products more environmentally friendly without significantly increasing their prices as they may be accustomed to the falling prices and constant progress that characterizes electronics manufacturing" [14]. Even so, the study concludes that "[w]orldwide, the trend is toward environmentally friendly electronics manufacturing," and green attributes can play a role in consumer preference for one device over another even when they do not command a premium price.

Participants in our study expressed optimism that corporations and other organizations could provide valuable solutions or environmentally friendly options, but they were also deeply concerned that the growing awareness and popularity of green values would increasingly result in unethical attempts to overstate claims of environmental responsibility by corporations and other businesses. Participants were highly sensitized to such "green-washing," and were strongly negative about products that were positioned as green but were in fact only partially or superficially green.

GREEN NETWORKING

Most social scientists are likely to agree that all of the attention, interest, excitement, and investment around environmentalism does not necessarily constitute a community. Some posit that community can be approached as a value [15], comprised of qualities such as solidarity, commitment, mutuality, reciprocity, and trust [16, 17]. More functionalist approaches see 'community' as constituent of two related components: 1) that the members of a group have something in common with each other; and 2) that the thing held in common distinguishes them in a significant way from the members of other possible groups [18]. Using these definitions as a guide, the concept of community necessarily involves an exclusionary act, in which the benefits of belonging to a particular group are denied in some way to non-members. In addition to this exclusionary position, shared expectations about the way people should behave 'in-group'—and how well integrated these behaviors are within conceptions of identity for its members—are also critical to understanding whether or not a group understands itself as a community.

For the participants in our study, it is difficult to argue that they collectively constitute a community. Beginning with Cohen's exclusionary principle of community, some participants in our study pointed to varying levels of commitment to sustainable building and living in which they tended to position themselves. This did not typically translate to exclusionary behavior, however. For example, Edward, an architect in the Bay area, is quick to point out that although he feels many other architects fall short in green building, he doesn't want to "disqualify them."

> *Allison*: Are there any green architects in—in the area that you [interact with] or…?
>
> *Edward*: Yeah, there are a few.
>
> *Allison*: Yeah?
>
> *Edward*: Yeah. But I have to say—I mean, you know, everybody tries to be green and I don't want to, you know, um disqualify them or anything. I mean, I guess everybody's doing their best, right?

This inclusive sentiment is often accompanied by a remarkable enthusiasm for sharing information and different perspectives on sustainable living.

Whether intentional or not, our participants' drive and enthusiasm to share knowledge and (in some cases) facilitate the somewhat viral spread of interest in sustainable living into the 'mainstream' would seem to indicate an inherent resistance to forming a 'community'

as defined by the social sciences. This could be interpreted as an effort to avoid the ghetto-ization of sustainable living and to steer 'earth-friendly' practices in directions that were out of reach for 1960s counterculture-based environmentalism. In other words, if mainstreaming is the goal, it doesn't do you much good to form a clique.

Our research revealed a number of patterns within which participants tended to communicate about green home construction, remodeling, and related forms of sustainable living. These patterns typically begin with processes of learning by networking with others who have already taken similar action. For instance, Jerry and Kylie cite how a solar conference proved to be a particularly informative venue for them:

> *Allison*: So, where do you find out how stuff works?
>
> *Jerry*: Lots and lots and lots and lots of reading and I think our biggest, most interesting eye opener in the past year was we went to a big solar fair called SolWest…
>
> In all, it was a great technology event as far as being able to sit down with the people and talk about problems and solving this and going to seminars and that kind of thing. It was a huge eye opener for why this technology is being accepted across the board. There were 280 or 300 visitors to SolWest. And like I told people at the solar home party after this whole thing was over, I said, "If you guys want to go somewhere and see an interesting cross-section, go to SolWest." There were like 280 visitors at SolWest and you could probably find 200 different reasons why people were interested in solar power.

Likewise, the Web tends to play a critical role in green networking. Here, Janet talks about how her non-profit environmental organization benefits from online communications:

> *Jay*: Does the Internet play an increasing role in the [non-profit's] networking?
>
> *Janet*: Well it does. Increasingly people come to the tour having found us on the net and then they tend…they will become members. Not everyone who comes on a tour joins, but a proportion of them do and that does build the network.

Instances like these—where synergy between online networks and traditional organizations builds momentum and reaches new audiences—are increasingly common. In fact, authors like Bill McKibben see the Internet as a key facilitator of a new kind of environmental action:

> "In days of yore, if you were concerned about, say, global warming, you might write a letter to your congressman. You might research the presidential candidates to figure out which one was most aggressive about climate change, and then you might mail him a check. But the chance to work together with people around the country on a common cause was mostly reserved for 'organizations'—for environmental groups, say, with big buildings in Washington, calendars and boards of directors.
>
> In the Internet Age, though, new models emerge…a certain kind of organizing no longer requires years of groundwork. It requires a good idea and a well-written e-mail…The fight against global warming requires all kinds of technology—solar panels and windmills, but also servers and routers" [19].

In addition to new forms of alliance-building, the Web also provides opportunities for participatory environmentalism in which blogs, forums, mailing lists, social networking sites, and other Web venues increasingly serve as platforms where individuals can personalize their own forms of environmentalism. The breadth and depth of environmentally-focused blogs alone (c.f., www.bestgreenblogs.com/) bears witness to the fact that there is a great deal of interest in varying forms of personal expression in this area.

Our research also revealed personalized forms of environmentalism in which the home itself is used as a vehicle for "green" expression. For many of our research participants, street protests or sit-ins aren't realistic options for expressing their concern about global climate change or other environmental issues. However, modifying a home to be more sustainable offers immediate personal gratification coupled with a tangible demonstration of environmental values. And, like other forms of expression, home modifications are also increasingly tied to the Web and the resources it offers.

These forms of individual expression align well with the historical tradition of North Americans' understanding of nature and environmentalism, and the relationship of these to North America's frontier-oriented history. These beliefs begin with the assumption that individual action (often enabled through networks) is the most direct way to affect change. Silveira puts it in terms of entitlement:

> "…the notion of Nature itself is laden with religious and moral overtones. From early on, American frontier ideology has equated the settling of the frontier with the rise of America's vigorous and independent democracy. Whether it be the right of a child to grow up without the increased

> risk of getting asthma from poor air quality, or the right of indigenous populations to fish in ancestral fishing grounds without increased risks of cancer from toxins in the fish, rights rhetoric can be used by environmentalists to both personalize and broaden the appeal of their message." [20]

Beyond Silveira's notion of environmentalism linked to entitlement, we found that our research participants also practice many other forms of expression that reflect a unique intersection between environmental priorities and North American values closely tied to the historical and cultural landscape of the United States. Perhaps most commonly, we witnessed a sense of optimism that is suggestive of the utopian narratives that describe the nation's settlement.

For example, our participants' critiques of the 'mainstream' often engage utopian imaginaries that are reminiscent of the nation's founding narratives, in which immigrants fleeing various forms of oppression from (historically European) cultural establishments find a place to start over, and in the process re-invent themselves. Although this latter narrative typically overlooks some of the tragedies that befell Native Americans in the process, it also includes an appreciation of the environment that is firmly rooted in the mythical and visual interpretations of the American West, where some of the country's most widely renowned early environmentalists (Emerson, Thoreau, Muir, Roosevelt, Leopold, etc.) found their inspiration.

PERCEPTIONS OF COMPUTING DEVICES

Having situated our research in terms of social shifts in environmentalism, green consumerism, the increasing importance of social networking, and the historical referents these phenomena exhibit, we now turn to a more specific consideration of technologies in the lives of our research participants. Many participants in this study expended tremendous effort and funds to behave in environmentally responsible ways and to make environmentally responsible purchasing decisions. They frequently performed sophisticated and complex analyses to determine their preferred course of action, often drawing on many facts and heuristics to make decisions about transportation, home energy use, water use, food purchases, etc. [21]. However, they had invested significantly less energy in considering the environmental impacts of computing devices and consumer electronics, and their analyses were much less sophisticated in this area relative to others.

> *Susan lives in affordable rental housing in the Los Angeles area. Susan recently purchased her first new car. Prior to the purchase, she and her friends did research and performed detailed calculations about the expense of a hybrid car versus a conventional gas-powered car. Susan and her friends ultimately concluded that, over the life of the car, it would cost Susan $2000 more to have a hybrid car rather than a conventional gas-powered car. Susan decided that it was worth it to her to spend $2000 extra to do what was right for the environment. She purchased the hybrid car and was very proud of her decision. We asked Susan what considerations will go into her next computer purchase. She said she usually just considers the speed of the computer—she wouldn't know how to think about the purchasing decision for a computer in an environmental way, she just doesn't have that knowledge.*

In some cases, participants appeared to draw simple analogies from other areas to try to reason about environmentally responsible behaviors and purchasing decisions for computing devices and consumer electronics. For example, one participant spoke at length about packaging and the amount of styrofoam in the packing box as a dominant concern in her purchasing decisions for stereos and other similar devices. In addition, participants often had a sense that they should recycle electronic products, although their understanding of the recycling process and potential environmental impacts of different recycling options seemed extremely limited when compared with the complex issues raised in publications such as Grossman's *High Tech Trash* [22]. Even participants who were quite technical did not have as much knowledge as one might expect about issues such as the recycling of electronic products.

Participants also often had a strong inclination to keep computers turned off and to have as few electronic devices as possible (in keeping with conservationist heuristics). Participants often optimized energy use of their computers by shutting them down entirely when they were going to be away from their computers for even short periods of time. Some participants were very concerned about "phantom loads" (energy that is consumed by devices when they are off or in stand-by mode). A common strategy for addressing phantom loads was to plug devices into a surge protector and then cut off power to the devices when they were not in use by turning off the surge protector. This suggests some opportunities for device design—features such as fast wake-up or other energy management functions would plainly be appealing to this audience, although careful consideration would have to be given to the energy consumption of the device while in stand-by mode.

Many participants evidenced a pragmatic attitude toward computing devices, expressing that computers had poor environmental qualities but that they were necessary or that their use of the computer was an overall benefit to the environment because they were using it to promote green causes.

> *Cory*: We both use computers every day in our jobs…We use them as tools. And so even though we know they're toxic, we know there's all sorts of things in them that aren't sustainable, they help us do our work, which is promoting sustainability, so we move on.

In addition to being used for green-related work such as designing green buildings or running a green business, computers were very commonly used as a resource for searching the Internet for information about green living and green purchasing.

> *Adam*: [The water reclamation system] is all automated, it is all mechanically automated.
>
> *Jay*: How did you find out about that?
>
> *Adam*: Let's see, there was a Web site that I was researching—when I was researching the systems I came across this product. It's manufactured in…the Midwest and there is a lot more water harvesting in the agricultural region in Texas and so forth, and they use it more for irrigation, but they come up with these little devices. It's fun to apply it here.
>
> *Christopher*: For one thing, in and of itself [the Web is] more sustainable than having people mail you stuff all the time. Catalogs and whatever… And I view it as being more current. You can find the latest information. You can find more impartial information. You don't have to wait for a brochure from a company who's trying to sell you something to send you this brochure. You can get more original research off the Web if you search carefully. And you can look up things about, you know, lifecycle analysis on different materials and so on, like that.

Overall, while computers were seen as a useful tool for green work-related activities, the purchasing choice and day-to-day operation of the computing devices themselves was seen as a somewhat marginal issue from a green perspective. One possible explanation as to why participants had a relative lack of awareness or interest in this area is that the historic counterculture roots of environmentalism did not embrace electronic devices. Therefore, early proponents of environmentally responsible activities may not have reflected in detail on electronic devices, in contrast to their more sophisticated consideration and dissemination of information in regard to (for example) food. As awareness of environmental issues has reached a wider audience, this early anti-technology bias may have left a gap in knowledge and education. We believe this gap may present an opportunity for education and (re)positioning of computing devices and electronics relative to environmental concerns.

STRATEGIES FOR GREEN COMPUTING

Technology is understood to have a complex relationship with environmental issues. As discussed above, while traditional environmentalism argued for a reduction in technology, more recent environmental perspectives tend to recognize positive opportunities for simultaneous advances in environmental responsibility, quality of life, and economic growth. Accordingly, while computing technology can be viewed as "part of the problem" via issues such as e-waste and energy consumption, computing technology can simultaneously be viewed as "part of the solution" by enabling better use of resources in a wide range of domains (e.g., telecommuting, optimized route planning for fleets of trucks, smart buildings). These two perspectives correlate with two broad categories of an overall green strategy for computing technology:

Minimize: Minimize computing technology's contribution to the problem in both products and operations, e.g., through novel energy-efficient products, product design for reuse, reduced use of materials, recycling, etc.

Enable: Maximize computing technology's contribution to the solution by enabling green applications of computing technology, e.g., green design, smart buildings, dematerialization, etc.

