

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE**

VLSI TECHNOLOGY LLC,

Plaintiff,

v.

INTEL CORPORATION,

Defendant.

C.A. No. _____

JURY TRIAL DEMANDED

VLSI TECHNOLOGY LLC'S COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff VLSI Technology LLC ("VLSI"), by and through its undersigned counsel, pleads the following against Intel Corporation ("Intel") and alleges as follows:

THE PARTIES

1. Plaintiff VLSI is a Delaware limited liability company duly organized and existing under the laws of the State of Delaware. The address of the registered office of VLSI is Corporation Trust Center, 1209 Orange St., Wilmington, DE 19801. The name of VLSI's registered agent at that address is The Corporation Trust Company.

2. VLSI is the assignee and owns all right, title, and interest to U.S. Patent Nos. 6,366,522 ("the '522 Patent"), 6,633,187 ("the '187 Patent"), 7,292,485 ("the '485 Patent"),

7,606,983 (“the ‘983 Patent”), 7,725,759 (“the ‘759 Patent”) and 7,793,025 (“the ‘025 Patent”) (collectively, the “Asserted Patents”).

3. On information and belief, Defendant Intel is a corporation duly organized and existing under the laws of the State of Delaware, having its principal place of business at 2200 Mission College Blvd., Santa Clara, CA 95054.

JURISDICTION AND VENUE

4. This is an action arising under the patent laws of the United States, 35 U.S.C. § 1 *et seq.* Accordingly, this Court has subject matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

5. This Court has personal jurisdiction over Intel because Intel is incorporated in Delaware. Intel also manufactures products that are and have been used, offered for sale, sold, and purchased in the District of Delaware.

6. Under 28 U.S.C. §§ 1391(b)-(d) and 1400(b), venue is proper in this judicial district because Intel is incorporated in this district, has committed acts of infringement within this judicial district giving rise to this action, and does business in this district.

FIRST CLAIM

(Infringement of U.S. Patent No. 6,366,522)

7. VLSI re-alleges and incorporates herein by reference Paragraphs 1-6 of its Complaint.

8. The ‘522 Patent, entitled “Method and apparatus for controlling power consumption of an integrated circuit,” was duly and lawfully issued on April 2, 2002. A true and correct copy of the ‘522 Patent is attached hereto as Exhibit 1.

9. The ‘522 Patent names Marcus W. May and Daniel Mulligan as co-inventors.

10. The '522 Patent has been in full force and effect since its issuance. VLSI owns by assignment the entire right, title, and interest in and to the '522 Patent, including the right to seek damages for past, current, and future infringement thereof.

11. The '522 Patent "relates generally to integrated circuits and more particularly to controlling power consumption by an integrated circuit." Ex. 1 at 1:6-8.

12. The '522 Patent explains, for example, that "a need exists for a method and apparatus that adjust[s] the system clock and/or the supply voltage based on the processing capabilities of an integrated circuit and the application being performed to conserve power." *Id.* at 1:45-48.

13. The '522 Patent further explains that it "provides a method and apparatus for controlling power consumption of an integrated circuit." *Id.* at 2:7-9.

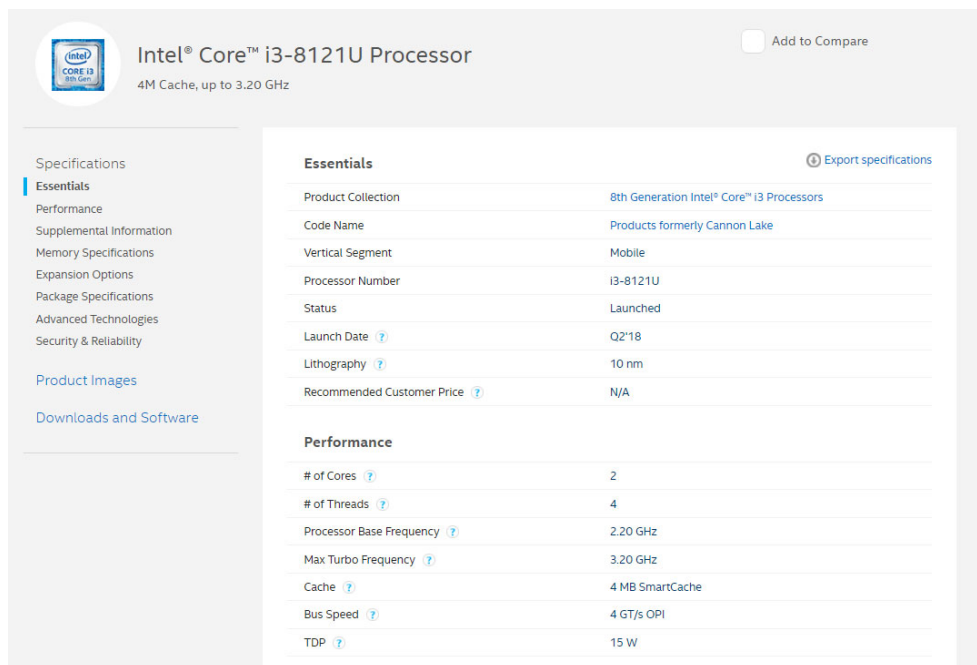
14. VLSI is informed and believes, and thereon alleges, that Intel has infringed and unless enjoined will continue to infringe one or more claims of the '522 Patent, in violation of 35 U.S.C. § 271, by, among other things, making, using, offering to sell, and selling within the United States, supplying or causing to be supplied in or from the United States, and importing into the United States, without authority or license, Intel products that use "Speed Shift" technology with a fully integrated voltage regulator ("FIVR") in an infringing manner.

15. For example, the '522 accused products embody every limitation of at least claim 9 of the '522 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

["9. A method for controlling power consumption of an integrated circuit, the method comprises the steps of:"]

16. Intel Cannon Lake processors are operated using a method of controlling power consumption of an integrated circuit.

17. For example, Intel Cannon Lake processors include an integrated circuit. *See, e.g.*, <https://ark.intel.com/products/136863/Intel-Core-i3-8121U-Processor-4M-Cache-up-to-3-20-GHz-?q=Core%20i3-8121U> [hereinafter “ARK”].



Intel® Core™ i3-8121U Processor
4M Cache, up to 3.20 GHz

Specifications

- Essentials
- Performance
- Supplemental Information
- Memory Specifications
- Expansion Options
- Package Specifications
- Advanced Technologies
- Security & Reliability
- Product Images
- Downloads and Software

Essentials

Product Collection	8th Generation Intel® Core™ i3 Processors
Code Name	Products formerly Cannon Lake
Vertical Segment	Mobile
Processor Number	i3-8121U
Status	Launched
Launch Date	Q2'18
Lithography	10 nm
Recommended Customer Price	N/A

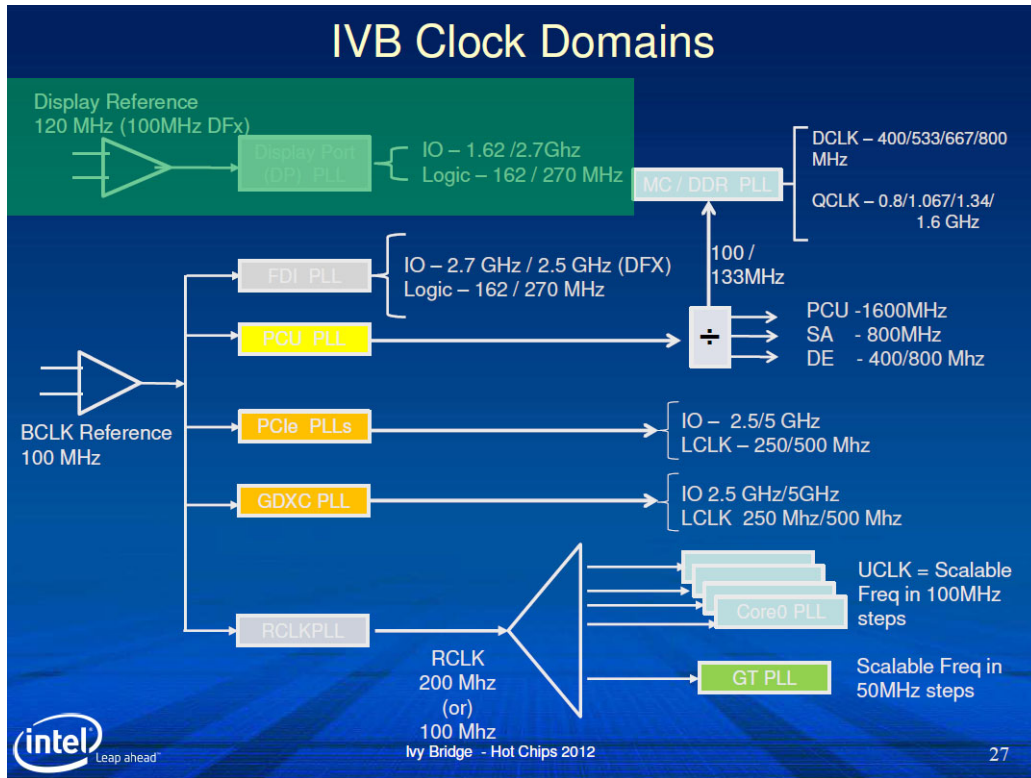
Performance

# of Cores	2
# of Threads	4
Processor Base Frequency	2.20 GHz
Max Turbo Frequency	3.20 GHz
Cache	4 MB SmartCache
Bus Speed	4 GT/s OPI
TDP	15 W

[“producing a system clock from a reference clock based on a system clock control signal;”]

18. Intel Cannon Lake processors are operated using a method that comprises producing a system clock from a reference clock based on a system clock control signal.

19. For example, the Intel power control architecture produces a number of controllably scalable system clocks, including core clocks, from a reference clock based on system clock control signals. *See, e.g.*, Power Management of the Third Generation Core Micro Architecture formerly codenamed Ivy Bridge [hereinafter “PM”] at 27.



