Microelectronics Reliability 52 (2012) 953–957

Contents lists available at SciVerse ScienceDirect

Microelectronics Reliability

journal homepage: www.elsevier.com/locate/microrel

Nvidia's GPU failures: A case for prognostics and health management

Michael Pecht

Prognostics and Systems Health Management Center, City University of Hong Kong, Hong Kong
CALCE Electronics Products and Systems Center, University of Maryland, College Park, MD 20742, USA

ARTICLE INFO

Article history:
Received 25 July 2011
Received in revised form 22 December 2011
Accepted 22 August 2011
Available online 24 December 2011

ABSTRACT

Prognostics and health management (PHM) is an enabling discipline consisting of technologies and methods to assess the reliability of a system in its actual life cycle conditions to determine the advent of failure and mitigate system risks. This paper presents a case study of a failure that had a significant economic impact on the computer industry and its customers, and then discusses how PHM implementation could have resulted in a dramatic difference in the outcome.

1. Introduction

On December 7, 2010, Nvidia announced in its quarterly filing with the US Securities and Exchange Commission that it had taken a cumulative net charge of $475.9 million in connection with its efforts to fix the many problems linked to its graphics processing units (GPUs) and media and communications processors (MCPs) [1]. Since at least 2007, when consumers began observing and reporting failures in their computers to companies including Hewlett–Packard, Asus, Toshiba, Dell and others, serious concerns have been raised about the effectiveness of Nvidia's reliability program and whether it caused the delays in understanding the failures, the root cause, and the most efficient and effective corrective actions [2–4]. In ensuing litigation, the claimants contended that the delays resulted in the continued production and sale of defective units [5–10]. While the actual costs to computer companies are unknown, total costs well in excess of what has been reported so far would not be unexpected.

Knowledge of Nvidia's GPU problems is available from a variety of publicly available sources, including Nvidia's own filings with the US Securities and Exchange Commission (SEC), filings with the US District Courts handling a variety of claims stemming from the GPU failures, and, to a lesser degree, online consumer complaints and journalistic reporting. The units concerned (semiconductor chips) were incorporated in many models of computers (mostly notebooks) manufactured by at least three original design manufacturers (ODMs) for 10–12 original equipment manufacturers (OEMs) [4,11]. Nvidia reported that due to a "weak die/packaging material set," the affected chips tended to fail [1,11–13]. In some news reports [14–16] the failure was attributed to cracked solder bumps, but the actual root cause has been the subject of much discussion and debate [9,17–20]. Hypotheses about the possible root cause have ranged from defects in the semiconductor device made by TSMC (a Taiwanese semiconductor company), to the underfill materials made by Namics (a Japanese company), to the substrate materials, to the bump metal layers, to package quality in assembly, and to the various "use conditions" that the units are subject to within the computers themselves [21]. In its May 21, 2010, Form 10-Q filing with the SEC, Nvidia stated that it has "not been able to determine with certainty a root cause for these failures" [22]. On June 3, 2010, Hewlett–Packard (HP) claimed in a declaration filed in court, that operation of the Nvidia part through a narrow temperature range—between 58 °C and 62 °C for some units and between 62 °C and 67 °C for other units, was the cause of failures experienced by some customers [23][1]. On June 4, 2010, attorneys for consumers suing Nvidia claimed that "the common materials and design defect in all Nvidia GPUs is the presence of Namics' 8439-1 underfill. Nvidia's choice of Namics as the underfill material was wholly inappropriate because the maximum operating temperature for GPUs far exceeds the operating temperature for NAMICS" [24]. However, it is not clear whether this indeed was the real problem. For example, HP apparently changed the BIOS, which controls the internal fan [25,26], to reduce the failure rate, which suggests it could also be a thermal management problem with their computers.

2. Background – the first signs of trouble

By 2007, consumers began reporting problems and turned to Internet postings and complaint boards to express their frustration, claiming they received slow and often inadequate responses from the product suppliers [9,27–29]. The reported problems and complaints included: no video, unexpected shut-downs, excessive heat,

[1] Most OEMs have in-built LCD and video diagnostic tests that apparently did not catch this issue. Also temperature monitoring may not have helped if the temperatures only ranged from 58 °C to 68 °C. However, since there is no record of the temperatures or the number of temperature cycles including mini-cycles of temperature, there is no way to make any assessments.

E-mail address: pecht@calce.umd.edu

0026-2714/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.microrel.2011.11.017

random characters, failure to reboot, vertical or horizontal lines, inability to recognize available wireless connection, no power, failure of LED lights, and multiple images. Many, if not most, of the problems were said to be intermittent and apparently could not be repeated on warranty returns, even when tested by the most sophisticated failure analysis companies. Such events are referred to in the industry as "no fault found" [30,31].

Due to the variety of reported problems and the claimed delay by the ODMs and OEMs in focusing their attention on Nvidia, it appears that it took several months to identify even the possible failure mechanisms that could produce the reported symptoms. During this time and even later, it appears there was insufficient data to permit corrective action. There has even been speculation as to whether there might have been a desire by the ODMs to simply push past their consumers' warranty expirations and hope for the best, but there is no evidence of this [9]. In any event, the lack of a means with which to obtain accurate real-time data (i.e., no *in situ* diagnostics or failure prognostics) undoubtedly contributed to this delayed understanding, which in turn would have led to the continued production and supply of computers with either problematic Nvidia chips or possibly poorly integrated systems containing Nvidia chips.