Inspired by the actions and attitudes of our participants as well as by our review of secondary research materials, we have developed a framework of opportunities for green actions in these categories. This framework appears in Table 1.

Table 1: Strategies to (1) *Minimize* harmful environmental impacts of products and operations, and (2) *Enable* green behavior through computing

Minimize	*Reduce*	*Optimize X (optimizations to use strictly less of X, or increase the efficiency of X)* Examples: Reduce emissions from manufacturing, turn off lights, use energy efficient devices
	Reconstitute	*Replace X with Y (generally game-changing or more radical innovations)* Examples: solar panels, biodegradable packaging, hybrid cars, eco-friendly household products
	Reuse	*Reuse X as X* Examples: extend useful life of objects
	Recycle	*Extract materials from X* Examples: extract gold, e-recycling, composting
	Reparation	*Compensate for X by also doing Y* Examples: carbon offsets
	Re-think	*Reconsider or increase awareness (indirectly prompt Reduce/Reconstitute/Reuse/Recycle/Reparation)* Examples: make new policy, education, eco-labeling, Energy Star
Enable	*Design and Invent*	*Use computing technology to design and/or invent green objects and green processes* Examples: design green buildings, invent sources of renewable energy, design fuel-efficient aircraft
	Optimize and Automate	*Use computing technology to calculate and/or execute optimizations in real-world systems* Examples: efficient route planning for delivery trucks, smart buildings, smart agriculture, smart appliances
	Minimize Miles	*Use computing technology to support remote activities to minimize travel of people and goods* Examples: telecommuting, teleconferencing, online purchasing
	Monitor and Model	*Use computing technology to monitor and model the state of the environment* Examples: environmental quality sensing, climate modeling
	Operate and Sell	*Use computing technology to sell green goods or services, or run a green non-profit or business* Examples: solar panel installers, carbon offset vendors
	Educate and Persuade	*Use computing technology to promote green behavior and share information locally and globally* Examples: green blogs, citizen action tools, ambient displays of energy use, personal digital coach for green behaviors

First, Table 1 lists green strategies to minimize harmful outputs and the use of resources: Reduce, Reconstitute, Reuse, Recycle, Reparation, and Re-think. These strategies underlie a wide range of environmentally responsible behaviors. To illustrate the generality of these principles we include examples of both actions that might be taken by a technology company (e.g., reducing emissions) as well as actions that might be taken by people in their daily lives (e.g., turning off the lights).

Second, Table 1 lists strategies to enable green processes, products, and actions: Design and Invent, Optimize and Automate, Minimize Miles, Monitor and Model, Operate and Sell, and Educate and Persuade. In the descriptions

and examples in the table, we focus in particular on how information and communication technologies can enable each of these strategies. Note that these Enable strategies implicitly enact Minimize strategies. For example, a green building can be designed to reduce energy use, reuse water, etc.

To date, the bright green perspective has tended to emphasize areas such as alternative energy sources and transportation. However, computing technology has much to offer in this arena as well. A comprehensive strategy for green computing will both Minimize and Enable, by making computing technologies that are themselves environmentally conscious and by deploying technologies to meet environmental challenges. Many of these directions are already being pursued, and many exciting opportunities lie ahead.

ACKNOWLEDGMENTS

We gratefully acknowledge our participants for their kindness in welcoming us into their homes. We thank Ryan Aipperspach, Paul Aoki, Sally Augustin, Eric Brewer, Shellie Collier, Herman D'Hooge, Sue Faulkner, Brooke Foucault, Diane Loviglio, Steven Russo-Schindler, and Christy Snyder for valuable comments and assistance.

REFERENCES

[1] http://www.thegreengrid.org

[2] http://www.climatesaverscomputing.org

[3] US Department of Energy/Energy Information Administration, *Emissions of Greenhouse Gases in the United States, 1995*, US Government Printing Office, Washington, D.C., pg. 8, 1996.

[4] US Department of Energy/Energy Information Administration, *Residential Energy Consumption Survey*. US Government Printing Office, Washington, D.C., 2001.

[5] Brundtland Commission, *Report of the World Commission on Environment and Development, General Assembly Resolution 42/187*, United Nations, 1987.

[6] McDonough, W., *The Hannover Principles: Design for Sustainability*, William McDonough Architects, Charlottesville, VA, 1992.

[7] Euromonitor International, "Sustainability: Its Impact on Global Consumption to 2010," Strategy Briefing, November 3, 2005.

[8] http://www.rmls.com

[9] Nauman, M., "Clean and Green Findings," *San Jose Mercury News*, September 27, 2007.

[10] Yi, M., "Alternative Energy Lighting It Up: Big Venture Capitalists Pony Up in Sizzling Market That May Get Boost from Prop 87," *San Francisco Chronicle*, November 4, 2006.

[11] Mintel Research, *Green Living–US*, September 2006.

[12] Chongchua, P., "Homeowners Pay More for 'Green' Homes," *Realty Times*, February 19, 2007.

[13] Organic Trade Association, *The Past, Present, and Future of the Organic Industry*, 2005.

[14] Saphores, J.-D., Nixon, H., Ogunseitan, O., and Shapiro, A., "California Households' Willingness to Pay for 'Green' Electronics," *Journal of Environmental Planning and Management*, Volume 50, Number 1, pps. 113–133, 2007.

[15] Frazer, E., *The Problem of Communitarian Politics: Unity and Conflict*, Oxford University Press, Oxford, 1999.

[16] Coleman, J., *Foundations of Social Theory*, Harvard University Press, Cambridge MA, 1990.

[17] Putnam, R., *Bowling Along: The Collapse and Revival of American Community*, Simon and Schuster, New York, 2000.

[18] Cohen, A., *The Symbolic Construction of Community*, Tavistock, London, 1985.

[19] McKibben, B., "The Power of the Click," *Los Angeles Times*, October 16, 2007.

[20] Silveira, S., "The American Environmental Movement: Surviving Through Diversity," *Boston College Environmental Affairs Law Review*, Winter 2001.

[21] Woodruff, A., Hasbrouck, J., and Augustin, S., "A Bright Green Perspective on Sustainable Choices," *Proceedings CHI 2008*, to appear.

[22] Grossman, E., *High Tech Trash: Digital Devices, Hidden Toxics, and Human Health*, Island Press, Washington, 2006.

AUTHORS' BIOGRAPHIES

Jay Hasbrouck is a Social Anthropologist in Domestic Designs and Technologies Research within the User Experience Group at Intel Corporation. Jay joined Intel in January of 2006 and has conducted fieldwork in Mexico, Egypt, Germany, South Korea, Brazil, and the United States. He holds a Ph.D. in Social Anthropology and an M.A. in Visual Anthropology, both from the University of

Southern California. His research interests include sustainability and the role of technology, social networks, visual culture, domesticity and cultural landscape. His e-mail is jay.hasbrouck at intel.com.

Allison Woodruff holds a Ph.D. in Computer Science from the University of California, Berkeley, an M.S. in Computer Science and an M.A. in Linguistics from the University of California, Davis, and a B.A. in English from California State University, Chico. Her research interests include sustainable technologies, domestic technologies, computer-mediated communication, and information visualization. She currently works at Intel Research Berkeley. Her e-mail is allison.woodruff at intel.com.

Novel Wastewater Reclamation Technology Meets Environmental and Business Challenges

John Currier, Technology and Manufacturing Group, Intel Corporation
Boris Eliosov, Technology and Manufacturing Group, Intel Corporation
Vyacheslav Libman, Technology and Manufacturing Group, Intel Corporation
Dan Enloe, Technology and Manufacturing Group, Intel Corporation
Don Crandall, Technology and Manufacturing Group, Intel Corporation

Index words: membrane bioreactor, environmental, wastewater, ultrapure water

ABSTRACT

In the relentless pursuit of Moore's Law, Intel's Technology Development cycle results in a new semiconductor production process every two years. As this continuous cycle of new products moves from conception to market, so do the innovative environmental solutions designed to meet Intel's philosophy of environmental excellence. Intel's success in design-for-environment (DFE) is widely recognized. In 2007 Intel was rated the top U.S. technology company to be named to the Dow Jones Sustainability Index—for the seventh straight year. In normal process technology development cycles, Intel's environmental systems are developed in parallel with the semiconductor manufacturing process. However, Intel is a global company and this proactive process sometimes needs adjusting to accommodate unique differences in local conditions.

Such a situation occurred recently for a new Intel High-Volume Manufacturing (HVM) facility that was under construction in an arid region. The technology development team faced a triple challenge of finding a technology solution that would meet water conservation needs, new environmental standards, and a compressed project schedule. The project site was being required to treat the entire wastewater stream of both new and existing factories to achieve a very high level of wastewater effluent quality and to meet water reuse and conservation priorities of the local authorities. Meeting the schedule while designing and implementing a system capable of achieving effluent quality targets that would also maximize water reuse became a paramount task, since the business permit for operating the Fab would not be issued without this infrastructure. With only half of the normal development time available and no prior Intel experience with the proposed wastewater treatment technology, a task force was formed from multiple business groups to achieve successful completion of the project.

In this paper we focus on the following aspects of the aforementioned situation:

1. How the team incorporated risk-taking with other Intel best-known-methods to address both schedule and technology challenges.
2. The intensive technical study that resulted in the definition and optimization of the technology.
3. Value Engineering that resulted in design and construction of this system at costs normally associated with systems that are an order-of-magnitude smaller.
4. Creating greater water recycle opportunities by treating wastewater to such high standards.
5. The additional benefits of being able to treat other waste streams that are currently segregated.
6. The future challenges and opportunities associated with completion of this project.

INTRODUCTION

Intel announced plans to construct a new manufacturing facility emphasizing an ambitious schedule for construction. As with all technology-based manufacturing scenarios, time-to-market was paramount, and even small delays could affect the technology's financial success.

Seamless technology transfer from development to manufacturing is a trademark of Intel. Being proactive in identifying areas of risk and resolving roadblocks has allowed Intel to extend Moore's Law for four decades. The process of developing technology at a dedicated site, then copying it (using the Intel *Copy Exactly* process) during transfer to other manufacturing sites, has allowed

Intel to be a truly virtual factory, resulting in multiple factories worldwide producing identical products with standardized processes. This philosophy also applies to the design of environmental systems and promotes proliferation of high environmental performance including water conservation.

In normal product development cycles, Intel's environmental engineers are allowed up to four years to identify problems and implement solutions. To meet future regulatory requirements and Intel's global environmental standards, process environmental goals are developed well in advance of the technology. Environmental requirements vary due to differences in geography, population density, industrial infrastructure, local regulations, and global initiatives. Despite these differences, Intel proliferates environmentally consistent and compliant technology worldwide, designed to meet requirements for the most sensitive or restrictive site. This approach allows flexible business expansion at any site with assurance of environmental compliance. When this process identifies a need for higher environmental performance, a technology selection process is followed to ensure the environmental systems match the technology roadmap. A simplified process for technology selection is outlined in Figure 1.

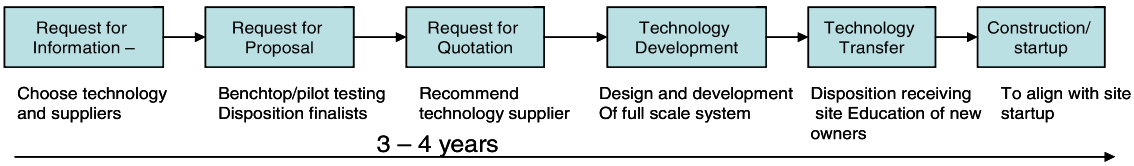

Figure 1: Technology selection process—steps normally do not overlap

This process does not easily accommodate unique site-specific conditions separate from the manufacturing process. Such a constraint occurred in this project when an unexpected change in the environmental requirements threatened the start-up date of the plant. In this case the construction schedule for the semiconductor plant outpaced the schedule for constructing the local municipal wastewater infrastructure. To avoid project delays, Intel developed its own solution for wastewater treatment and conservation. This coincided with the efforts of the local authorities to adopt more stringent wastewater quality standards intended to increase wastewater reclamation rates.

The new performance guidelines mandated use of effluent for unrestricted irrigation. Every drop of water discharged from the plant would ultimately be reused, either internally (by the plant) or externally. This means all water not evaporated and not used by building occupants could potentially be reused. Table 1 outlines these wastewater quality standards.

Table 1: Effluent quality requirements

Parameter	Quality standards for irrigation (mg/l)
COD	100
BOD	10
TSS	10
Ammonia	20
Total nitrogen	25

At this point, implementing technology selection, designing, and building an onsite treatment plant became a critical path for the startup of the new plant. Figure 2 shows the longer schedule for the municipal system compared with the schedule necessary for the onsite option.