20. On information and belief, these and the other pertinent portions of Intel’s Ivy Bridge generation (as well as the Sandy Bridge and 4th Generation Intel Core SoCs discussed below) were carried over to the Cannon Lake processor in a manner that is materially the same with respect to the infringement analysis presented in this example.

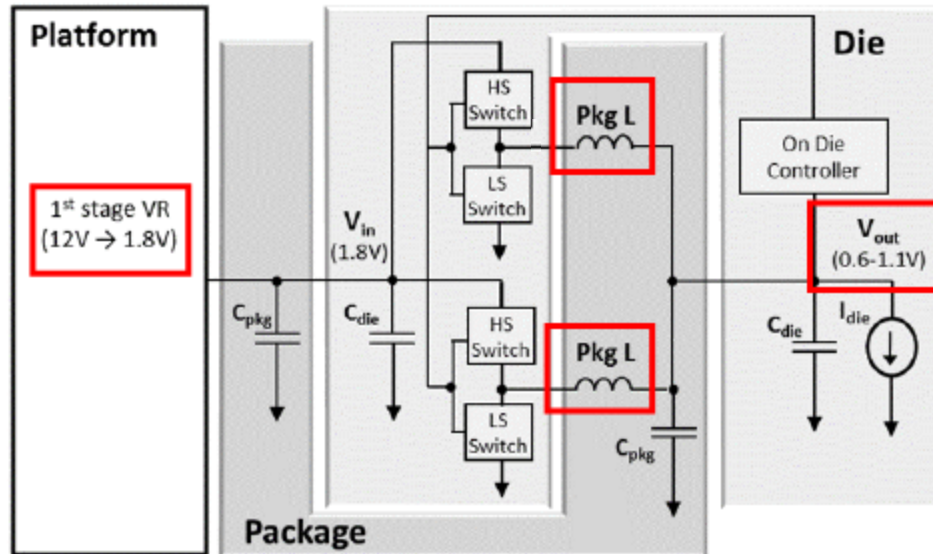
21. As a further example, reverse engineering of an Intel Cannon Lake processor also shows the production of a number of controllably scalable system clocks, including core clocks, from a reference clock based on system clock control signals.

Processors Information	
Socket 1	ID = 0
Number of cores	2 (max 2)
Number of threads	4 (max 4)
Name	Intel Cannon Lake
Codename	Skylake
Specification	Intel(R) Core(TM) i3-8121U CPU @ 2.20GHz
Package (platform ID)	(0x7)
CPUID	6.6.3
Extended CPUID	6.66
Core Stepping	
TDP Limit	15.0 Watts
Tjmax	105.0 °C
Core Speed	3094.7 MHz
Multiplier x Bus Speed	31.0 x 99.8 MHz
Base frequency (cores)	99.8 MHz
Base frequency (ext.)	99.8 MHz
Stock frequency	2200 MHz
Max frequency	3200 MHz

["regulating at least one supply from a power source and an inductance based on a power supply control signal; and"]

22. Intel Cannon Lake processors are operated using a method that comprises regulating at least one supply from a power source and an inductance based on a power supply control signal.

23. For example, on information and belief the Fully Integrated Voltage Regulator (FIVR) on Intel Cannon Lake Processors regulates an output supply (*e.g.*, V_{out}) from a power source (*e.g.*, 1st stage VR) and an inductance (*e.g.*, Pkg L) based on power supply control signals. *See, e.g.*, “FIVR – Fully Integrated Voltage Regulators on the 4th Generation Intel Core SoCs” [hereinafter “FIVR”] at 2.



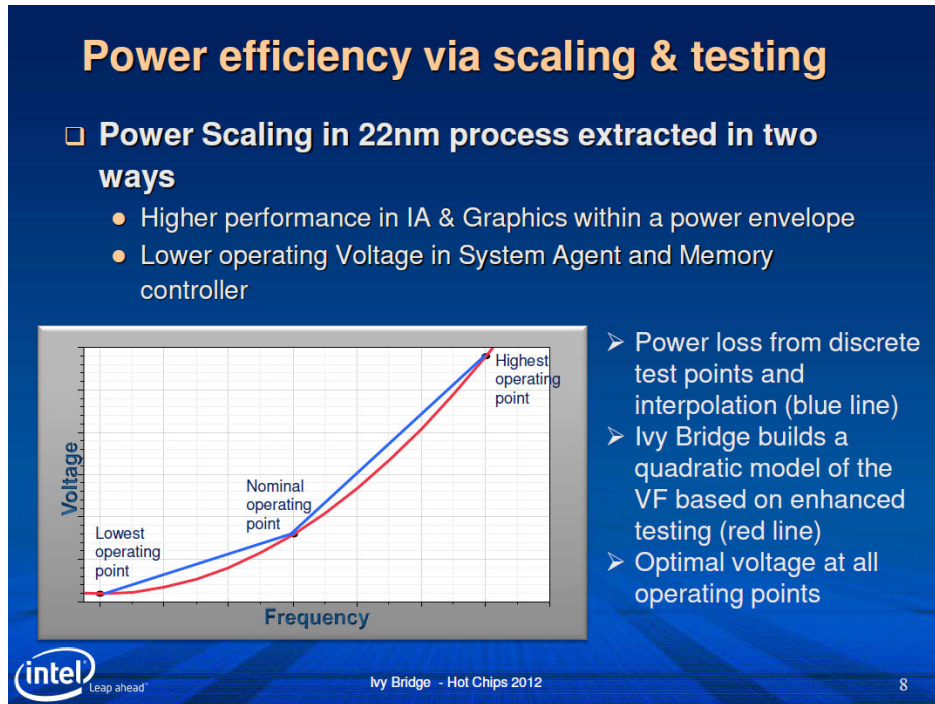
24. See also, e.g., FIVR at 1: “Each FIVR is independently programmable to achieve optimal operation given the requirements of the domain it is powering. The settings are optimized by the Power Control Unit (PCU), which specifies the input voltage, output voltage, number of operating phases, and a variety of other settings to minimize the total power consumption of the die.”

25. See also, e.g., <https://forums.anandtech.com/threads/fivr-returning-with-cannonlake.2484038/#lg=xfUId-1-1549490270&slide=0> (discussing “FIVR at uncore for Cannonlake SoC” and “Cannonlake architecture ... with FIVR”).

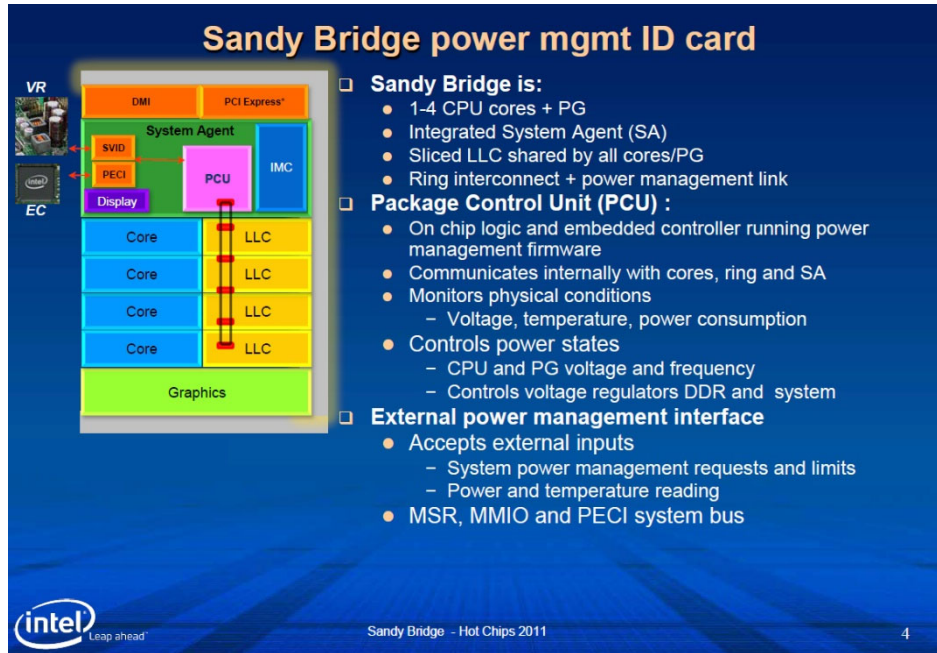
[“producing the system clock control signal and the power supply control signal based on a processing transfer characteristic of a computation engine and “]

26. Intel Cannon Lake processors are operated using a method that comprises producing the system clock control signal and the power supply control signal based on a processing transfer characteristic of a computation engine.

27. For example, the PCU produces power supply control signals based on a quadratic model of the VF requirements, a model that includes one or more processing transfer characteristics of the computation engine. *See, e.g.,* PM at 8.



28. *See also, e.g.,* “Power management architecture of the 2nd generation Intel Core microarchitecture, formerly codenamed Sandy Bridge” at 4.



29. Additionally, the PCU produces system clock control signals based on maximum frequencies, maximum turbo frequencies and TDP, additional processing transfer characteristics of the computational engine. *See, e.g., ARK.*

Performance	
# of Cores ?	2
# of Threads ?	4
Processor Base Frequency ?	2.20 GHz
Max Turbo Frequency ?	3.20 GHz
Cache ?	4 MB SmartCache
Bus Speed ?	4 GT/s OPI
TDP ?	15 W

[“processing requirements associated with processing at least a portion of an application by the computation engine.”]

30. Intel Cannon Lake processors are operated using a method that comprises producing the system clock control signal and the power supply control signal based on processing requirements associated with processing at least a portion of an application by the computation engine.

31. For example, the processor P-state, which includes an operating frequency and voltage, is set by the PCU and the Hardware-Controlled Performance State system based on processing requirements deemed appropriate for the applied workload. *See, e.g.*, Intel Software Developer’s Manual [hereinafter “SDM”] at 3158-3159.