By the middle of 2007, it appears that HP began taking some action to investigate the growing number of field failures of computers that had Nvidia's GPUs. A joint defense team (formed to deal with the impending avalanche of lawsuits) apparently comprising HP, Nvidia, Dell, Apple and Quanta (a Taiwanese OEM) began consulting with one another, and probably with outside experts, to understand the root cause of the failures [10,23]. From mid-2007 through mid-2008 and beyond, the group was involved in many series of tests, experiments, and analyses of the affected GPUs, boards, and systems [9,10].

On July 2, 2008, Nvidia stated in a filing with the SEC that "it would take a $150–$200 million charge against the cost of revenue to cover anticipated customer warranty, repair, return, replacement and other consequential costs and expenses arising from a weak die/packaging material set in certain versions of our previous generation MCP and GPU products" [12] Nvidia was also publicly stating that it could not determine a root cause [12,22]. Nevertheless, Nvidia did state, "while we have not been able to determine a root cause for these failures, testing suggests a weak material set of die/package combination, system thermal management designs, and customer use patterns are contributing factors" [12]. In other words, Nvidia claimed that the failures should be attributed to its suppliers, the computer companies, and even the consumers. At the same time, the OEMs continued to manufacture and the ODMs continued to sell computers with Nvidia MCPs and GPUs during this lengthy and ongoing failure analysis process. HP extended its warranty an additional year to meet the growing customer complaints and apparently changed the BIOS to attempt to alleviate the problem [25,26].

Against this background, it seems clear that there was no effective reliability program of *in situ* diagnostics and prognostics and systems health management (PHM)[2]. Without this, Nvidia, the ODMs, and OEMs were left to wade through delayed, sparse, and inconsistent consumer data. There could be no uniform standard for collection and analysis of this hodgepodge of consumer data. Months went by without anyone—from component suppliers to ODMs—apparently being able to understand the patterns that existed within the data, and it seems no one discerned that those patterns were the early warning signs of a growing wave of failures about to crash upon them. The scope of the problem only continued to expand as the defective products continued to make their way into the stream of commerce during this period of delayed realization. Instead of an immediate solution, products continued to be sold, thus presenting consumers with a significant risk of varied and intermittent failure symptoms that could require months (if not years) to understand.

3. Analysis and findings

Nvidia's first announcement that it had a problem with its GPUs came on July 2, 2008, when it filed a Form 8-K with the SEC reporting that it would take a charge of $150–$200 million dollars as a result of the problematic GPUs [12]. The share price of Nvidia's stock dropped dramatically on release of that news which, as one might imagine, would not have pleased Nvidia's shareholders.

Upon learning that Nvidia had known that there was a problem with its GPUs for more than a year, Nvidia's shareholders filed multiple class action lawsuits resulting in a September 9, 2008 consolidated complaint naming as defendants Nvidia, its President and CEO Jen-Hsun Huang, and its Chief Financial Officer Marvin Burkett [7]. That lawsuit, filed in the US District Court for the Northern District of California, is still pending, although Mr. Burkett has now been dropped as a defendant. The securities class action dramatically demonstrates the potential costs of delayed understanding of a problem. If the Nvidia GPU problem had been identified earlier, perhaps Nvidia would have been able to make a public disclosure sooner and avoid the securities litigation. Effective use of PHM is the key to the early identification of data anomalies and the understanding of problems likely to ensue. Without these PHM tools, the delay in gathering, analyzing, and understanding the data will invariably lead to increased costs and loss.

In May 2010, Nvidia announced that it had taken a charge of nearly $400 million simply to satisfy the demands of its customers—the ODMs [4,22]. Nvidia has entered into settlement agreements with many, if not all, of its customers whereby it may well have agreed to bear some of the financial burden of extended warranties and other customer support services [4]. If so, it seems unlikely that these costs will reflect the complete costs Nvidia has incurred, or will incur, in dealing with consumer litigation against it, consumer litigation against its customers, or its shareholders' class action. Nvidia, the ODMs, and the OEMs have not announced the number of computers affected by the Nvidia GPUs, but, judging by the volume of complaints posted on ODM websites, Internet blogs, and complaint boards, as well as those reported in the press, it seems that a large number of computers may have been affected and the problem appears to be industry-wide. This could have contributed to the surge in consumer litigation [18,19]. In late 2008 and early 2009, approximately 10 consumer class actions were filed in courts around the US alleging that various ODMs, as well as Nvidia, breached statutory consumer warranties and are liable under theories of strict liability, negligence, and breach of statutory consumer protection laws [24]. The US District Court handling the consolidated consumer class actions arising out of the Nvidia GPU failures upheld the adequacy of the certain claims made directly against Nvidia, even though Nvidia never dealt directly with the consumers bringing the litigation [32]. In sum, the Court held that a consumer may sue a component supplier directly!

The ODMs and Nvidia have long sought to deflect the consumer litigation by arguing that even if they knew that computers containing Nvidia's GPUs were likely to fail prematurely, the consumers' remedies were limited to the terms of the ODMs'

[2] Prognostics is the process of assessing the extent of deviation or degradation of a product from its expected normal operating conditions, and then, based on continuous monitoring, predicting the future reliability of the product. If one can monitor key performance and environmental parameters, then this data can be used in conjunction with precursor reasoning algorithms and stress-and-damage models to enable prognostics. By being able to determine when a product will fail, procedures can be developed to provide advance warning of failures, reduce life cycle costs, and improve the design and qualification of fielded and future systems.