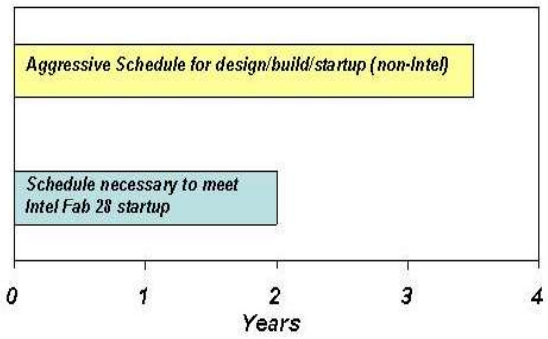

Figure 2: Technology implementation schedule

THE PROCESS

In contrast to the normal development process, the technology development effort would have to be modified and compressed by half to successfully meet the schedule.

The site team moved into action and partnered with the development groups responsible for transferring the new technology to this project. This taskforce was a critical part of bringing in the technology. Engineers across Intel were chosen both for their expertise and availability. Wastewater experts, construction experts, environmental professionals, analytical/laboratory staff, commodities experts, project managers, and other support groups participated in this task force. The team's mission was simple: enable the site to meet both the business permit requirements and the construction schedule. The taskforce first obtained top management support for this accelerated process to validate the business urgency and help ensure success.

SCHEDULE RULES

One of the first decisions made by the team was to define the project priorities. In this case, all parties agreed that *Schedule* was the primary driver, followed by performance, and finally budget. This hierarchy was extremely important in defining the team's business practices.

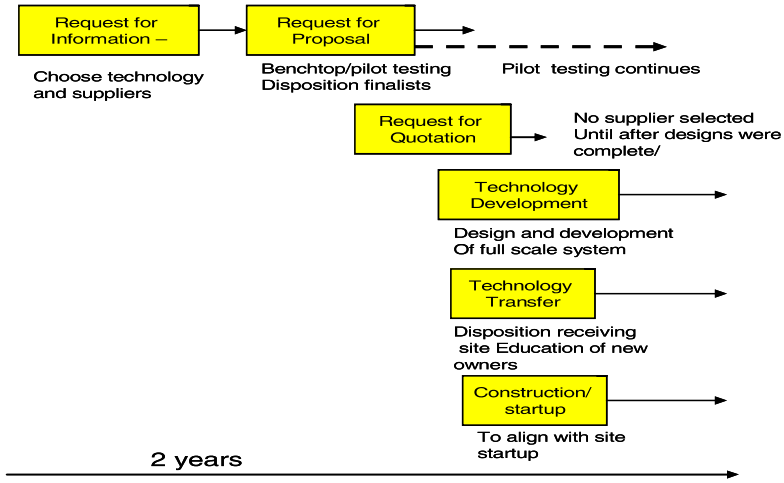

Figure 3: Modified selection process

Figure 3 illustrates the technology selection process that was used for this project. In the following section we review the parallel activities and the risk-taking decisions that drove this process.

Request-for-Information

By engaging in the parallel activities outlined above, the task force was able to quickly resolve many issues. The purpose of the Request-for-Information phase was to gather as much information as possible about the technologies available to treat Intel's wastewater. One of the key tenets that the project team used during this phase was redundancy. Because the schedule would be compressed, and the time to test multiple iterations of technology was minimal, it was important to obtain the information as quickly as possible. Multiple third-party consultants and extensive literature surveys were used, in addition to Intel's experts, to ensure that all technology options were being explored. One of the first activities was to characterize semiconductor wastewater for the constituents of concern. The outcome showed that in addition to the process wastewater containing organics

and nitrogen, the wastewater from one of the Fab functional areas and the slurry solids from another area could require pretreatment. Figure 4 outlines the potential treatment and segregation streams.

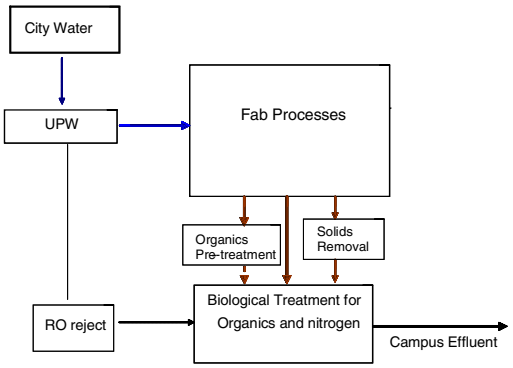

Figure 4: Segregation and pretreatment options for semiconductor wastewater

The next phase was to narrow down the options for treatment. Because of the restrictive solids limitations, only advanced biological treatment options were chosen. After additional analyses with vendors, only one technology seemed feasible to meet required timelines: the membrane bioreactor (MBR). See Figure 5 for a schematic of this MBR technology [8].

Request for Proposal/Request for Quotation

With viable MBR technologies available from two suppliers, the taskforce made the "at risk" decision to continue through pilot testing and design with both suppliers. With two options being developed, the development team intended to again use redundancy to compensate for the compressed schedule. This initial cost in both support (manpower) and rental fees was appreciable, but it allowed the team to gain valuable insight into the two suppliers' true capabilities regarding support, technical ability, operational needs, performance, and product quality. At the same time, Intel chose to operate each pilot system with "in-house" staff to ensure unbiased results while also minimizing operational costs. Part of qualifying the two systems involved placing analytical capability "in house", resulting in rapid turnaround of analytical results for wastewater samples. This allowed us to quickly fine-tune and respond to process excursions thereby gaining significant cost savings while enabling development of operating procedures for full-scale system implementation.

As the pilot plants were being assembled and shipped, the taskforce realized that the scheduled startup would be delayed if the major milestones outlined in Figure 3 weren't reached in parallel. Intel then asked both MBR technology suppliers to initiate designs even before feasibility testing and vendor selection was complete. In normal cycles, feasibility is confirmed, quotations are solicited, a vendor is chosen, and only then is design initiated. This new process would prove to have several benefits. In addition to providing competition, this allowed the vendors to supply accurate and quality bids that would be incomplete without the knowledge of the design details.

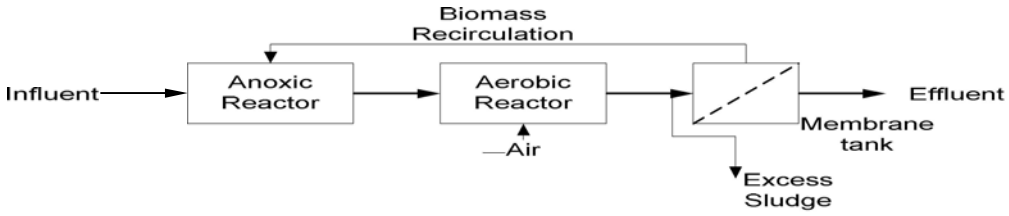

Figure 5: Membrane Bioreactor (MBR) process flow

As the vendors initiated design, both pilot systems confirmed that the MBR technology could treat Intel wastewater to the required effluent quality. One key barrier to validating the technology was that Intel's chemical "recipes" were still being developed. To reduce risk, one pilot system was left running throughout the technology development activities. An additional benefit from leaving one pilot system in operation would be to optimize design parameters and test further technology applications.

Request for Quotation (RFQ)

This phase of technology selection was another highlight for this team. With feasibility testing complete, and design proceeding, the task force was confident in the ability of each vendor to supply accurate bids. Each vendor was graded on technology development capabilities, initial design, bid preparation, and construction competence.

It should be noted that receipt of bids was a key milestone in this project. This marks the transfer of project

ownership from the Technology Development Group to the Project Team.

The supplier evaluation was performed by both the technology development task force and the project construction team. In addition, Intel's third-party experts were asked to review the "MBR Technology" scoring, to ensure accuracy, and to provide the critical second opinion. When the evaluation was complete, the supplier was chosen and announced. With the design work nearly complete at the time of supplier selection, the project moved quickly into construction. However, the Technology Development phase (including the pilot) continued to present new challenges that constantly altered the design and operations in the coming months.

Technology Development

At the time the contract was awarded to the vendor, the technology development task force had been working for nine months, and a pilot system had been running for six months. However, at this time, only feasibility had been proven, and there were still many unanswered questions. Among these was the ability of the chosen MBR technology to remove certain organics and solids without pretreatment.

At this point, technology development, transfer, and construction became concurrent activities. Due to the complexity of this process and the continually changing effluent composition, technology development was viewed as the key to success. Analyzing and characterizing the multiple organic compounds, modeling the outputs, setting up Design of Experiments for the pilot, developing analytical capabilities, and many other activities were keystones of the technology development process.

Unlike most projects (where design is not final), the construction team was already preparing the site for building even before the supplier selection was finalized. Late-breaking technology decisions would be incorporated into the design real-time and adjustments made accordingly. To offset the expected change orders, the project team initiated an aggressive Value Engineering effort to minimize cost while maintaining all the necessary treatment requirements.

Changes to the full-scale system based on technology development outcomes are expected to continue throughout the construction and start-up phase.

MEMBRANE BIOREACTOR PERFORMANCE AND APPLICATIONS TO FURTHER WATER REUSE

Background

The semiconductor industry uses ultrapure water (UPW) to rinse impurities from the wafer surface. Once contaminated with dissolved and suspended solids, organic and ammonia nitrogen, and multiple organic compounds, most semiconductor manufacturers discharge their wastewater offsite to municipal treatment plants, where these contaminants are removed. Although some companies reclaim some of this water for facilities applications (e.g., scrubbers, cooling towers), most UPW is treated as a "once-through" opportunity, limiting possibilities for water conservation.

MBR technology is not new and has been proven in many applications for organic, suspended solids and nitrogen removal. The technology development team quickly noted that effluent from the pilot plant was producing water comparable in quality to city water and might be suitable to replace city water in the UPW system feed. This was consistent with the data in literature that indicated that the MBR was quickly gaining recognition as the most effective process for treatment of industrial wastewater. In many cases, the MBR has been proven to produce very high-quality effluent suitable for almost unrestricted recycling. In fact, some U.S. states (e.g., Oregon) are currently revising their rules for reclaimed water to classify MBR effluent as Level IV reclaimed water, the highest purity classification available [13].

Semiconductor Wastewater Characteristics

Data in the literature indicates that the typical industrial effluent from semiconductor plants has the following characteristics, listed in Table 2.

Table 2: Characteristics of semiconductor plant effluents

Parameter	Concentration	Reference
pH	6-9	[14]
Biological oxygen demand (BOD)	50-200 mg/L	[14], [6]
Chemical oxygen demand (COD)	200-300 mg/L	[15]
Nitrogen	10-70 mg/L	[14]
Suspended solids (SS)	20-50 mg/L	[14], [15]

Of the organic compounds typically used in semiconductor manufacturing, the compound used in the new plant had been reported to pose a problem for biological treatment. Although the organic compound can be biodegraded by certain strains of bacteria [7, 10] its biodegradation in conventional wastewater treatment systems was reported to be slow and incomplete [5, 11]. For this reason, both total organics degradation and the removal efficiency for this particular organic compound were evaluated.

Organic matter (shown as COD and BOD in Table 2) has been demonstrated to severely reduce the performance of RO membranes [3]. Therefore, effective removal of organics is also critical if treated effluent is to be reused.

MBR Pilot Operation

The MBR pilot was operated at an Intel Technology Development location for 14 months. The pilot system treated 1.8 gpm of the combined effluent from the factory and consisted of aerobic and anoxic biological reactors followed by the membrane module. Pretreatment was accomplished using an ultrafilter (UF) module utilizing hollow-fiber membranes operated under vacuum. Additionally, 100 mg/L of the problematic organic compounds were added to the waste stream over the entire period of operation to simulate worst-case conditions.

Analytical

The pilot was equipped with an on-line system for measuring turbidity in the final effluent. Suspended solids, oxidizable organics, and nitrogen (ammonia nitrogen, organic nitrogen, nitrates, and nitrites) were measured in both waste stream and treated effluent. In addition, Total Organic Carbon (TOC) was periodically measured using a TOC analyzer. Tests were performed in accordance with Standard Methods [1]. For the compound in question, concentration in the treated effluent was measured by a vendor using a proprietary method.

RESULTS

Treated Effluent Quality

Biological Oxygen Demand (BOD) tests are typically used to measure organic matter content in treated wastewater. However, these parameters cannot be directly translated into a concentration of organics if the exact chemical composition is not known. For this reason, TOC measurements were also conducted.

Analyses of data from the pilot operations indicated stable and effective removal of nitrogen compounds, organic matter, and suspended solids to levels lower than required (Table 1). In addition to nearly complete removal of organics and TKN, complete biodegradation of the other organic compound was confirmed by direct measurements. In all samples the concentration was below the detection limit of 4 mg/L.

Several projects have recently reported the feasibility of reuse of biologically treated municipal wastewater. Most of the systems included microfiltration (MF) followed by RO [12, 9]. A range of treated wastewater parameters from the literature, for which MF/RO treatment was successfully implemented, is shown in Table 3.