14.4 HARDWARE-CONTROLLED PERFORMANCE STATES (HWP)

Intel processors may contain support for Hardware-Controlled Performance States (HWP), which autonomously selects performance states while utilizing OS supplied performance guidance hints. The Enhanced Intel Speed-Step® Technology provides a means for the OS to control and monitor discrete frequency-based operating points via the IA32_PERF_CTL and IA32_PERF_STATUS MSRs.

In contrast, HWP is an implementation of the ACPI-defined Collaborative Processor Performance Control (CPPC), which specifies that the platform enumerates a continuous, abstract unit-less, performance value scale that is not tied to a specific performance state / frequency by definition. While the enumerated scale is roughly linear in terms of a delivered integer workload performance result, the OS is required to characterize the performance value range to comprehend the delivered performance for an applied workload.

When HWP is enabled, the processor autonomously selects performance states as deemed appropriate for the applied workload and with consideration of constraining hints that are programmed by the OS. These OS-provided hints include minimum and maximum performance limits, preference towards energy efficiency or performance, and the specification of a relevant workload history observation time window. The means for the OS to override HWP’s autonomous selection of performance state with a specific desired performance target is also provided, however, the effective frequency delivered is subject to the result of energy efficiency and performance optimizations.

32. Reverse engineering analysis shows that Intel Cannon Lake processors include the HWP feature.

Processors Information	
Socket 1	ID = 0
Number of cores	2 (max 2)
Number of threads	4 (max 4)
Name	Intel Cannon Lake
Codename	Skylake
Specification	Intel(R) Core(TM) i3-8121U CPU @ 2.20GHz
Package (platform ID)	(0x7)
CPUID	6.6.3
Extended CPUID	6.66

[...]

IBRS	supported, disabled
IBPB	supported
STIBP	supported, disabled
RDCL_NO	no
IBRS_ALL	not supported
Turbo Mode	supported, enabled
Max non-turbo ratio	22x
Max turbo ratio	32x
Max efficiency ratio	4x
Speedshift	Autonomous

33. Intel has had knowledge of the ‘522 Patent at least since the filing of this Complaint, and if it did not have actual knowledge prior to that time, it was willfully blind to the existence of the ‘522 Patent based on, for example, its publicly-known corporate policy forbidding its employees from reading patents held by outside companies or individuals. For example, in *Intel Corp. v. Future Link Sys., LLC*, 268 F. Supp. 3d 605, 623 (D. Del. 2017) the court noted the patent owner’s observation that “Intel’s own engineers concede that they avoid reviewing other, non-Intel patents so as to avoid willfully infringing them.” As a further example, former Intel employees, including Intel’s long-time Chief Architect Robert Colwell, have admitted that this policy’s purpose is to “avoid possible triple damages for ‘willful infringement.’” As still another example, on information and belief, Intel has been sued for infringing patents previously assigned to NXP while this policy was in place, including for infringing a patent naming Marcus W. May (also an inventor on the ‘522 Patent) as an inventor. *See, e.g., VLSI v. Intel Corporation*, Civil Action No. 18-0966-CFC (D. Del.). Yet despite this notice, Intel proceeded to infringe other patents on inventions developed in the same area by Mr. May. Intel should have known that its conduct was infringing both prior to and following the filing of this case.

34. VLSI is informed and believes, and thereon alleges, that Intel actively, knowingly, and intentionally has induced infringement of the ‘522 Patent by, for example, controlling the

design and manufacture of, offering for sale, selling, supplying, and otherwise providing instruction and guidance regarding the above-described products with the knowledge and specific intent to encourage and facilitate infringing uses of such products by its customers both inside and outside the United States (as used in this pleading, “customers” refers to both direct and indirect customers, including entities that distribute and resell the accused products, alone or as part of a system, and end users of such products and systems). For example, Intel publicly provides documentation, including datasheets available through Intel’s publicly accessible ARK service and software developer’s manuals, instructing customers on uses of Intel’s products that infringe the methods of the ‘522 Patent. *See, e.g.,* <http://ark.intel.com>. On information and belief, Intel’s customers directly infringe the ‘522 Patent by, for example, making, using, offering to sell, and selling within the United States, and importing into the United States, without authority or license, products containing the above-described Intel products.

35. VLSI is informed and believes, and thereon alleges, that Intel has contributed to the infringement by its customers of the ‘522 Patent by, without authority, importing, selling and offering to sell within the United States materials and apparatuses for practicing the claimed invention of the ‘522 Patent both inside and outside the United States. For example, the above-described products constitute a material part of the inventions of the ‘522 Patent and are not staple articles or commodities of commerce suitable for substantial noninfringing use. On information and belief, Intel knows that the above-described products constitute a material part of the inventions of the ‘522 Patent and are not staple articles or commodities of commerce suitable for substantial noninfringing use. On information and belief, Intel’s customers directly infringe the ‘522 Patent by, for example, making, using, offering to sell, and selling within the United States,

and importing into the United States, without authority or license, products containing the above-described Intel products.

36. As a result of Intel's infringement of the '522 Patent, VLSI has been damaged. VLSI is entitled to recover for damages sustained as a result of Intel's wrongful acts in an amount subject to proof at trial.

37. To the extent 35 U.S.C. § 287 is determined to be applicable, on information and belief its requirements have been satisfied with respect to the '522 Patent.

38. In addition, Intel's infringing acts and practices have caused and are causing immediate and irreparable harm to VLSI.

39. VLSI is informed and believes, and thereon alleges, that Intel's infringement of the '522 Patent has been and continues to be willful. As noted above, Intel has had knowledge of the '522 Patent, or, at minimum, has been willfully blind to the existence of the '522 Patent. Moreover, after service of this Complaint, Intel undisputedly had knowledge of its infringement of the '522 Patent. Intel has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for VLSI's patent rights. Thus, Intel's infringing actions have been and continue to be consciously wrongful.

40. Based on the information alleged in this claim, as well as the information alleged in the Second Claim through the Sixth Claim *infra*, VLSI is informed and believes, and thereon alleges, that this is an exceptional case, which warrants an award of attorney's fees to VLSI pursuant to 35 U.S.C. § 285.

SECOND CLAIM

(Infringement of U.S. Patent No. 6,633,187)

41. VLSI re-alleges and incorporates herein by reference Paragraphs 1-40 of its Complaint.

42. The '187 Patent, entitled "Method and apparatus for enabling a stand alone integrated circuit," was duly and lawfully issued October 14, 2003. A true and correct copy of the '187 Patent is attached hereto as Exhibit 2.

43. The '187 Patent names Michael R. May and Marcus W. May as inventors.

44. The '187 Patent has been in full force and effect since its issuance. VLSI owns by assignment the entire right, title, and interest in and to the '187 Patent, including the right to seek damages for past, current, and future infringement thereof.

45. The '187 Patent states that it "relates generally to integrated circuits and more particularly to enabling a stand-alone integrated circuit." Ex. 2 at 1:7-9.

46. The '187 Patent states that integrated circuits "include a large amount of circuitry in a very small area." *Id.* at 1:12-13. The '187 Patent further explains that when a "power converter is on-chip with the digital circuitry and the power converter requires a clock signal to produce a supply voltage, a difficulty arises in enabling such a stand-alone integrated circuit." *Id.* at 1:34-39. Thus, the patent explains, there was a need "for a method and apparatus for enabling a stand-alone integrated circuit that includes an on-chip power converter." *Id.* at 1:40-42.

47. VLSI is informed and believes, and thereon alleges, that Intel has infringed and unless enjoined will continue to infringe one or more claims of the '187 Patent, in violation of 35 U.S.C. § 271, by, among other things, making, using, offering to sell, and selling within the United States, supplying or causing to be supplied in or from the United States, and importing into the

United States, without authority or license, Intel products containing an infringing fully integrated voltage regulator (“FIVR”).

48. For example, the ‘187 accused products embody every limitation of at least claim 1 of the ‘187 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

[“1. A method for enabling a stand-alone integrated circuit (IC), the method comprises the steps of:”]

49. Intel Broadwell processors use a method for enabling a stand-alone integrated circuit.

50. For example, Intel Broadwell processors comprise a stand-alone integrated circuit that includes a FIVR. *See, e.g.*, 5th Generation Intel Core Processor Family Datasheet Vol. 1 [hereinafter “Datasheet”] at 80.

7.1 Integrated Voltage Regulator

A feature to the processor is the integration of platform voltage regulators into the processor. Due to this integration, the processor has one main voltage rail (V_{CC}) and a voltage rail for the memory interface (V_{DDQ}), compared to six voltage rails on previous processors. The V_{CC} voltage rail will supply the integrated voltage regulators which in turn will regulate to the appropriate voltages for the cores, cache, system agent, and graphics. This integration allows the processor to better control on-die voltages to optimize between performance and power savings. The processor V_{CC} rail will remain a VID-based voltage with a loadline similar to the core voltage rail (also called V_{CC}) in previous processors.

[“a) establishing an idle state that holds at least a portion of the stand-alone IC in a reset condition when a power source is operably coupled to the stand-alone IC;”]

51. Intel Broadwell processors use a method that comprises establishing an idle state that holds at least a portion of the stand-alone IC in a reset condition when a power source is operably coupled to the stand-alone IC.

52. For example, Broadwell processors initialize into an idle state that holds the cores in a reset condition, preventing them from running while a power source is operably coupled to the stand-alone IC until the correct power sequencing signals are received. *See, e.g.*, Datasheet at 77.

6.8 Power Sequencing Signals

Table 33. Power Sequencing Signals

Signal Name	Description	Direction / Buffer Type
PROCPWRGD	The processor requires this input signal to be a clean indication that the V _{CC} and V _{DDQ} power supplies are stable and within specifications. This requirement applies regardless of the S-state of the processor. 'Clean' implies that the signal will remain low (capable of sinking leakage current), without glitches, from the time that the power supplies are turned on until the supplies come within specification. The signal must then transition monotonically to a high state.	I Asynchronous CMOS
VCCST_PWRGD	The processor requires this input signal to be a clean indication that the V _{CCST} and V _{DDQ} power supplies are stable and within specifications. This signal must have a valid level during both S0 and S3 power states. 'Clean' implies that the signal will remain low (capable of sinking leakage current), without glitches, from the time that the power supplies are turned on until the supplies come within specification. The signal must then transition monotonically to a high state."	I Asynchronous CMOS
PROC_DETECT#	Processor Detect: This signal is pulled down directly (0 Ohms) on the processor package to ground. There is no connection to the processor silicon for this signal. System board designers may use this signal to determine if the processor is present.	—

["b) receiving a power enable signal;"]

53. Intel Broadwell processors use a method that comprises receiving a power enable signal.

54. For example, Broadwell processors receive a power enable signal such as PROCPWRGD. *See, e.g.*, Datasheet at 77.

6.8 Power Sequencing Signals

Table 33. Power Sequencing Signals

Signal Name	Description	Direction / Buffer Type
PROCPWRGD	The processor requires this input signal to be a clean indication that the V _{CC} and V _{DDQ} power supplies are stable and within specifications. This requirement applies regardless of the S-state of the processor. 'Clean' implies that the signal will remain low (capable of sinking leakage current), without glitches, from the time that the power supplies are turned on until the supplies come within specification. The signal must then transition monotonically to a high state.	I Asynchronous CMOS
VCCST_PWRGD	The processor requires this input signal to be a clean indication that the V _{CCST} and V _{DDQ} power supplies are stable and within specifications. This signal must have a valid level during both S0 and S3 power states. 'Clean' implies that the signal will remain low (capable of sinking leakage current), without glitches, from the time that the power supplies are turned on until the supplies come within specification. The signal must then transition monotonically to a high state."	I Asynchronous CMOS
PROC_DETECT#	Processor Detect: This signal is pulled down directly (0 Ohms) on the processor package to ground. There is no connection to the processor silicon for this signal. System board designers may use this signal to determine if the processor is present.	—

[“c) enabling, in response to the power enable signal, an on-chip power converter of the stand-alone IC to generate at least one supply from the power source, wherein the enabling includes:”]

55. Intel Broadwell processors use a method that comprises enabling, in response to the power enable signal, an on-chip power converter of the stand-alone IC to generate at least one supply from the power source.

56. For example, Broadwell processors are started in response to the power enable signals, including PROCPWRGD. The system’s Power Control Unit (PCU) then turns on or off given “rails,” or regions of the chip powered by a particular on-chip power converter. These on-chip power converters then provide the correct supply voltages.

57. See, e.g., “FIVR – Fully Integrated Voltage Regulators on 4th Generation Intel Core SoCs” [hereinafter “FIVR”] at 4.

C. System Control

In order to minimize losses from FIVR, a modified version of the PCU [1] dynamically configures each FCM based on the current activity level of the domain. The PCU turns each rail on or off based on activity, and specifies an output voltage target to support the desired frequency. It also optimizes the settings discussed in section II.A for the anticipated operating conditions. These settings include the number of active phases (i.e. phase shedding to improve light load performance), the compensator settings (to maintain optimal transient response as the number of phases changes), and the timing of switch drivers (to ensure zero voltage switching at light loads). This allows each FIVR domain to operate at near peak efficiency across a wide range of load conditions from retention to Turbo. An example of the benefit this provides is shown in Section IV.A.

58. On information and belief, these and the other pertinent portions of the FIVR in 4th Generation Intel Core chips were carried over to the Broadwell processor in a manner that is materially the same with respect to the infringement analysis presented in this example.

["generating a clock signal;" and "generating power converter regulation signals based on the clock signal;"]

59. Intel Broadwell processors use a method that comprises generating a clock signal and generating power converter regulation signals based on the clock signal.

60. For example, the Frequency Control Module in each FIVR domain of the Intel Broadwell processor generates a clock signal using, *e.g.*, a "triangular waveform synthesizer" PLL and generates power converter regulation signals, for instance, in a comparator based on this clock signal. *See, e.g.*, FIVR at 3.

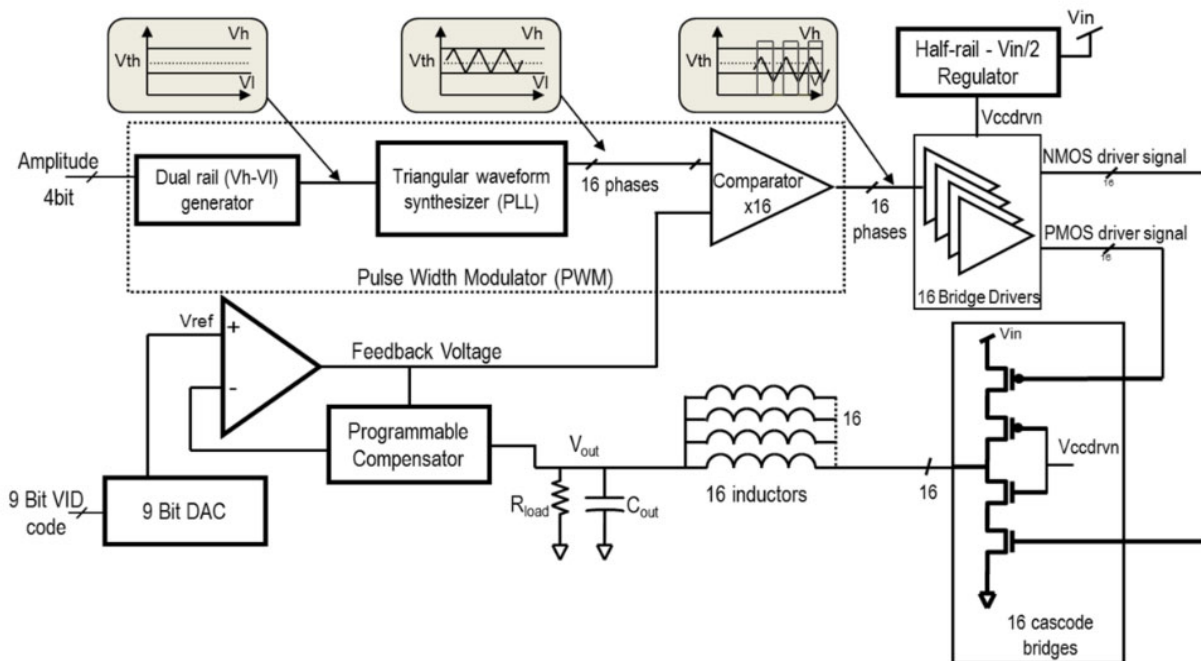


Figure 2. Simplified block diagram of the circuitry for a single representative FIVR domain

["enabling a band-gap reference that is used in generating the power converter regulation signals; and"]

61. Intel Broadwell processors use a method that comprises enabling a band-gap reference that is used in generating the power converter regulation signals.


62. For example, Broadwell processors, which are manufactured in 14-nm CMOS technology, on information and belief use a band-gap reference (BGREF) in generating their power converter regulation signals by measuring “the digital VCC power-up.” See, e.g., “An Accurate Bandgap-Based Power-on-Detector in 14-nm CMOS Technology” [hereinafter “Bandgap”] at 1.

An Accurate Bandgap-Based Power-on-Detector in 14-nm CMOS Technology

Joseph Shor, *Senior Member, IEEE*, and Dror Zilberman

Abstract—A novel power-on-detector (POD) circuit is reported in 14-nm technology. This POD has a sigma accuracy of 12 mV, including process, voltage, and temperature variations with a power of 496 μ W and an area of 0.02 mm². The POD uses an open-loop offset-cancelled bandgap-based circuit to achieve this high accuracy. It is also capable of detecting two operating voltages simultaneously and can measure the digital VCC power-up at supply levels lower than the diode voltage.

Index Terms—Analog, bandgap reference (BGREF), micro-processors, power-on-detector (POD) circuit.



I. INTRODUCTION

63. See also, e.g., FIVR at 3.

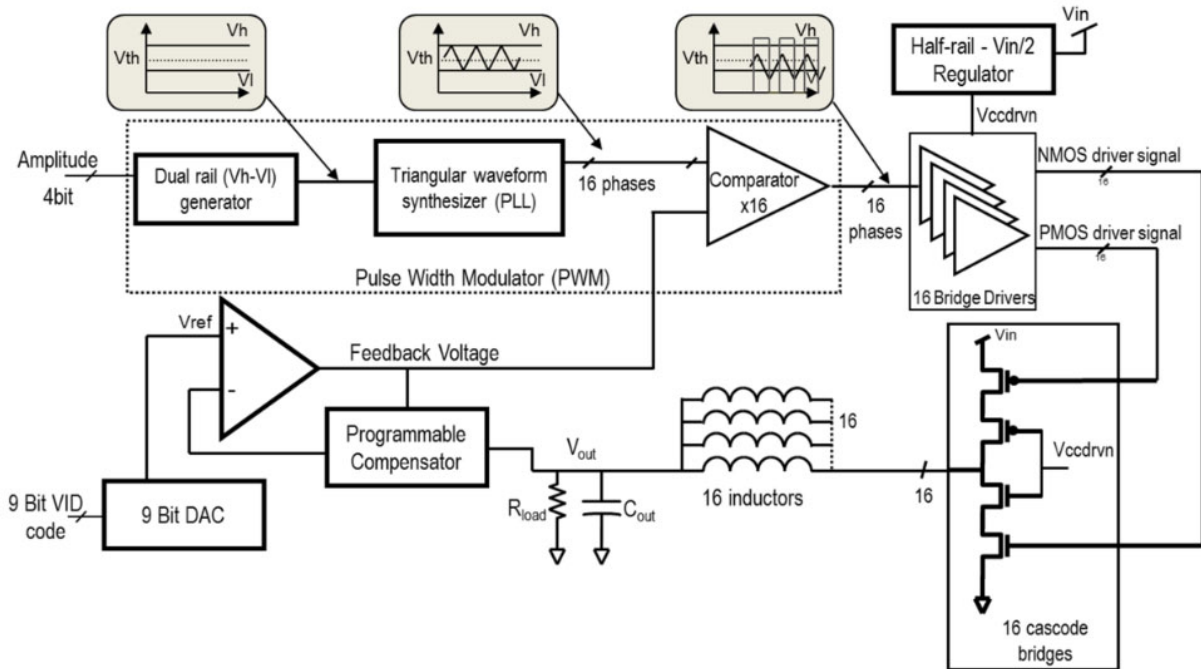


Figure 2. Simplified block diagram of the circuitry for a single representative FIVR domain

64. See also, e.g., Bandgap at 4, stating that this circuit was produced “extensively in high-volume manufacturing.”

["d) when the at least one supply has substantially reached a steady-state condition, enabling functionality of the stand-alone IC."]