Table 3: Treated wastewater quality parameters

Parameter	Values reported in literature	Values observed in this study
BOD, mg/L	2-12	Below 0.5
TOC, mg/L	11-40	1.5-8
Turbidity, NTU	2.3-33	Less than 0.2
TDS, mg/L	490-1200	NA
Conductivity, µS/cm	820-2000	NA
Organic N, mg/L	1.4-9.7	Less than 0.4

MF pretreatment was needed to remove suspended solids that were present in the treated municipal wastewater, as evidenced by the effluent turbidity. Data in Table 3 suggest that combined effluent of a semiconductor Fab can be treated to achieve a quality exceeding that reported in other successful water reuse projects. Low turbidity values observed in this study are due to the fact that MBR utilizes UF for biomass separation, and that effectively removes all suspended solids and potentially eliminates the need for RO pretreatment. In fact, RO was successfully applied directly to the MBR effluent in several previous studies [2, 4].

Cost Considerations

Water reuse concepts ranging from irrigation to UPW recycle as described in this paper can be implemented at competitive cost, providing both monetary and resource conservation benefits. The feasibility of this approach can only be determined on a case-by-case basis. The following considerations must be taken into account in the cost/benefits analysis:

Costs

- Capital cost of the MBR.
- MBR operating cost.
- UPW system upgrade if the existing system does not have the capacity to operate at reduced recovery rate.
- Higher volume of RO reject.

Benefits

- Environmental protection and water conservation.
- Greatly reduced dependency on local infrastructure.
- Reduced water bill and wastewater treatment fee.
- Reduced use of RO pretreatment systems, since MBR effluent can be directly fed to RO.
- In some cases, MBR may eliminate the need for certain waste treatment systems (e.g., ammonia, solvents, waste neutralization), thus fully or partially off-setting the capital cost.
- Reduced dependence on availability and quality of local water supply.

CONCLUSION

In summary, environmental challenges that threatened to delay the startup of a new semiconductor facility drove the technology development team to define a new technology selection process as well as a new treatment technology. The new technology application met all environmental requirements and will enable reuse of all factory wastewater. The team had a clear charter that prioritized schedule, thus enabling informed risk-taking to speed development, and redundancy to minimize risk.

It has been proved that MBR can adequately treat semiconductor wastewater to the standards proposed by local authorities to enable unrestricted reuse of the wastewater for irrigation. In addition to implementing the technology to enable startup, many side benefits are possible as a result of the high water quality that will be discharged from MBR applications. In arid geographies where Intel operates several facilities, the unlimited reuse opportunities for water are even more important.

ACKNOWLEDGMENTS

We acknowledge the following for their contributions to this project: Steven Kinsler, John Morgan, Jerry Meyers, Kathleen Fiehrer, Minh Lu, Sandra Chadwick, Ariel Rosenthal, Dana Grooper, Ronen Wohlfarth, Gershon Avraham, John Ormando, Leo Kenny, Ted Reichelt, and Chris Lee.

REFERENCES

[1] APHA, AWWA, WPCF, *Standard Methods for the Examination of Water and Wastewater*, 20th Edition, American Public Health Association, American Water Works Association, Water Pollution Control Federation, Washington, D.C., 1998.

[2] Comerton, A.M., Andrews, R.C., and Bagley, D.M., "Evaluation of an MBR-RO System to Produce High Quality Reuse Water: Microbial Control, DBP Formation and Nitrate," *Wat. Res*., 39, 3982–3990, 2005.

[3] Conner, W., Liu, J., and Yee, J., "Comparative Evaluation of MBR/RO and PAC Enhanced MBR/RO Treatment of Refinery Oily Wastewater," in *Proceedings of the 80th Annual Water Environment Federation Technical Conference*, San Diego, CA, October 13–17, 2007.

[4] DeCarolis, J., S., Adham, J., Oppenheimer, W.H., Pearce, L., and Wasserman, L., "Integration of MBR Technology and Desalting Membranes for Water Reuse," *Membrane Technology Conference & Exposition*, Phoenix, AZ, March 6–9, 2005.

[5] Den, W., Ko, F.-H., and Huang, T.-Y., "Treatment of Organic Wastewater Discharged from Semiconductor Manufacturing Process by Ultraviolet/hydrogen Peroxide and Biodegradation," *IEEE Transactions on Semiconductor Manufacturing*, 15, 540–551, 2002.

[6] Hashimoto, N. and Sumino, T., "Wastewater Treatment Using Activated Sludge Entrapped in Polyethylene Glycol Prepolymer," *Journal of Fermentation and Bioengineering*, 86, 4, 424–426, 1998.

[7] Lin, S.H. and Yang, C.R., "Chemical and Physical Treatments of Chemical Mechanical Polishing Wastewater from Semiconductor Fabrication," *Journal of Hazardous Materials*, 108, 1–2, 103–109, 2004.

[8] Metcalf and Eddy, Inc., *Wastewater engineering: Treatment, and Reuse*, 4th ed., Revised by G. Tchobanoglous , F.L. Burton, and H.D. Stensel, McGraw-Hill, New York, 2003.

[9] Patel, M., Knoell, T., Sevenandt, W., and Lacy, S., "70 MGD Indirect Potable Ruse: Optimization of Orange County's Groundwater Replenishment System," in *Proceedings of the 80th Annual Water Environment Federation Technical Conference*, San Diego, CA, October 13–17, 2007.

[10] Tsuzakai, M. and Nakamoto, S., "Biodegradation of Tetramethylammonium, an Ingredient of a Photoresist Developer, by Bacteria Isolated from Activated Sludge," *NEC Research and Development*, 97, 8–12, 1990.

[11] Urakami, T.; Araki, H.; and Kobayashi, H., "Isolation and Identification of Tetramethylammonium-biodegrading Bacteria," *Journal of Fermentation and Bioengineering*, 70, 1, 41–44, 1990.

[12] Wallis-Lage, C., Freeman, S., Steichen, M., Bates, J., Peries, I., and McEvoy, J., "Reuse Downunder: Membranes Provide Drought Relief for Queensland," in *Proceedings of the 80th Annual Water Environment Federation Technical Conference*, San Diego, CA, October 13–17, 2007.

[13] Wilson, S., Wilson, D., and Melcer, H., "Evaluating Applicable Treatment Technologies for an Integrated and Sustainable Food Processing Facility," in *Proceedings of the 80th Annual Water Environment Federation Technical Conference*, San Diego, CA, October 13–17, 2007.

[14] Electronics Manufacturing, *Pollution Prevention and Abatement Handbook*, World Bank Group, 1998.

[15] You, S.-H., Tseng, D.-H., and Guo, G.-L., "A Case Study of the Wastewater Reclamation and Reuse in Semiconductor Industry," *Resource Conservation and Recycling*, 32, 1, 73–81, 2001.

AUTHORS' BIOGRAPHIES

John Currier is a Senior Environmental Engineer in the Technology Development EHS Group. During his 20 years of experience in the ultrapure water and wastewater field, John has spent the last seven at Intel. Currently his primary responsibility is setting worldwide environmental goals that allow expansion of Intel's manufacturing operations while minimizing adverse environmental impacts. John holds a B.S. degree in Chemical Engineering from North Carolina State University and an M.S. degree in Engineering Management from the University of Tennessee. His e-mail is john.r.currier at intel.com.

Vyacheslav (Slava) Libman is a Staff Engineer and Strategic Technology Area Lead in Ultrapure Water and Environmental Engineering, responsible for technology development projects in these areas. Slava has over 17 years of experience in environmental engineering, out of which 9 years are with Intel. Slava received a PhD degree in Environmental Engineering from the Technion, Israel Institute of Technology. His e-mail is vyacheslav.libman at intel.com.

Dan Enloe is a Construction Materials Purchasing Manager for Intel Corporation, working there since 1984. He has published technical papers on ion implant technology and has two U.S. Patents in EUV lithography. He holds a B.S.E.E. degree from the US Naval Academy at Annapolis, has completed Navy post-graduate work in Nuclear Power, is qualified in nuclear submarines, and serves as a Captain in the U.S. Navy Reserve. His e-mail is dan.enloe at intel.com.

Don Crandall is a Senior Process Engineer for Intel Corporation, working there since 1991. He has held a variety of positions in the development and startup of water, waste, and wastewater treatment systems and is currently working in the Logic Facilities Technology Development group. He holds a B.S.E. degree from Arizona State University and is a registered professional chemical engineer. His e-mail is don.d.crandall at intel.com.

Boris Eliosov is an Environmental Process Engineer in the Flash Memory Group. He joined Intel in 2005, with 12 years of experience in water and wastewater treatment systems. Boris holds a PhD degree in Environmental Engineering from the Technion, Israel Institute of Technology and is a licensed professional engineer. His e-mail is Boris.Elilosov at intel.com.

more information on SPEC benchmarks, please see http://www.spec.org

THIS PAGE INTENTIONALLY LEFT BLANK

Dynamic Data Center Power Management: Trends, Issues, and Solutions

David Filani, Digital Enterprise Group, Intel Corporation
Jackson He, Digital Enterprise Group, Intel Corporation
Sam Gao, Digital Enterprise Group, Intel Corporation
Murali Rajappa, Digital Enterprise Group, Intel Corporation
Anil Kumar, Digital Enterprise Group, Intel Corporation
Pinkesh Shah, Digital Enterprise Group, Intel Corporation
Ram Nagappan, Digital Enterprise Group, Intel Corporation

Keywords: dynamic power management, data center, data center operations, energy efficiency, platform management policy, power limiting, power capping

ABSTRACT

In this paper we examine the challenges of increasing data center power consumption and higher energy costs in the face of ever-increasing computing needs. An examination of how power is allocated to computing resources in data centers shows that current methods do not result in optimal use of available data center power and space. We identify requirements that server platforms must address to solve data center power problems, and we offer a solution that includes a platform resident Policy Manager (PM). The PM monitors power and thermal sensors and enforces platform power and thermal policies. We explain how the PM can be used as the basis of a data center power management solution. We present results from a Proof of Concept (PoC) implementation, and we conclude by showing that a policy-based approach is powerful for maximizing power allocation within a given power envelope and increasing server density in data centers.

INTRODUCTION

One of the biggest challenges for data center operators today is the increasing cost of power and cooling as a portion of the total cost of operations. As shown in Figure 1, over the past decade, the cost of power and cooling has increased 400%, and these costs are expected to continue to rise. In some cases, power costs account for 40-50% of the total data center operation budget. To make matters worse, there is still a need to deploy more servers to support new business solutions (Figure 2). Data centers are therefore faced with the twin problem of how to deploy new services in the face of rising power and cooling costs. In a recent survey of data centers (Figure 3), 59% identify power and cooling as the key factors limiting server deployment.

If these trends continue, the ability of data centers to deploy new services will be severely constrained. To overcome this constraint, data centers have three choices: expand power and cooling capacity, build new data centers, or employ a power management solution that maximizes the usage of existing capacity. The first two choices can be very expensive because they involve capital expenditure for purchasing and installing expensive new power delivery equipment. For this reason, the power management approach bears close examination, and this approach is the focus of the rest of our paper. For previous work in this area, the reader is referred to Felter et al. [4] who examine the benefits of dynamic power budget allocation, Femal [5] who examines the benefits of monitoring and coordinating power distribution to achieve higher application throughput, [6] where a framework to monitor power is discussed, and Bianchini [7] who presents a survey of energy management techniques by type of server system.