65. Intel Broadwell processors use a method that comprises, when the at least one supply has substantially reached a steady-state condition, enabling functionality of the stand-alone IC.

66. For example, Broadwell processors use a “power-on detector” to enable functionality of the IC when the supply voltages have reached their steady state operating level.

See, e.g., Bandgap at 1.

COMPUTER systems require power-on-detector (POD) circuits to determine that the supply voltages have reached the correct operating level [1]–[5]. This is especially true when memory elements are involved. In present microprocessors, there are usually multiple supply domains, both analog and digital, each of which can have its own operating point [6], [7]. The accuracy of the POD will determine the minimum VCC level of the microprocessor or system-on-chip, so this is an important parameter.

67. Intel has had knowledge of the ‘187 Patent at least since the filing of this Complaint, and if it did not have actual knowledge prior to that time, it was willfully blind to the existence of the ‘187 Patent based on, for example, its publicly-known corporate policy forbidding its employees from reading patents held by outside companies or individuals, as already described above. As still another example, on information and belief, Intel has been sued for infringing patents previously assigned to NXP while this policy was in place, including for infringing a patent naming Marcus W. May (also an inventor on the ‘187 Patent) as an inventor, as also already explained above. Yet despite this notice, Intel proceeded to infringe other patents on inventions

developed in the same area by Mr. May. Intel should have known that its conduct was infringing both prior to and following the filing of this case.

68. VLSI is informed and believes, and thereon alleges, that Intel actively, knowingly, and intentionally has induced infringement of the '187 Patent by, for example, controlling the design and manufacture of, offering for sale, selling, supplying, and otherwise providing instruction and guidance regarding the above-described products with the knowledge and specific intent to encourage and facilitate infringing uses of such products by its customers both inside and outside the United States. For example, Intel publicly provides documentation, including datasheets available through Intel's publicly accessible ARK service and software developer's manuals, instructing customers on uses of Intel's products that infringe the methods of the '187 Patent. *See, e.g.,* <http://ark.intel.com>. On information and belief, Intel's customers directly infringe the '187 Patent by, for example, making, using, offering to sell, and selling within the United States, and importing into the United States, without authority or license, products containing the above-described Intel products.

69. VLSI is informed and believes, and thereon alleges, that Intel has contributed to the infringement by its customers of the '187 Patent by, without authority, importing, selling and offering to sell within the United States materials and apparatuses for practicing the claimed invention of the '187 Patent both inside and outside the United States. For example, the above-described products constitute a material part of the inventions of the '187 Patent and are not staple articles or commodities of commerce suitable for substantial noninfringing use. On information and belief, Intel knows that the above-described products constitute a material part of the inventions of the '187 Patent and are not staple articles or commodities of commerce suitable for substantial noninfringing use. On information and belief, Intel's customers directly infringe the

'187 Patent by, for example, making, using, offering to sell, and selling within the United States, and importing into the United States, without authority or license, products containing the above-described Intel products.

70. As a result of Intel's infringement of the '187 Patent, VLSI has been damaged. VLSI is entitled to recover for damages sustained as a result of Intel's wrongful acts in an amount subject to proof at trial.

71. To the extent 35 U.S.C. § 287 is determined to be applicable, on information and belief its requirements have been satisfied with respect to the '187 Patent.

72. In addition, Intel's infringing acts and practices have caused and are causing immediate and irreparable harm to VLSI.

73. VLSI is informed and believes, and thereon alleges, that Intel's infringement of the '187 Patent has been and continues to be willful. As noted above, Intel has had knowledge of the '187 Patent, or, at minimum, has been willfully blind to the existence of the '187 Patent. Moreover, after service of this Complaint, Intel undisputedly had knowledge of its infringement of the '187 Patent. Intel has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for VLSI's patent rights. Thus, Intel's infringing actions have been and continue to be consciously wrongful.

THIRD CLAIM

(Infringement of U.S. Patent No. 7,292,485)

74. VLSI re-alleges and incorporates herein by reference Paragraphs 1-73 of its Complaint.

75. The '485 Patent, entitled "SRAM having variable power supply and method therefor," was duly and lawfully issued November 6, 2007. A true and correct copy of the '485 Patent is attached hereto as Exhibit 3.

76. The '485 Patent names Olga R. Lu, Lawrence F. Childs, and Craig D. Gunderson as co-inventors.

77. The '485 Patent has been in full force and effect since its issuance. VLSI owns by assignment the entire right, title, and interest in and to the '485 Patent, including the right to seek damages for past, current, and future infringement thereof.

78. The '485 Patent states that it "relates generally to memories, and more particularly, to a static random access (SRAM) memory having a variable power supply and method therefor." Ex. 3 at 1:6-9.

79. The '485 Patent explains that "[s]tatic random access memories (SRAMs) are generally used in applications requiring high speed, such as memory in a data processing system." *Id.* at 1:13-15. The '485 Patent explains "increasing [certain] ratios improves cell stability. However, improving stability comes at the expense of lower write performance." *Id.* at 1:36-38. The patent further explains, "there is a need for a SRAM having improved cell stability while also having improved write margins." *Id.* at 1:42-43.

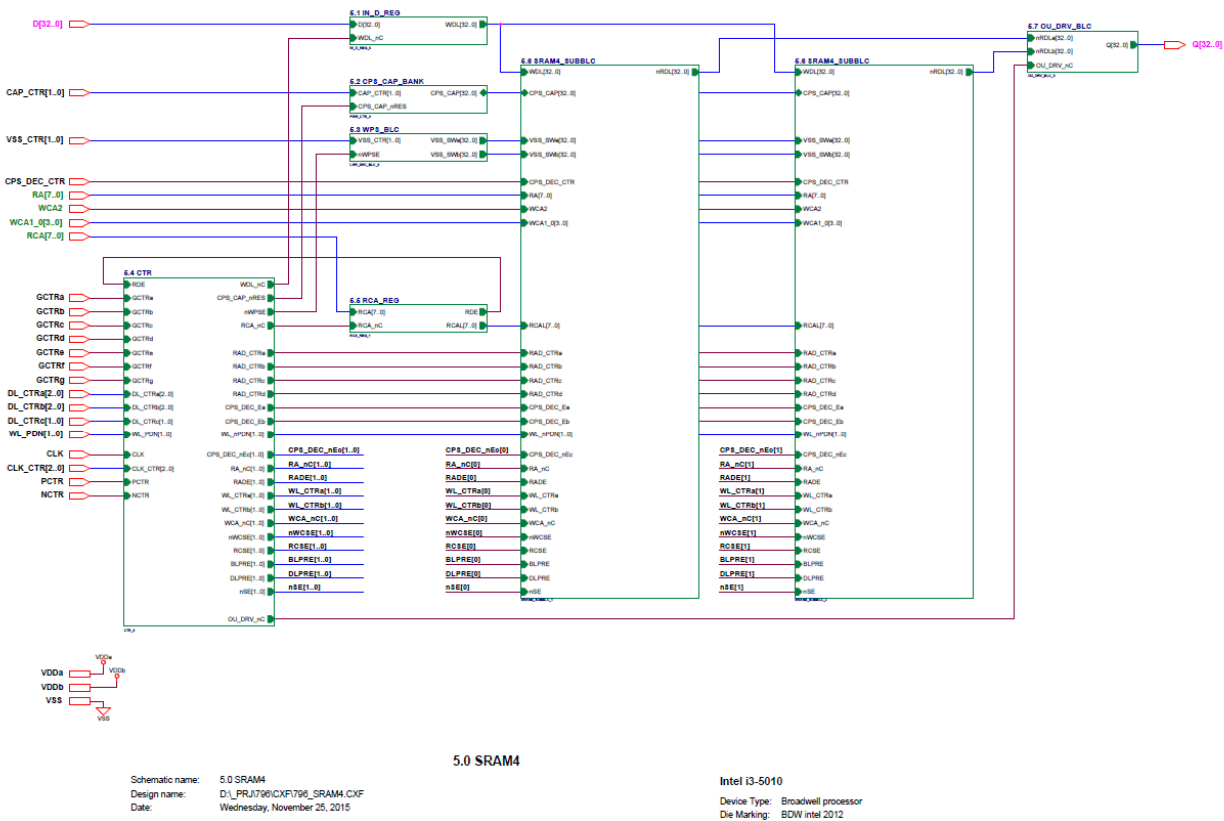
80. VLSI is informed and believes, and thereon alleges, that Intel has infringed and unless enjoined will continue to infringe one or more claims of the '485 Patent, in violation of 35 U.S.C. § 271, by, among other things, making, using, offering to sell, and selling within the United States, supplying or causing to be supplied in or from the United States, and importing into the United States, without authority or license, Intel products that use infringing write-assist technology in static random access memory ("SRAM") arrays.

81. For example, the ‘485 accused products embody every limitation of at least claim 12 of the ‘485 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

“12. A method, comprising: providing a memory comprising:”

82. Intel Broadwell processors operate using a method comprising providing a memory.

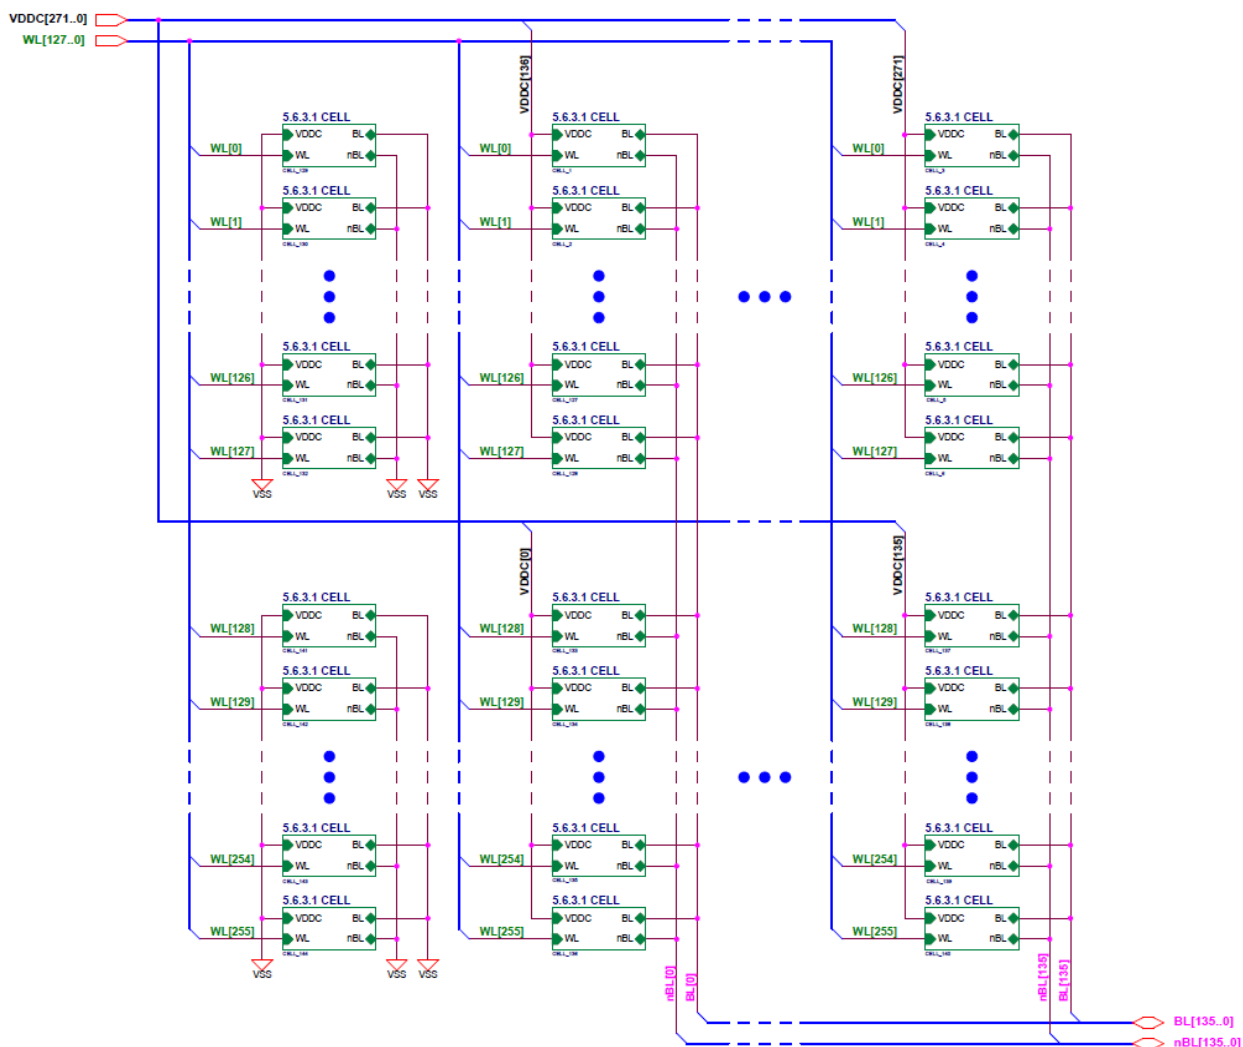
83. For example, reverse engineering shows that Broadwell processors include several SRAM arrays, including the array shown below as SRAM4:



“a memory array comprising a first line of memory cells and a second line of memory cells;”

84. Intel Broadwell processors operate using a method comprising providing a memory array comprising a first line of memory cells and a second line of memory cells.

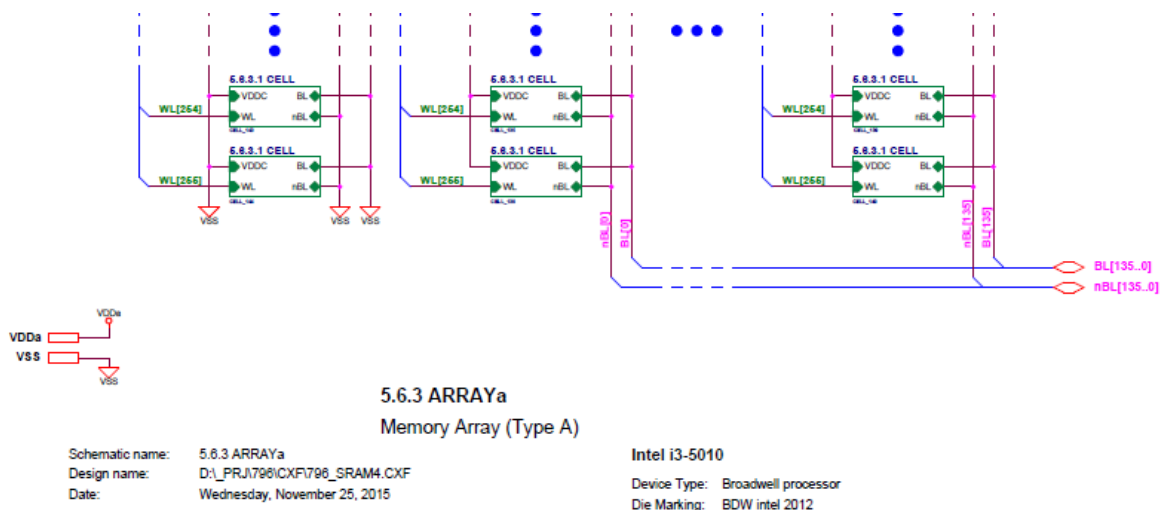
85. For example, reverse engineering shows that Broadwell’s SRAM arrays are constructed using a plurality of columns:



[“a first power supply terminal;”]

86. Intel Broadwell processors operate using a method comprising providing a first power supply terminal.

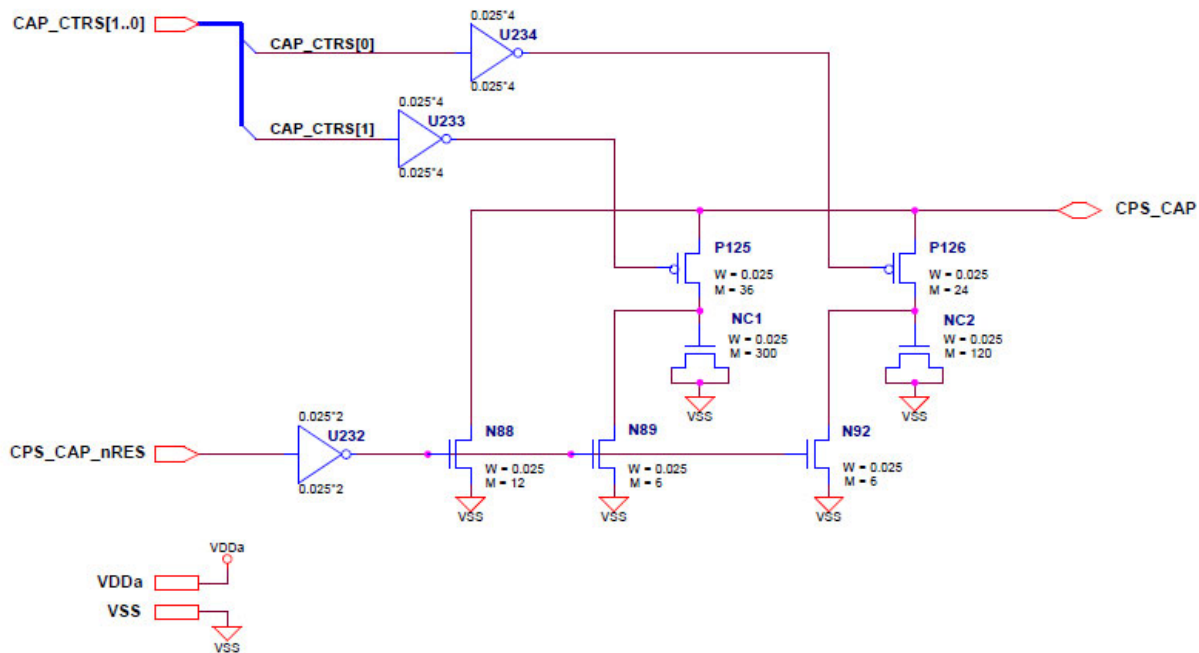
87. For example, reverse engineering shows that Broadwell includes a first power supply terminal (e.g., VDDa):



["a first capacitance structure includes a plurality of dummy cells;"]

88. Intel Broadwell processors operate using a method comprising providing a first capacitance structure that includes a plurality of dummy cells.