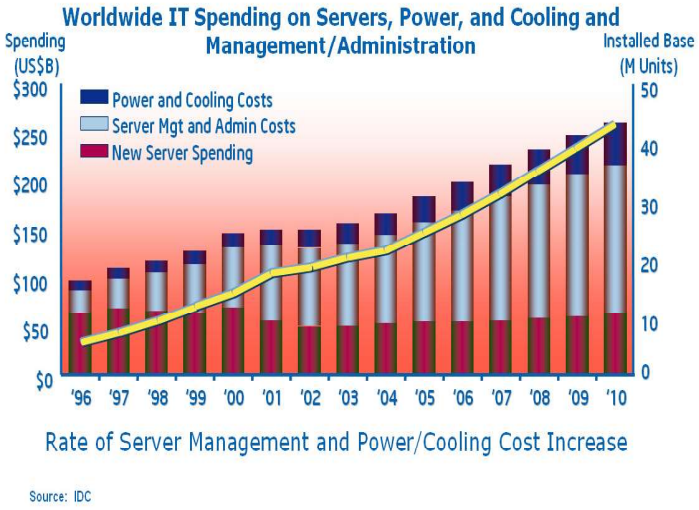

Figure 1: IDC Report of data center cost structure and trend

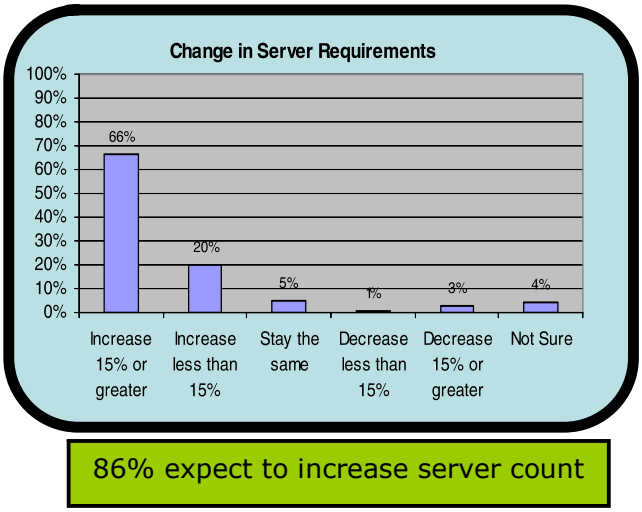

Figure 2: Expected growth in server count

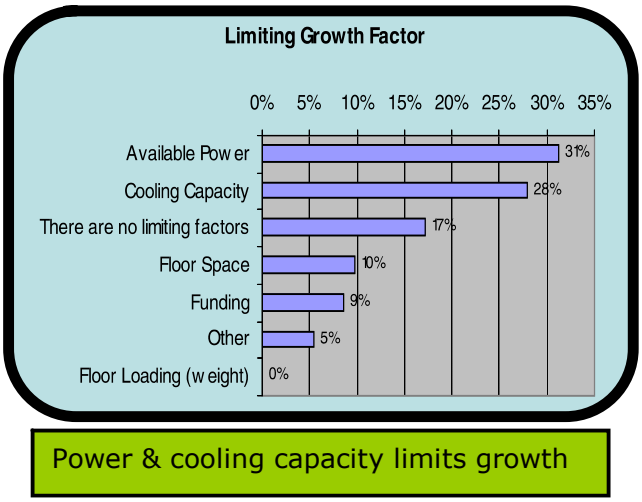

Figure 3: Factors limiting server growth

Our paper is organized as follows. In the next section we describe the current data center power allocation approach and the resulting problems; we follow this by proposing a new method for dealing with data center power allocation and describe the resulting server requirements; we then describe the role and functions of a platform resident Policy Manager (PM) and show how it addresses these requirements. Finally, we present the results of a PM Proof of Concept (PoC) and the benefits it offers data centers.

CURRENT POWER ALLOCATION METHODS

A typical data center power distribution hierarchy is designed to deliver a fixed amount of power to the room and then to each rack. The challenge of the data center operator is to determine the number of servers for each rack while ensuring that the overall rack (hence room) power consumption does not exceed the limit. To do this, the operator must make certain assumptions about the maximum power consumption per server. For most data centers, there are two ways of determining this: 1) using server nameplate power value, and 2) using a derated nameplate value.

The server nameplate value, which is marked on the server by the manufacturer, is the maximum possible power value that the server can consume. Actual power consumption is typically much less that the nameplate power. Most data center operators are aware that typical server power consumption never reaches the nameplate value, and one way for them to increase server density is to derate the nameplate power by a certain percentage—depending on the workload that is deployed on the server.

While derating, as opposed to using nameplate value, can improve server density, it is obvious that both methods are not optimal. The result is the following:

1. Under-utilization of available power for computing: A static allocation of power based on the worst-case scenario planning leads to inefficiencies and does not maximize the use of available power capacity.
2. Under-population of rack space: This is a direct result of the static allocation approach described above. The wastage of rack space is more severe when racks are populated using the nameplate power.
3. The two factors above result in higher energy costs than actually necessary, primarily from the cost of cooling the room that is sparsely populated with servers.
4. Unnecessary power cooling capacity expansion: Although existing power capacity is not being utilized effectively (as described in 1, 2, and 3 above), if the data center needs to deploy new services, the operator has no choice but to expand power/cooling capacity or even build a new data center at very high costs.

From the analysis above, we can see that the current power allocation approach results in wastage and high total cost of ownership. A better approach is needed.

DYNAMIC POWER MANAGEMENT APPROACH

The fundamental problem with the current approaches just described is that they are not based on measurements of actual power consumption. The data center operator has no visibility into how much power each server (and hence the rack and room) is consuming at any given time, neither does he/she understand the power consumption pattern over time. Without that visibility, the operator (or data center management software) is unable to decide how much power to allocate to servers/racks based on actual need. A better approach is to allocate power and populate racks using the following steps:

1. Monitor actual power consumption to understand average and peak power utilization for the server/rack.
2. Dynamically allocate power to groups of servers to maximize power/space utilization while staying within the power constraints determined by the power and cooling capacities.
3. Dynamically reallocate power when necessary to accommodate shifts in power needs of servers.

To implement this dynamic power allocation approach, server platforms must address key power management requirements.

POWER MANAGEMENT REQUIREMENTS FOR SERVER PLATFORMS

In this section, we identify the server requirements that address the issues of data center power and cooling efficiency, power allocation, and power provisioning.

Power Measurement Requirements

Actual power consumption for each server must be measurable at any point in time. This allows a power-monitoring module to collect power consumption data for each server and aggregate power-usage values at the rack and data center levels over a period of time.

Real-time power monitoring will allow data center managers to see the trend of power usage over time. This allows her to identify key values such as minimum, typical, and peak usage. This information is very useful in planning for future expansions, in identifying where there might be power and cooling constraints, and in locating areas where new servers can be deployed without violating power and cooling constraints. In addition, the minimum, typical, and peak values can be used to determine an appropriate power policy for each rack and hence each server in the rack. System power can be monitored by communicating with power supplies that support the Power Management Bus (PMBus) interface [13].

It is also desirable to monitor the power consumed by server subsystems: CPU, memory, fans, disks, etc. As shown in [8] an understanding of subsystem power consumption characterization for workloads can be valuable for power adaptation. Hence subsystem power monitored values can be used for intelligent fine-grained power control and optimization algorithms.

Power monitoring eliminates the usage of nameplate or derated power value to determine the number of servers for each rack. Visibility into actual power consumption values allows IT personnel to determine the optimal number of servers to deploy per rack.

Power Control Requirements

As mentioned above, the ability to monitor server power consumption is in itself useful for power and cooling capacity planning purposes. But another key reason for monitoring power consumption is to determine appropriate power policies that can be set for servers in a data center. Such policies can be enforced autonomously

in the platform, by following the platform as a service model described in [2].

One important power policy is power capping. In their study of power usage at a large Internet services data center at Google, Fan et al. [1] found that power capping offers two advantages: first, it acts as a safety valve by protecting the power distribution hierarchy against overdraw; and secondly, it enables effective usage of the available power, thereby increasing rack population. As such, dynamic power capping is a primary power control requirement that must be addressed by a power management solution.

Power Usage Reporting Requirements

To monitor power and dynamically allocate/reallocate power in a data center, a standard interface must exist between data center management software and the server being managed to do the following:

- Monitor actual power consumption over a planning period to understand historical usage patterns for capacity planning.
- Provide current as well as peak power, minimum power, and average power over an interval.
- Notify the higher-level management system if the power policy cannot be enforced.
- Send alerts to higher-level management if a certain power threshold is reached.

Regulation Requirements

The American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) publish a guideline for providing the data needed for designing and provisioning a data center [9, 10]. A PM's power usage reporting capability should support this requirement. A PM can be an instrumental piece of monitoring carbon credit/generation. As carbon generation may have caps and limits in the future, a PM will be a fundamental tool to monitor and track power use, which is correlated to carbon generation. The Green Grid—a consortium of information technology companies and professionals seeking to improve energy efficiency in data centers around the globe—is developing the most advanced metrics for data center efficiency [11, 12], and the power management technology should support these reporting requirements.

SOLUTIONS: A POLICY-DRIVEN APPROACH TO POWER MANAGEMENT

A typical data center is managed by deploying management systems for monitoring the health of the computing infrastructure and performing various management functions including error detection, failure resolution, server provisioning, and service deployment to servers. We can extend this management infrastructure to enable the dynamic power management approach described above by introducing a new component: a server-resident PM that meets the requirements for power monitoring and control (see Figure 4).

The PM is responsible for enforcing power management directives that can be specified for each server. This is done by monitoring the appropriate power and thermal sensors on the servers, and by controlling the appropriate effectors that allow the PM to control the servers' power consumption as directed.

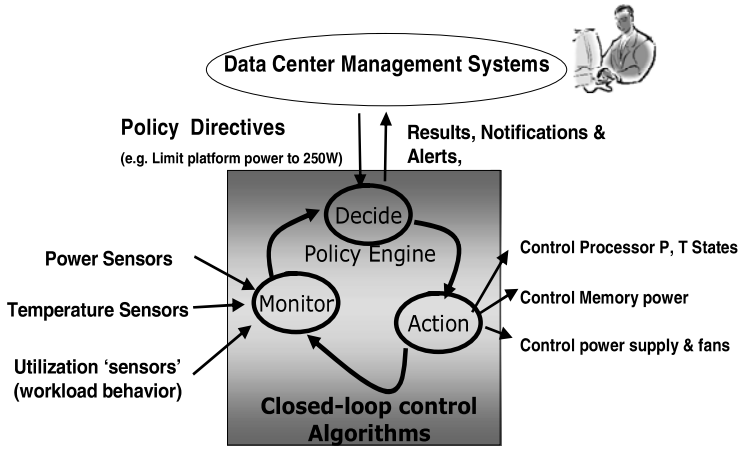

Figure 4: Power Management architecture

Intel's current series of server processors contains mechanisms that help control the power consumed by the processors. In addition, memory controllers provide ways of controlling power consumed by the memory modules. These can be used by the PM as low-level *effectors* for power control.

The interaction model between the PM and the data center management software is shown in Figure 4. The PM receives a policy directive from the Data Center Management System software. The policy directive may specify a power limit that is to be maintained for the server.

The PM has three major components operating in a closed loop manner: monitoring, policy engine, and control. The monitoring module is responsible for monitoring the sensors to determine if a new set of settings of the effectors needs to be enforced. This is done primarily by comparing the actual values of power and thermal against the policy limits. The policy engine decides a new set of controls based on the deviation of the actual power consumed from the limit and the utilization of various components. The control module then enforces the new settings as determined by the policy engine using mechanisms appropriate for that particular effector.

The sensors used by the monitoring module are the platform power sensors (including system and component powers) and the thermal sensors (system inlet/outlet temperature and component temperatures). Sensors providing information about the utilization of the components are used to determine the behavior of the workload.

The control module uses the effectors described earlier to effect power control on processor and memory. There are also additional effectors for components provided by the OEM (like power supply, fans, etc.) that are used by the PM to achieve greater power control.

When the power consumption exceeds the power limit, the PM will choose a set of settings for various effectors (processor, memory controller, etc.) that effectively reduce the power consumption of the platform. When power consumption goes below the limit, any restrictive controls previously placed on the processor, memory, and other effectors are relaxed. This is done in an iterative manner to account for the fact that the choice of effector settings may not have been accurate given the uncertainties in the behavior of the workload and the uncertainties in the parameters used by the PM to determine these settings. If the PM exhausts the use of all the controls at its disposal and still finds power consumption above the specified limit, it can generate alerts.

The PM exposes an abstracted interface for interacting with the external management software. It also requires interfaces for communicating with monitored components (e.g., temperature sensors, platform power measurements) and controlled subsystems (e.g., processor and memory modules).

The PM can be deployed in various ways:

- As a firmware running on a dedicated microcontroller in the server.
- As part of the baseboard management controller (BMC) that also performs other system management functions in the server.
- As an in-band agent in the operating system.
- As a combination of the above.

While the PM provides the capabilities described in the previous paragraphs, it depends on external management software to specify the policy parameters for it to operate effectively. The PM exposes an interface to the management software for this purpose. The interface includes commands to read power consumption and thermal data. The interface also allows management software to specify commands to set and get power control policies (e.g., set power limit) and to receive alerts from the PM.

The external interface is exposed as extensions to industry-standard server management protocols such as Intelligent Platform Management Interface (IPMI) [14] and Web Services Management (WS-Man) [15]. While IPMI is widely used in the industry today for server management, we anticipate WS-Man to gain more acceptances in the future.

CASE STUDIES AND EARLY RESULTS

To demonstrate the value of a policy-based dynamic power management approach using a platform-resident PM, we implemented the PM and conducted two sets of experiments: one set at a pilot data center, and the other in our internal labs. The rest of this section describes these experiments and the early results that demonstrated the value of this approach.

We conducted a proof of concept (PoC) at a top Internet portal customer's data center as a pilot project. The objective of this PoC is to maximize the number of servers allowed in a single rack within a given power envelope while maintaining maximum application performance.

Table 1 lists the use cases we developed for the PoC. At the beginning of test, we recorded the nameplate value of the servers (~350W) which the customer uses to populate their data centers today. We then installed servers with

PMs in the rack and measured the actual maximum power consumption of each server at peak search workload using the PM (~310W). We then used the observed maximum power value as the baseline for setting power limits for the servers and for determining number of servers per rack. We allowed a 10% headroom, to make sure that the power consumption at the rack level did not exceed the power envelope for prolonged period of time (10 min. or longer). The PM automatically adjusts power consumption toward this target, while continuing to deliver maximum performance for the given workload.

The data center management system communicates with the PM using IPMI [14] to continuously monitor actual power consumption of each server. It then aggregates power measurements at the rack level to make sure that the rack-level power envelope was not violated. The data center manager is used to set power limits dynamically for each server as desired to achieve the IT management policy. If the PM cannot maintain the limit set, or the data center manager observes a trend towards violating the rack-level power envelope, it resets the limits appropriately to ensure that the rack power envelop is not violated. With the interaction of the data center manager and the PM, the customer can safely achieve the maximum number of servers for a given rack-level power budget, thereby increasing the density of servers on a rack.

Table 1: Rack-level power optimization use cases

Use Cases	Description
Get power consumption on each server	Using the Policy Manager to dynamically gather point time power consumption from each server on the rack.
Estimate total power consumption of a rack	Estimate rack-level power consumption by summing up node-level power consumption; display on and notify to console as appropriate.
Optimize rack-level policy within a given power envelope and server workload	At rack level, analyze the power consumption of each server, overall power consumption, rack-level power envelope, and targeted performance goals (utilization, response time, query queue length, etc.) as well as other factors important to Baidu to determine the optimal power distribution policy. Baidu will set the policy and optimization strategy based on their work load and priority.
Set policy to servers on the rack	From the console, set policy to each rack in terms of particular power budget target that the server has to observe.
Node-level monitoring and tracking against policy	Leveraging Intel Node Server features to adjust server power consumption to the target set by the policy within 60 seconds and maintain at the target until further notice.
Node-level alert and notification	Use Policy Manager to detect and send alert when a server fails to reach policy target in 60 seconds or maintain the target during operation.
Alert handling and mitigation	Once an alert is received, the console needs to automatically decide on a course of action to mitigate the risk–ignore, set a new policy, or shut down the troubled server. …

The initial result from the PoC described above is summarized in Table 2. The result shows the performance measurement and power reduction observed for a single server in the PoC. The server has two Quad-Core Intel® Xeon® processors configured with 16 GB memory and a PM. The server was running actual search workload at the customer site in a near-production test environment.

It is interesting to note that when the workload is around 1,500 concurrent searches (above average workload) and the PM imposes a power cap at around 270W, the server CPU utilization and throughput virtually remain the same, i.e., ~67% and ~4.7ms per search respectively. This means that with a PM and proper power limit, we could save 40W from a server without performance loss when the CPU is not fully loaded. This is a 13% power reduction without performance loss. Under this circumstance, we could add one more server to a rack that is populated with 6 servers within the same power envelop.

It is important to understand that the value of PM is dependent on the actual application running on the server, the typical workload, and configuration of the server itself. For each combination of server and workload it is running, the user should determine the desired control points that reduces power consumption with minimal or no performance impact.

Table 2: PM Test Result on a Single Node

PM Setting	Platform Power Consumption	Workload	CPU Utilization	Search Time
No PM	310W	1,468 concurrent searches	67.81%	4.79ms
PM Power Capping 265 W	270W	1,514 concurrent searches	67.83%	4.69ms

In the second set of experiments which we conducted in our labs, we further explored the value of a PM by setting policies at different levels. We populated a rack with Intel® Bensley servers equipped with a Quad-Core Intel Xeon processor configured with a PM. We integrated servers under test with a management console, so that the management console could get real-time server power-consumption data and define policies to set a power-limiting target, while maintaining best possible performance at the limit.

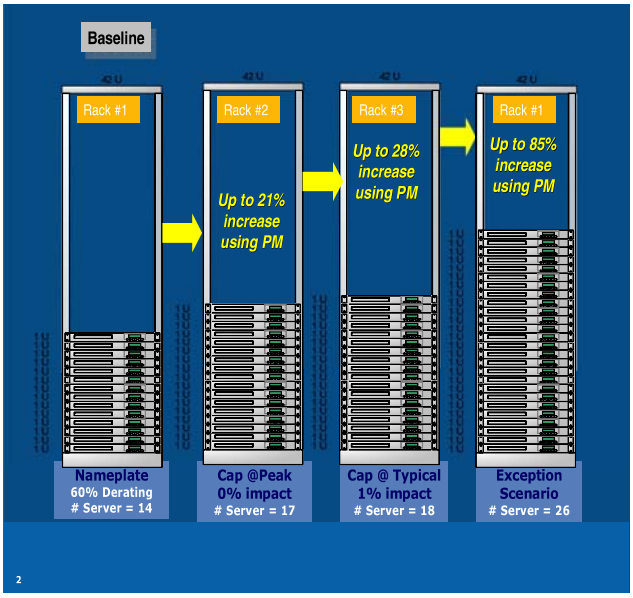

Figure 5: Policy Manager case study

Four different test scenarios were considered:

1. Populate the rack with nameplate power consumption (current customer practice).
2. Populate the rack based on power capping at maximum performance power measured for given search workloads (~280W).
3. Populate the rack based on power capping at ~98% of maximum power measured for given search workloads (~275W).
4. Populate the rack based on power capping at 90% of maximum power measured for given search workloads (~250W).

As shown in Figure 5, servers with a typical configuration, running representative workloads are populated in Rack #1. The rack is provisioned with a nameplate power with a derating factor of 60%.

Using the PM, as shown in Rack #2, by provisioning servers with power corresponding to maximum performance power, additional servers can be populated in a rack up to 21% more. Similarly, using the PM, Rack #3 shows an increase in server density of 28% when servers are provisioned with power that impacts peak performance 1% of the time. An exception scenario shown in Rack #4 is applicable when the power is severely constrained, for e.g., due to bad weather. In this case, the PM will limit power to the individual servers and hence to the whole data center, but it will allow the data center to operate in a stable environment. Each additional server defers data center capital expenditure by ~$2,000 [3].

Another example of the value of the PM is demonstrated in Figure 6. We show the power consumption and throughput measures with and without the PM for utilizations between 60% and 100% for WebBench load, which is a benchmark for Web traffic. For this workload, when the PM has a policy that limits the power to 20% below at 100% utilization, the impact to throughput is only 10%. The numbers are better at 80% utilization, where for a power reduction of 18%, the throughput is only 3%. It should be noted that typical servers in data centers run at utilizations well below 60% and as seen from the chart, the performance impact for a given power reduction is even less. Therefore, using the PM, we could safely craft power-capping policies to limit server platform power consumption with little impact on the peak performance of applications.

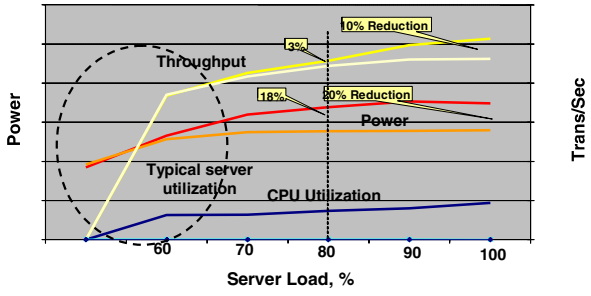

Figure 6: Power/performance with the Policy Manager

CONCLUSION

Current trends in data center power reveal a fundamental need for a power management capability on the platform that can be used to monitor power consumption and enforce power policies.

We have described an embedded policy manager (PM) as the foundational capability for a dynamic policy-based power management approach. The initial results of our implementation show that a dynamic policy-based power management approach using a PM can be used to increase server density within a rack power envelop, and to reduce power consumption with minimal performance impact.

In our future work, we will explore additional use cases and policies, and further investigate the benefits of fine-grained power-control methods.

ACKNOWLEDGMENTS

Special thanks to our colleagues Mike Patterson for providing regulatory requirements and Derek Collier for insightful suggestions during the review of this paper.

REFERENCES

[1] X. Fan, W. Weber, L. Barroso G; Bramlett, "Power Provisioning for a Warehouse-sized Computer," in *Proceedings of the ACM International Symposium on Computer Architecture,* June 2007, San Diego, CA.

[2] E. Castro-Leon, M. Chang, J. Hahn-Steichen, J. He, J. Hobbs, G. Yohanan, "*Service Orchestration of Intel-Based Platforms Under a Service-Oriented Infrastructure," Intel Technology Journal, Vol. 10, Issue 4*, November 9, 2006.

[3] J. Koomey, K. Brill, P. Turner, J. Stanley, B. Taylor, "A Simple Model for Determining True Total Cost of Ownership for Data Centers," *Uptime Institute White Paper*, October 2007.

[4] W. Felter, K. Rajamani, T. Keller, C. Rusu, "A Performance-Conserving Approach for Reducing Peak Power Consumption in Server Systems," in *Proceedings of the 19th annual international Conference on Supercomputing*, 2005, pps. 293–302.

[5] M. E. Femal, V. W. Freeh, "Safe Overprovisioning: Using Power Limits to Increase Aggregate Throughput," *Power-Aware Computer Systems*, Lecture Notes in Computer Science, Springer Berlin, Vol. 3471, 2005.

[6] J. Moore, J. Chase, K. Farkas, and P. Ranganathan, "Data Center Workload Monitoring, Analysis, and Emulation," *Eighth Workshop on Computer Architecture Evaluation using Commercial Workloads*, February 2005.

[7] R. Bianchini, R. Rajamony, "Power and energy management for server systems," *IEEE Computer magazine*, Nov. 2004, pps. 68–76.

[8] W. L. Bircher, L. K. John, "Power Phase Variation in a Commercial Server Workload," in *Proceedings of the 2006 international symposium on Low power electronics and design*, ACM, pps. 350–353.

[9] "Data Center thermal guidelines," at www.ashrae.org

[10] "Report to Congress on Server and Data Center Energy Efficiency," at http://www.energystar.gov/ia/products/downloads/EPA_Server_Report_Public_Draft_23Apr.pdf

[11] "The Green Grid Data Center Power Efficiency Metrics: PUE and DciE," at http://www.thegreengrid.org/gg_content/

[12] "Existing Metrics, Guidelines and Programs Affecting Data Center and IT Energy Efficiency," at http://www.thegreengrid.org/gg_content/

[13] PMBus standard information, at http://pmbus.info/

[14] Intelligent Platform Management Interface (IPMI), http://www.intel.com/design/servers/ipmi/spec.htm

[15] Web Services for Management (WS-Man), http://www.dmtf.org/standards/wsman/

AUTHORS' BIOGRAPHIES

David Filani is a Principal Engineer at Intel's Digital Enterprise Group. His current interests are platform power optimization and management, and he leads a team of architects focused on developing platform power management technologies. He holds a B.S. and M.S. degree in Computer Science from Arizona State University. He can be reached at david.filani at intel.com.

Jackson He is a Lead Architect in Intel's Digital Enterprise Group End user Platform Integration Group. He received PhD and MBA degrees from the University of Hawaii. His current focus is platform manageability and dynamic resource management in data centers. His e-mail is jackson.he at intel.com.

Sam Gao is a Program Manager at Intel's Digital Enterprise group in charge of directing server power management technology products. He received his MBA degree from the University of Chicago and an M.S. in Computer Science and Physics from Rutgers University. He is a graduate of Tsinghua University (Beijing, China) with a Bachelor's degree in Physics.

Murali Rajappa is a Platform Architect in Intel's Digital Enterprise Group Architecture and Planning Group. His

current focus is platform-level power management for next-generation computers with special emphasis on servers. His e-mail is muralidhar.rajappa at intel.com.

Anil Kumar is a Platform Architect in Intel's Digital Enterprise Group Architecture and Planning Group. His current focus is platform-level power management for next-generation computers with special emphasis on servers. He holds a PhD degree from the University of Connecticut and an MBA degree from Southern Methodist University. His e-mail is anil.k.kumar at intel.com.

Pinkesh Shah is a Platform Architect in the Digital Enterprise Group at Intel Corporation. He is involved in power management technology pathfinding for Intel's future server and client products. He received a PhD degree from the University of Arizona, an MSEE degree from Columbia University, and a BSEE degree from the New Jersey Institute of Technology. His e-mail is pinkesh.j.shah at intel.com.

Ram Nagappan is a Platform Power Architect in the Digital Enterprise Group at Intel Corporation where he is responsible for pathfinding for power and thermal solutions for Intel's next-generation server, client, and communications products. He received his Masters of Science degree in Engineering from Arizona State University. He can be reached at ramkumar.nagappan at intel.com.