89. For example, reverse engineering shows that Intel Broadwell processors include a cell power switch capacitor comprising a plurality of dummy transistors (e.g., NC1 and NC2):



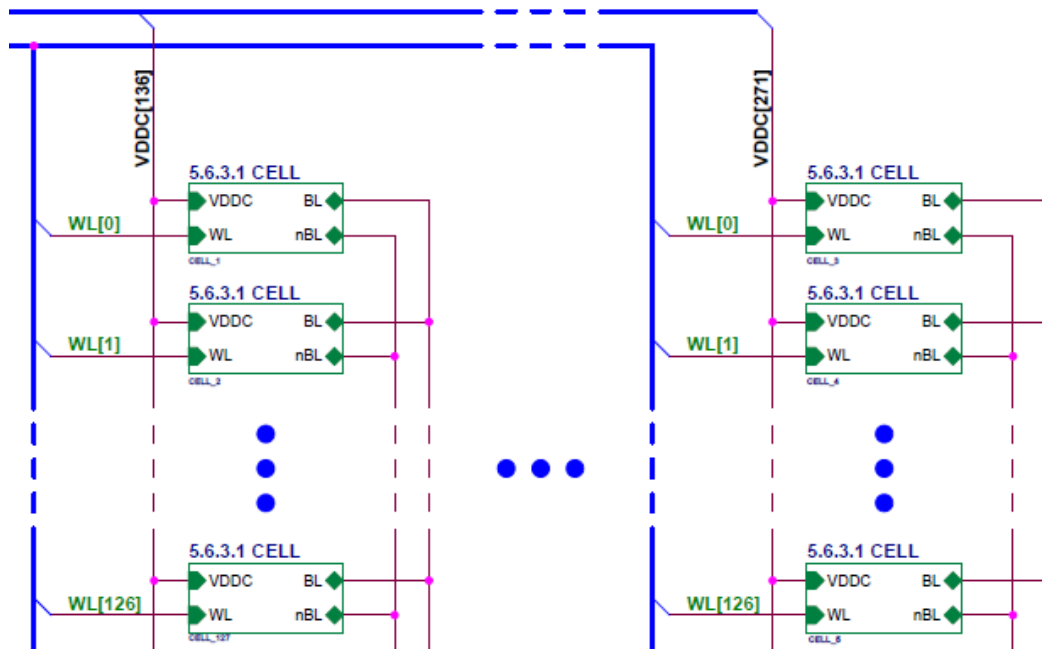
5.2.2 CPS_CAP

Cell Power Switch Capacitor

[“a first power supply line coupled to the first line of memory cells; and”]

90. Intel Broadwell processors operate using a method comprising providing a first power supply line coupled to the first line of memory cells.

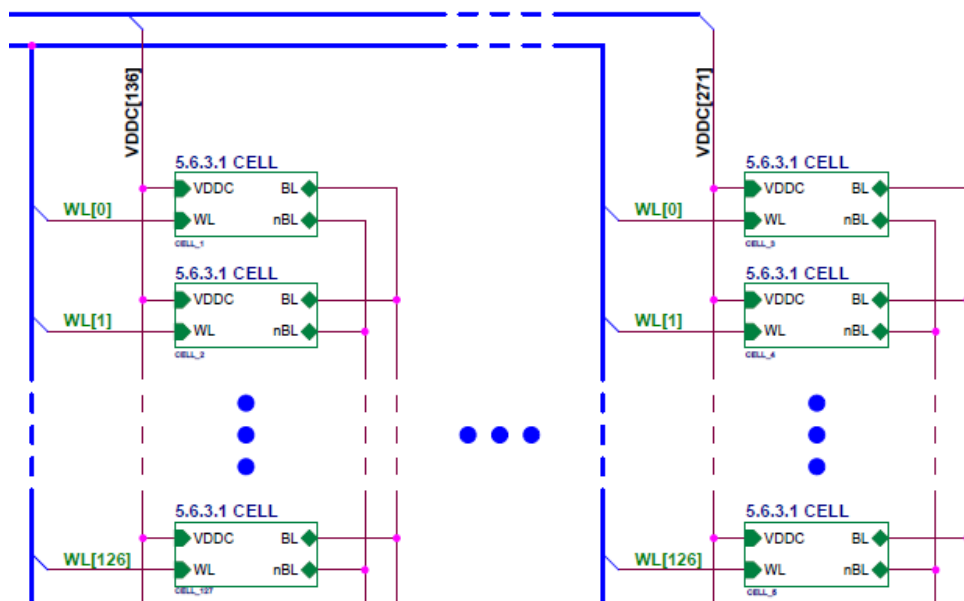
91. For example, reverse engineering shows that Broadwell includes a first power supply line (e.g., VDDC[136]), coupled to the first line of memory cells:



[“a second power supply line coupled to the second line of memory cells;”]

92. Intel Broadwell processors operates using a method comprising providing a second power supply line coupled to the second line of memory cells.

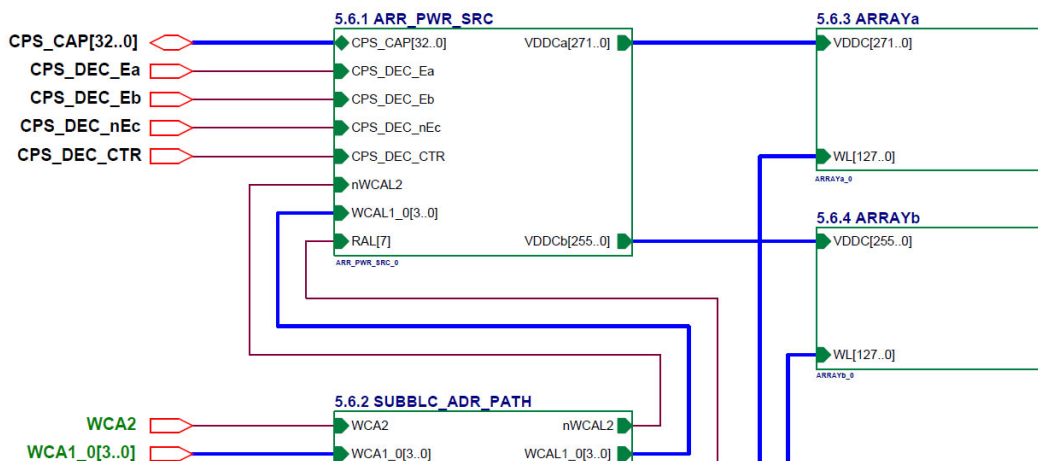
93. For example, reverse engineering shows that Broadwell includes a second power supply line (e.g., VDDC[271]) coupled to the second line of memory cells:



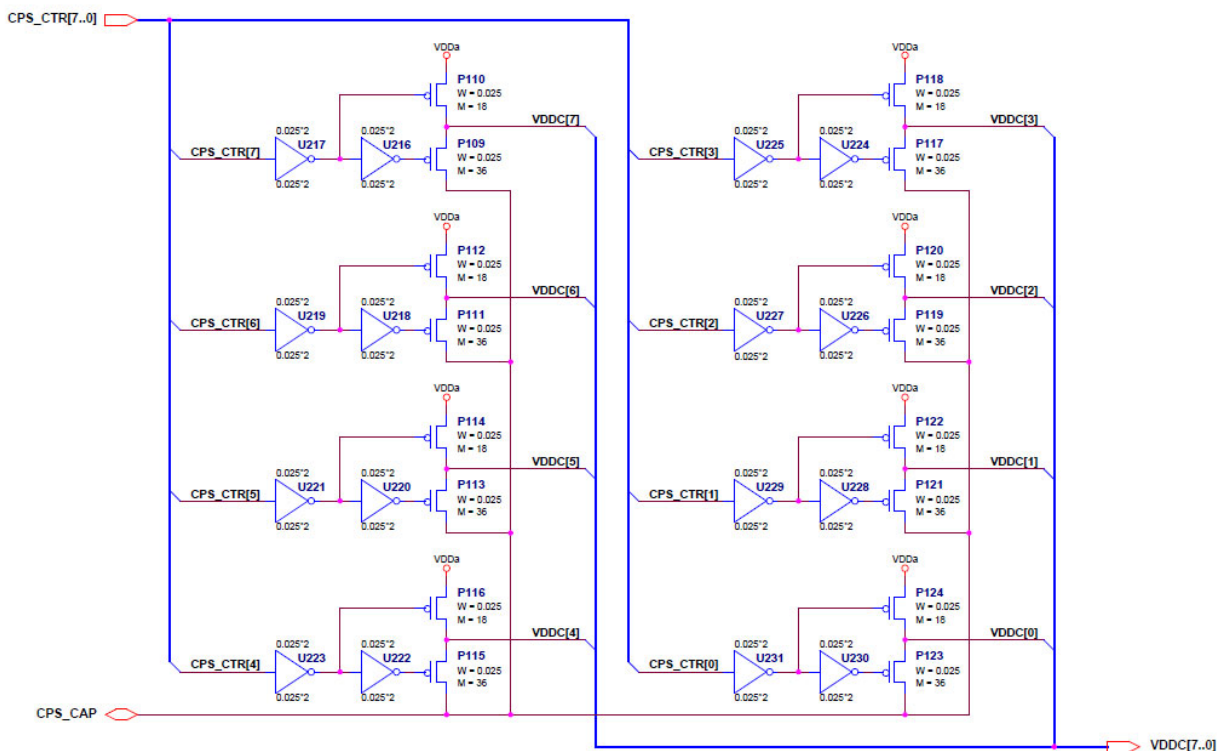
[“a switching circuit that has transistors that connected between the first power supply terminal, the first power supply line, the second power supply line and the first capacitance structure”]

94. Intel Broadwell processors operate using a method comprising providing a switching circuit that has transistors that connected between the first power supply terminal, the first power supply line, the second power supply line and the first capacitance structure.