THIS PAGE INTENTIONALLY LEFT BLANK

Evaluation Process for Semiconductor Fabrication Materials that are Better for the Environment

Victor Fan, Technology and Manufacturing Group, Intel Corporation
David Harman, Technology and Manufacturing Group, Intel Corporation
Jim Jewett, Technology and Manufacturing Group, Intel Corporation
Bob Leet, Technology and Manufacturing Group, Intel Corporation
Dawn Speranza, Technology and Manufacturing Group, Intel Corporation

Index Words: chemical, environment, Fab, health, materials, safety, selection, semiconductor

ABSTRACT

The successful future of the semiconductor industry is dependent upon the development of an array of new and novel materials to make semiconductor devices. In parallel, there is a growing sentiment in public policy and regulatory forums for a highly conservative precautionary approach to approving the introduction of new chemicals into commerce. Intel's business model of *Copy Exactly*! requires new material development very early in the technology development cycle. Therefore, Intel's approach to chemical development, and selection of materials, must incorporate consideration of the environmental impact of its policies and materials extremely early.

Prior to the selection and purchase of fab manufacturing materials at Intel, procedures are in place at the development and selection stages to integrate consideration of the materials' impact on the environment, to collaborate with suppliers and others to reduce that impact, and at the same time to meet Intel's technology needs. In this paper, we provide an overview of what drives the procedures used by Global Fab Materials, we look at how the procedures are used, and we look at the successful results through a case study.

INTRODUCTION

Technology Trend for Material Development

For more than four decades the semiconductor industry has been successfully producing one of the smallest and most effective devices that man has made, in step with Moore's Law [1] which basically predicts a doubling of the number of transistors per chip every two years. Today's modern semiconductors have nearly one billion transistors on them.

It takes over 400 individual steps of manufacturing and testing to make such a device. An individual semiconductor chip (with a postage-stamp-size of manufactured silicon inside, called a die) has not substantially changed in size relative to those seen in the mid 1980s, but the circuitry has increased in complexity 1,500 times. The Intel 386 chip had 275,000 transistors on a 1-micron feature size. Today's Intel® Penryn quad-core chip contains 820 million transistors on a 45-nanometer feature size. This increased complexity is attributable to both the continuous reduction in the size of the transistors on the die that allows for more transistors on the die, and the increased number of layers and features created within the external package surrounding the die, that are necessary to dissipate the intense heat away from the die.

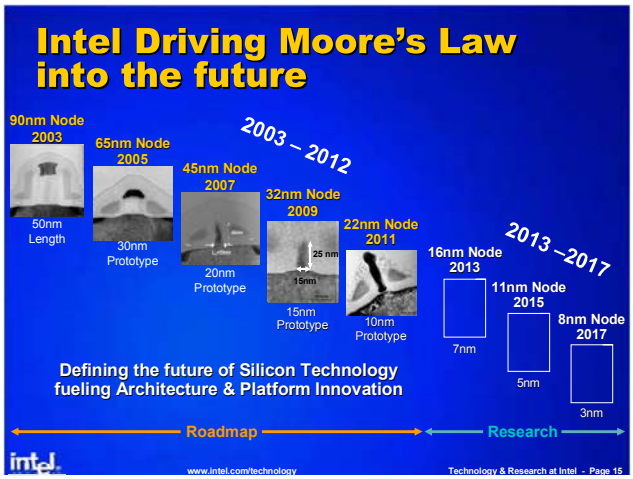

Figure 1: Continuum of reduced feature size

The shrinking feature size of modern semiconductors has created the need to be atomically precise in manufacturing and to have chemicals that are ultra pure. When one examines a 45-nm feature size transistor (see Figure 1), the channel (the area where electrons flow from one leg of

the transistor to the other) is as little as 20 nm, or approximately 100 silicon atoms across. With such a relatively small number of atoms responsible for the overall performance of each transistor, it is critical that they be manufactured without defects or impurities and provide exactly the technical attributes required for the semiconductor to function.

Control of the manufacturing process at the atomic scale has lead to a new generation of materials. Examples of these include material changes driven by the need to reduce the particle size and types of chemical interactions in the chemical mechanical polishing (CMP) of the silicon wafers[1]. As the feature size continues to decrease, the size of the mechanical polish particles has to decrease as well in order to prevent destruction of the features. For example, transforming a rough-cut stone into a lustrous gem requires many polishing steps using finer and finer grit. Similarly, the chemicals used in the CMP process will likely change any time the material composition of a feature changes (e.g., using mineral spirits to clean up oil-based paint versus water to clean up latex-based paint). Another example of chemicals that have changed in the past is those used in the etching and deposition[2] manufacturing process steps. In many of these process steps gas-based chemicals are used in place of liquid-based chemicals, since more control can be achieved with the etch depth and deposition thickness in gas-phase chemical reactions. The last example of materials change is in the wafer patterning area, referred to as lithography[3]. The challenges for lithography, in light of continually smaller features, are several, and two are described here. First, the feature size is now smaller than the wavelength of the light beam being used to make it, which means either the light source must be changed and/or the engineers must play chemical tricks with the patterning material (called photoresist) that is layered on the silicon wafers. Second, the intensity of the light source and the harshness of subsequent etch process steps both play a role in the composition of the photoresist that is used for any one lithography step. Further complexity arises due to the fact that each step may require a different photoresist.

Future semiconductor devices are critically dependent on the ability of stable and reliable materials to support device operation. As the market continues to demand an increase in scaling (miniaturization) with no penalty in performance, the need for new materials with increased mobility—lower energy and higher speed—is greater. Since devices are comprised of several materials and interfaces, the properties of new materials and their ability to interface with the properties of other materials will require materials with dramatically improved or new properties. All of these changing materials needs have lead to an explosion in the use of differing atoms from the Periodic Table of Elements (see Figure 2).

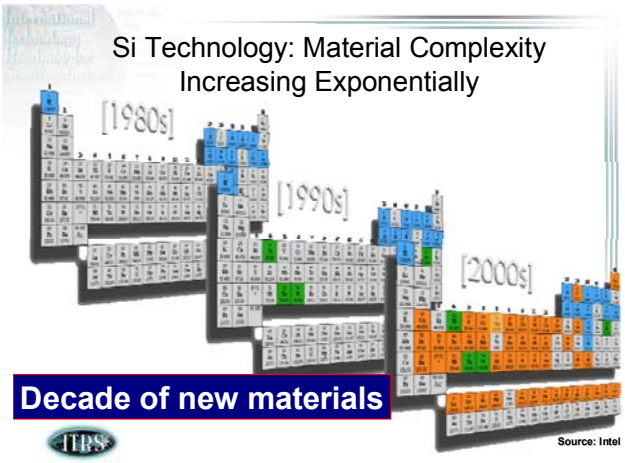

Figure 2: Increasing complexity of materials

PRECAUTIONARY POLICY

Thousands of chemicals have been developed and are used to manufacture the products we use in everyday life. Despite the widespread use of chemicals, until recently the prevailing thought was that exposure to most industrial chemicals was unlikely, especially outside the industrial environment. For years, regulatory policies toward such chemicals have presumed them to be safe with little or no information on their health implications. The U.S. EPA has found that even a basic level of toxicity information exists for less then ten percent of the approximate 2800 high-production-volume chemicals (those produced over one million pounds per year). The absence of information is often misinterpreted as evidence of safety. Growing evidence that some chemicals can potentially cause harm, and concern that current global chemical polices are not sufficient to predict or prevent potential harm to individuals or the environment, have led to major reforms in industrial chemical polices worldwide. The most prevalent perspective today is known as the Precautionary Principle, which in effect requires that precautionary action be taken before there is scientific certainty of cause and effect.

Seeking out and evaluating alternatives is preferred rather than asking what level of contamination is safe or economically optimal. The precautionary approach asks how to reduce or eliminate the hazard and considers all possible means of achieving that goal [2].

The precautionary perspective underscores a basic difference between hazard approach vs. risk approach. The first is absolute in terms of eliminating materials should they pose an "unacceptable" hazard. The risk approach evaluates whether or not a hazard will have an impact (e.g., human exposure) and determines how the impact can be mitigated.

"Copy Exactly!" Factory Strategy

Intel introduced its "Copy Exactly!" factory strategy in the mid-1980s and completed its adoption in 1996. Intel can credit "Copy Exactly!" with enabling the company to bring factories online quickly with high-volume practices already in place; hence, decreasing time to market and increasing production yields.

"Copy Exactly!" solves the problem of getting production facilities up to speed quickly by duplicating everything from the technology development facility to the volume-manufacturing factory. In particular, it means ensuring that the process devised at the development facility is fine-tuned not just for performance and reliability, but for high-volume production as well. (Background information on Intel's Copy Exactly! strategy can be found at [4].)

"Copy Exactly!" Versus Traditional Semiconductor Factory Strategy

In most semiconductor factories, equipment and processes used in research vary greatly from those used in high-volume manufacturing. At many companies, each new technology is brought to a technology development facility where a team of engineers precision-tune the process until it is perfected. Then the process is transferred to a high-volume manufacturing facility where a new set of engineers modifies the process so that it can be produced in large quantities.

The impact of Intel's Copy Exactly! policy on materials is that development and selection must occur much earlier in the technology development cycle, including the evaluation for Environmental Health and Safety (EHS) performance.

INTEL® GLOBAL FAB MATERIALS ORGANIZATION SYSTEM

Intel's Global Fab Materials (GFM) organization is responsible for early chemical development and procurement of semiconductor die manufacturing materials. Intel conducts a thorough EHS evaluation of chemicals prior to their use at Intel and implements state-of-the-art exposure control and environmental emissions abatement technology for managing their use. GFM deals with EHS considerations at the earlier stages of chemical development and selection from both industry and company perspectives. GFM's strategy is to enable continued global growth of our operations, while identifying and prioritizing new chemicals of EHS concern for further evaluation. We determine information requirements; and work on collaborating with our supply chain, consortia, and research and development organizations to balance the demands for new technology development, Copy Exactly!, and public policy concerns as part of materials development. Components of the GFM Materials EHS system include 1) External Research and Development, 2) Supplier Engagement, and 3) Integration with the overall Materials Risk Evaluation System.

EXTERNAL RESEARCH & DEVELOPMENT

Consortia Engagement

The mission of our external programs is to influence the direction of the research; then to extract the value from external R&D organizations and activities and to bring this back to Intel. These are some of the items that we look at:

- The current (and evolving) EHS trends and their potential impact on Intel.
- The response needed from university and fundamental research to address these EHS trends.
- The research needed to develop science and technology leading to simultaneous process performance/cost/EHS gain.
- Incorporating EHS principles into engineering and science education.
- Promoting Design for Environment and sustainability as a technology driver and business benefit.

On an industry level, Intel collaborates on EHS issues with other semiconductor manufacturers and with our suppliers of tools and chemicals in national associations, such as the International Sematech Manufacturing Initiative (ISMI), and the U.S. and European Semiconductor Industry Associations (SIA and ESIA, respectively). Broad industry needs with regard to EHS elements of new technology (including chemicals) are integrated into the International Technology Roadmap for Semiconductors (ITRS). The ITRS is a global industry 15-year roadmap that identifies technology requirements for the continued success of the semiconductor industry.

Intel also actively participates in the World Semiconductor Council (WSC), a global industry body whose efforts include pre-competitive cooperation on major EHS policy issues for the industry. Membership is composed of semiconductor trade associations from six leading global centers of manufacturing (EU, China, Taiwan, Japan, Korea, and the U.S.). Global collaboration at the WSC level has led to voluntary global industry agreements on the responsible use of chemicals. A key initiative has been the reduction of global warming by reducing perfluorinated compounds (PFCs) emissions by 10% lower than 1995 levels by 2010. A second example is

the full phase-out of perfluorooctyl sulfonates (PFOS) in non-critical applications while also continuing R&D to eventually phase out critical uses, where possible.

Environmentally Benign Materials Research (EBMR) is a key program that GFM-EHS uses to enhance the sustainability of current and future technologies. This program identifies critical material research needs for environmental purposes, such as alternatives to PFCs, and targets research to find solutions, or at a minimum, gain better knowledge about the nature of the problem. These projects can be accomplished internally with direct assistance from suppliers, or via industry consortia such as the ISMI, the Semiconductor Research Corporation's Engineering Research Center (SRC-ERC), and the semiconductor consortia in Belgium (**IMEC**), for example.

Intel has been a promoter of collaborative R&D efforts to create the science, technology, and educational methods to remain in a leadership position in promoting a safe and environmentally conscious supply chain.

Consortia External Supplier Engagement

Via some of these consortia efforts, Intel GFM has some indirect interaction with suppliers to assist in directing their efforts towards the industry's needs. An example of this is the SEMATECH Supplier Data Council. This team consisted of chemical suppliers and device makers whose mission was to determine how to obtain consistent timely EHS data for semiconductor chemicals. The team created a standardized methodology to guide the development of consistent EHS data by suppliers and a method to communicate the data to suppliers and downstream manufacturing users.

DIRECT SUPPLIER ENGAGEMENT

Addressing EHS concerns and issues successfully for the long-term sustainability of our industry requires close collaboration with Intel's suppliers. Furthermore, integrating EHS into the design of new chemicals is generally easiest in the early stages of development (see Figure 3). This is also one of the challenges stated in the ITRS. The specific challenge is a lack of timely information flowing to the technology teams about the EHS characteristics of new materials in order to help minimize the EHS impact of chemicals used.

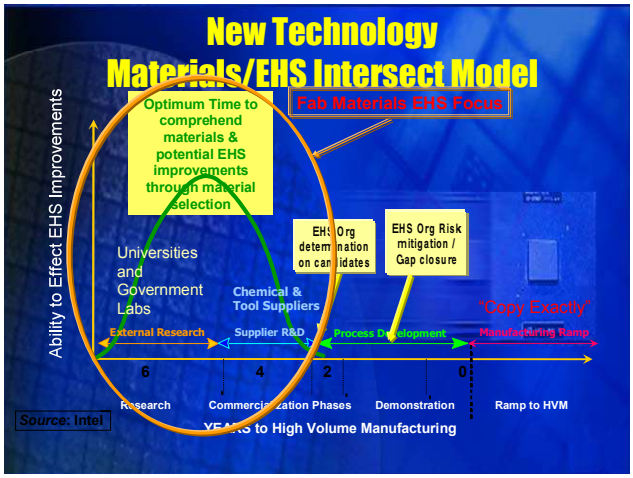

Figure 3: Integrating EHS into the design of new chemicals

One of the Intel GFM's goals is to close the chemical data gap by ensuring that Intel's chemical suppliers generate, distribute, and communicate information on chemical toxicity. Chemical suppliers are best equipped to acquire and provide this information and in some instances are required to do so by law. This is seen as necessary to meet the requirements of the European Union Registration, Evaluation, Authorization and potential Restrictions of Chemicals (REACH) legislation. REACH represents the EU's effort to address long-standing deficiencies in chemical information.

An electronic template was developed through collaboration between semiconductor manufacturers and suppliers that allows suppliers to provide specific EHS information data sets as a condition of the purchase of the chemicals. This information is then used to make risk-based chemical decisions to enable the use of the chemistry. In addition, we have certain basic supplier guidelines for developing new chemicals [3].

On an annual basis, Intel GFM EHS communicates with supplier executives on EHS expectations and exchanges information regarding EHS trends and issues at Intel Supplier Day. In 2006 for example, Intel provided pertinent information on the new REACH regulation that communicated the primary supplier requirements to ensure a continued supply of chemicals to Intel's European operations. We provided information on the new international mandate for a globally harmonized hazard classification and compatible labeling system, including material safety data sheets (MSDS) and easily understandable symbols. Intel also provided information on the new global restriction of hazardous substances regulations and Intel's supplier intellectual property management systems. Intel Material Supplier Day presentations are an effective mechanism to engage

supplier executives outside traditional supplier-to-customer relationships.

Intense competition for chemical market share makes the security of intellectual property (IP) a critical component of supplier management. Simultaneously, accurate and timely disclosure of EHS data for proprietary chemical ingredients is crucial for a comprehensive EHS risk assessment and avoidance of any current and future regulatory consequences. A mechanism was developed by Intel's GFM and EHS groups to improve the flow of EHS data to ensure the safe use of chemicals while simultaneously improving the protection of suppliers' IP. This mechanism, the Supplier EHS IP Management System (SEIMS), is a data management application for EHS-related documents containing supplier intellectual property. These are some of its key features:

- Design to enhance the security of suppliers' IP and limiting dissemination within Intel to only those with a need-to-know to make sure that chemicals can be used safely.
- Ensures appropriate and secure document management within Intel.
- Enables supplier document submission online to Intel, eliminating the inefficiency and security issues of manual document management.
- Intel personnel without a need-to-know are unable to view, physically control, or manipulate the supplier's IP.

Chemical suppliers play a vital role in our quest toward greener chemistry and sustainable manufacturing. On a proprietary basis, Intel identifies areas where specific EHS concerns with certain materials exist, and we coordinate closely with suppliers to develop material solutions where both technical and EHS needs can be met.

INTEGRATION WITHIN INTEL

The final parts of the GFM EHS system are a Regulatory Early Screening process, EHS elements in the general Materials Risk Assessment, and integration of the GFM results into Intel's overall EHS material evaluation and risk mitigation process. At this stage of technology development, a list of candidate materials is created that focuses the development and EHS systems on likely chemicals for the new generation of technology.

Early Screening

GFM developed and implemented a regulatory early screening procedure to proactively evaluate and mitigate associated risks identified through government authorities responsible for approving new chemicals for the industrial market. The scope of early screening covers all legacy materials as well as potential candidate chemicals for the new technology. This analysis is done in conjunction with Intel suppliers who develop the EHS and industrial application information portfolio for their materials. Since requirements vary with regard to both the location of chemical manufacturing and the final location for semiconductor manufacturing, this screening can include multiple schemes of data and information from various geographies (e.g., U.S. Toxics Substances Control Act, Korean Toxic Substances Control Act, etc.).

Materials Risk Assessment

Another key responsibility of Intel GFM-EHS is to ensure that regulatory or EHS constraints in the Fab materials supply chain are mitigated prior to chip high-volume manufacturing (HVM) proliferation. Failure to comply with regulatory requirements can lead to prohibition of materials shipment to the factories. Lack of proper EHS risk evaluation of materials before HVM use may result in last-minute risk mitigation actions being required as the new semiconductor technology is ready for delivery to HVM facilities.

A Materials Risk Assessment (MRA) program is utilized to determine and then mitigate potential risk in materials EHS readiness and regulatory compliance. Throughout the development lifetime of each new technology, any issues identified through the MRA evaluation are logged, mitigation actions are identified, and a completion plan is implemented to manage each issue identified.

Internal Collaboration with Intel's EHS Cross-site HVM Review Process

Intel policy requires a comprehensive EHS review of all chemicals identified for use at Intel. GFM development and procurement experts and evaluators in EHS departments across the company all work closely to achieve the Intel EHS mission. Materials evaluation is conducted before the first sample shipment to an Intel facility. The evaluation starts with a thorough chemical EHS characterization including, but not limited to, environmental impacts, human health and toxicity hazards, and occupational safety considerations. This review builds upon the information compiled during GFM's development and supply chain efforts, and it extends analysis to site-specific environmental, health, safety, and waste issues. The GFM system dovetails with the EHS organization's program at this stage, resulting in a complete EHS materials program from concept to technology end-of-life.

CASE STUDY ON PERFLUORINATED COMPOUNDS (PFCs)

PFCs are key materials used for plasma chamber cleaning in chemical vapor deposition (CVD) and for plasma dry etch. They are also a group of materials that have been identified as having high Global Warming Potentials (GWPs) in reference to the GWP of carbon dioxide (CO2). In some instances, these GWPs can be thousands of times higher that that of CO2. The high GWP of PFCs led to a great deal of focus on the reduction of their emissions beginning in the late 1990s, and that focus is expected to continue for many years.

The reduction of PFC emissions was the first effort by the WSC to establish a voluntary agreement across the semiconductor industry. The goal was a 10% reduction in absolute PFC emissions from 1995 levels, by 2010. This required extensive collaboration between semiconductor manufacturers, equipment suppliers, materials suppliers, and research from universities and consortia. To develop the agreement and a roadmap for reductions also required collaboration within Intel between Government Affairs, Legal, EHS, Technology Development, and GFM-EHS.

The reduction efforts focused on the high-end of the pollution prevention hierarchy where reduce, reuse, and replace were the priority. The search for alternatives tied into another GFM-EHS strategic program, Environmentally Benign Materials Research (EBMR). Significant progress has been achieved over the past 10+ years in reducing emissions through the evaluation and integration of these environmentally benign materials into the manufacturing process. Using the screening process on new potential materials allowed for the selection of the best alternative, without replacing one environmental problem for another. The development of replacements came from a wide range of sources from universities to suppliers. However, the integration of these materials into the manufacturing process required significant work due to the complexity of making a chip.

Currently, Intel Corporation is on track to meet the voluntary goal for 2010 and continues to not only focus on opportunities for reducing emissions but inserting distinct requirements for reduction into future technology development roadmaps. This continues to push us to seek out new alternatives and methods for reductions.

CONCLUSION

The Intel Materials' environmental, health and safety early screening and materials management program has been implemented to ensure stability of the Fab materials supply-chain through identification and management of EHS-associated risks. This supports both Intel and semiconductor industry leadership in developing environmentally preferable materials, protecting public health and the environment, and maintaining a safe and healthy workplace, while continuing to enable new technology.

ACKNOWLEDGMENTS

John Harland for his guidance, encouragement and efforts evaluating the many drafts of this article.

Our EHS-focused colleagues and managers from many organizations at Intel who participate in and support the efforts to meet the objectives of this process successfully.

Fellow travelers in the semiconductor industry—Manufacturers, Suppliers, Researchers and Educators—without whose common vision none of this would be possible.

REFERENCES

[1] See http://www.intel.com/technology/mooreslaw/index.htm

[2] Joel Tickner, Carolyn Raffensperger, and Nancy Myers, *The Precautionary Principle in Action: A Handbook*, at http://www.biotech-info.net/handbook.pdf

[3] https://supplier.intel.com/static/EHS/materials.htm

[4] Background information on Copy Exactly!, at http://www.intel.com/pressroom/kits/manufacturing/copy_exactly_bkgrnd.htm

AUTHORS' BIOGRAPHIES

Victor S. Fan is a Senior TD Materials Engineer with Strategic EHS Programs in Intel's Global Fab Materials (GFM). His experience with Intel includes chemical EHS characterization, business program establishment, and as an Intel representative to SEMATECH. Experience outside Intel includes engineering and management positions in petrochemical and health care industries. He received an M.S. degree from Harvard University and a PhD degree from Johns Hopkins University.

David Harman is a Senior Environmental Engineer in the Intel Technology Development EHS Group. He works on a variety of strategic and technical programs to support materials, equipment, and process development of new technologies. He also is the Intel representative to Sematech EHS programs. He has a B.S. degree in Chemical and Environmental Engineering from the University of Arizona.

Jim Jewett is a Principal Engineer and Manager of Global Fab Materials' (GFM) EHS group, whose responsibility is to ensure that Fab materials provide cost-effective environmental and health solutions in their

design and selection, while enabling Fab process technologies and maintaining a secure supply chain. In addition to the internal responsibilities, Jim chairs both the Chemical (EHS) Management program for the World Semiconductor Council and the International Technical Working group on EHS for the ITRS (International Technology Roadmap for Semiconductors). He is Intel's Industrial Advisory Board member to the SRC/SEMATECH Engineering Research Center for Environmentally Benign Semiconductor manufacturing. During his 23 years at Intel, Jim has remained focused on environmental programs while belonging to a variety of organizations including Corp EHS, Fab Manufacturing, Components Research and Materials.

Robert (Bob) Leet is a Senior TD Materials Engineer at Intel Corporation, in Portland, Oregon, USA. He holds a Bachelor's and Master's degree in Chemical Engineering from Arizona State University, his Master's degree focusing on Semiconductor Processing and Manufacturing. He has a specialty in Environmental Health and Safety business processes and information systems for Intel's Global Fab Materials organization. He is the holder of four patents and has four patents pending for semiconductor devices as well as waste treatment technology. He has published several industry-forum papers related to both corporate responsibility and government regulatory topics. Bob is a founding member of the EICC (Electronics Industry Citizenship Coalition) and participates in several working groups within the semiconductor and electronics industries.

Dawn E. Speranza has held various positions in the EHS field and is currently a Senior TD Materials Engineer and REACH program manager for the Intel Global Fab Materials EHS group. Prior to this, Dawn was on assignment at International SEMATECH where she managed various global chemical initiatives and EHS assessments of new materials and processes for advanced technologies. She also worked at the Intel Hudson facility in Construction Safety, Mergers and Acquisitions, Industrial Hygiene, and Occupational Safety. She has a B.S. degree from Tulane University and an M.S. degree from the University of Massachusetts at Amherst. Dawn is a certified Industrial Hygienist and Safety Professional. She currently holds the position of president elect on the SESHA (Semiconductor Environmental, Safety and Health Association) board of directors.

ENDNOTES

[1] Silicon wafers are a thin disc of pure silicon upon which up to a couple of thousand semiconductor die can be manufactured simultaneously. Afterwards, the discs are cut apart into a single die in preparation for the package-assembly and test process steps, before finally being shipped to end users.

[2] Etching is the process of removing material from predefined areas of the surface of a wafer. Deposition is the building up of material on the wafer surface.

[3] Lithography is the process of transferring an image from a pattern onto a surface by using light. In the manufacture of semiconductors, it is the process that predefines the device features on the silicon wafer prior to the etch process steps.

THIS PAGE INTENTIONALLY LEFT BLANK

For further information visit:

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