95. For example, reverse engineering shows that Broadwell includes a switching circuit, *e.g.*, ARR_PWR_SRC, that has transistors that connect a capacitance structure, the power supply lines, and the plurality of power supply lines:



96. Reverse engineering also shows connection of first power supply terminal (VDDa) to VDDc:



["selecting the second line of memory cells for writing;"]

97. Intel Broadwell processors operate using a method comprising selecting the second line of memory cells for writing.

98. For example, the column based voltage bias technique used in the SRAM selects particular columns for reading. *See, e.g.*, “A 0.6V 1.5GHz 84Mb SRAM Design in 14nm FinFET CMOS Technology” [hereinafter “SRAM”] at 1: “The CS-TVC operation begins with a self-timed pulse (TVC PULSE) aligned to the rising edge of WL that simultaneously disconnects the selected VCS region from VCC, disables the NMOS devices in the CS-TVC capacitor, and connects VCS to the pre-discharged GCSCAP node.”

[“coupling the first power supply terminal to the first power supply line;”]

99. Intel Broadwell processors operate using a method comprising coupling the first power supply terminal to the first power supply line.

100. For example, columns that are not being written to (including the first column) are connected to the first power supply terminal to their power supply lines. This power supply terminal is described as “VCS” by Intel. *See, e.g.*, SRAM at 1: “The CS-TVC operation begins with a self-timed pulse (TVC PULSE) aligned to the rising edge of WL that simultaneously disconnects the selected VCS region from VCC, disables the NMOS devices in the CS-TVC capacitor, and connects VCS to the pre-discharged GCSCAP node. Charge is balanced through a PMOS switch between the VCS region selected and GCSCAP, resulting in a temporary suppression of the VCS node to improve write margin. The falling edge of the TVC pulse completes the operation and restores the VCS voltage level to VCC.”

[“decoupling the second line of memory cells from the first power supply terminal;”]

101. Intel Broadwell processors operate using a method comprising decoupling the second line of memory cells from the first power supply terminal.

102. For example, columns that are being written to (including the second column) are disconnected from the first power supply terminal. *See, e.g.*, SRAM at 1: “The CS-TVC operation begins with a self-timed pulse (TVC PULSE) aligned to the rising edge of WL that simultaneously disconnects the selected VCS region from VCC, disables the NMOS devices in the CS-TVC capacitor, and connects VCS to the pre-discharged GCSCAP node. Charge is balanced through a PMOS switch between the VCS region selected and GCSCAP, resulting in a temporary suppression of the VCS node to improve write margin. The falling edge of the TVC pulse completes the operation and restores the VCS voltage level to VCC.”

[“coupling charge from the second power supply line to the first capacitance structure; and”]

103. Intel Broadwell processors operate using a method comprising coupling charge from the second power supply line to the first capacitance structure.

104. For example, columns that are being written to (including the second column) have charge coupled to the capacitance structure, here referred to as GCSCAP. *See, e.g.*, SRAM at 1: “The CS-TVC operation begins with a self-timed pulse (TVC PULSE) aligned to the rising edge of WL that simultaneously disconnects the selected VCS region from VCC, disables the NMOS devices in the CS-TVC capacitor, and connects VCS to the pre-discharged GCSCAP node. Charge is balanced through a PMOS switch between the VCS region selected and GCSCAP, resulting in a temporary suppression of the VCS node to improve write margin. The falling edge of the TVC pulse completes the operation and restores the VCS voltage level to VCC.”

[“writing a memory cell in the second line of memory cells.”]

105. Intel Broadwell processors operate using a method comprising writing a memory cell in the second line of memory cells.

106. For example, memory cells are written in columns that are being written to (including the second column). *See, e.g.*, SRAM at 1: “The CS-TVC operation begins with a self-timed pulse (TVC PULSE) aligned to the rising edge of WL that simultaneously disconnects the selected VCS region from VCC, disables the NMOS devices in the CS-TVC capacitor, and connects VCS to the pre-discharged GCSCAP node. Charge is balanced through a PMOS switch between the VCS region selected and GCSCAP, resulting in a temporary suppression of the VCS node to improve write margin. The falling edge of the TVC pulse completes the operation and restores the VCS voltage level to VCC.”

107. Intel has had knowledge of the ‘485 Patent at least since the filing of this Complaint, and if it did not have actual knowledge prior to that time, it was willfully blind to the existence of the ‘485 Patent based on, for example, its publicly-known corporate policy forbidding its employees from reading patents held by outside companies or individuals, as already described above. Intel should have known that its conduct was infringing both prior to and following the filing of this case.

108. VLSI is informed and believes, and thereon alleges, that Intel actively, knowingly, and intentionally has induced infringement of the ‘485 Patent by, for example, controlling the design and manufacture of, offering for sale, selling, supplying, and otherwise providing instruction and guidance regarding the above-described products with the knowledge and specific intent to encourage and facilitate infringing uses of such products by its customers both inside and outside the United States. For example, Intel publicly provides documentation, including datasheets available through Intel’s publicly accessible ARK service and software developer’s manuals, instructing customers on uses of Intel’s products that infringe the methods of the ‘485 Patent. *See, e.g.*, <http://ark.intel.com>. On information and belief, Intel’s customers directly

infringe the '485 Patent by, for example, making, using, offering to sell, and selling within the United States, and importing into the United States, without authority or license, products containing the above-described Intel products.

109. VLSI is informed and believes, and thereon alleges, that Intel has contributed to the infringement by its customers of the '485 Patent by, without authority, importing, selling and offering to sell within the United States materials and apparatuses for practicing the claimed invention of the '485 Patent both inside and outside the United States. For example, the above-described products constitute a material part of the inventions of the '485 Patent and are not staple articles or commodities of commerce suitable for substantial noninfringing use. On information and belief, Intel knows that the above-described products constitute a material part of the inventions of the '485 Patent and are not staple articles or commodities of commerce suitable for substantial noninfringing use. On information and belief, Intel's customers directly infringe the '485 Patent by, for example, making, using, offering to sell, and selling within the United States, and importing into the United States, without authority or license, products containing the above-described Intel products.

110. As a result of Intel's infringement of the '485 Patent, VLSI has been damaged. VLSI is entitled to recover for damages sustained as a result of Intel's wrongful acts in an amount subject to proof at trial.

111. To the extent 35 U.S.C. § 287 is determined to be applicable, on information and belief its requirements have been satisfied with respect to the '485 Patent.

112. In addition, Intel's infringing acts and practices have caused and are causing immediate and irreparable harm to VLSI.

113. VLSI is informed and believes, and thereon alleges, that Intel's infringement of the '485 Patent has been and continues to be willful. As noted above, Intel has had knowledge of the '485 Patent, or, at minimum, has been willfully blind to the existence of the '485 Patent. Moreover, after service of this Complaint, Intel undisputedly had knowledge of its infringement of the '485 Patent. Intel has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for VLSI's patent rights. Thus, Intel's infringing actions have been and continue to be consciously wrongful.

FOURTH CLAIM

(Infringement of U.S. Patent No. 7,606,983)

114. VLSI re-alleges and incorporates herein by reference Paragraphs 1-113 of its Complaint.

115. The '983 Patent, entitled "Sequential ordering of transactions in digital systems with multiple requestors," was duly and lawfully issued October 20, 2009. A true and correct copy of the '983 Patent is attached hereto as Exhibit 4.

116. The '983 Patent names Kevin Locker as the inventor.

117. The '983 Patent has been in full force and effect since its issuance. VLSI owns by assignment the entire right, title, and interest in and to the '983 Patent, including the right to seek damages for past, current, and future infringement thereof.

118. The '983 Patent "relates generally to methods and apparatuses for designing digital systems." Ex. 4 at 1:7-8.

119. The '983 Patent explains that "[m]any digital electronic devices include multiple autonomous or semi-autonomous functional modules, such as processors, that share access to common resources, such as memory." *Id.* at 1:19-22. The '983 Patent describes, for example,

providing “an improved transaction ordering policy in which individual occurrences of access requests can specify whether or not the associated transaction is to be performed in order.” *Id.* at 4:57-60.

120. VLSI is informed and believes, and thereon alleges, that Intel has infringed and unless enjoined will continue to infringe one or more claims of the ‘983 Patent, in violation of 35 U.S.C. § 271, by, among other things, making, using, offering to sell, and selling within the United States, supplying or causing to be supplied in or from the United States, and importing into the United States, without authority or license, Intel products that implement the Intel Quick Path Interconnect (“QPI”) Link Layer in an infringing manner.

121. For example, the ‘983 accused products embody every limitation of at least claim 11 of the ‘983 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

[“11. A method of processing information, the method comprising:”]

122. Broadwell Server processors, which support the Intel QPI Link Layer, use a method of processing information as shown below. *See, e.g.*, <https://ark.intel.com/products/91317/Intel-Xeon-Processor-E5-2699-v4-55M-Cache-2-20-GHz->.


Essentials		Export specifications
Product Collection	Intel® Xeon® Processor E5 v4 Family	
Code Name	Products formerly Broadwell	
Vertical Segment	Server	
Processor Number	E5-2699V4	
Status	Launched	
Launch Date ?	Q1'16	
Lithography ?	14 nm	
Recommended Customer Price ?	\$4115.00	
Performance		
# of Cores ?	22	
# of Threads ?	44	
Processor Base Frequency ?	2.20 GHz	
Max Turbo Frequency ?	3.60 GHz	
Cache ?	55 MB SmartCache	
Bus Speed ?	9.6 GT/s QPI	
# of QPI Links ?	2	
TDP ?	145 W	
VID Voltage Range ?	0	

["holding data;"]

123. Broadwell Server processors use a method of processing information comprising holding data.

124. For example, Broadwell Server processors are each connected via a memory interface to a memory controller. *See, e.g.,* An Introduction to the Intel QuickPath Interconnect [hereinafter “QPI”] at 8.

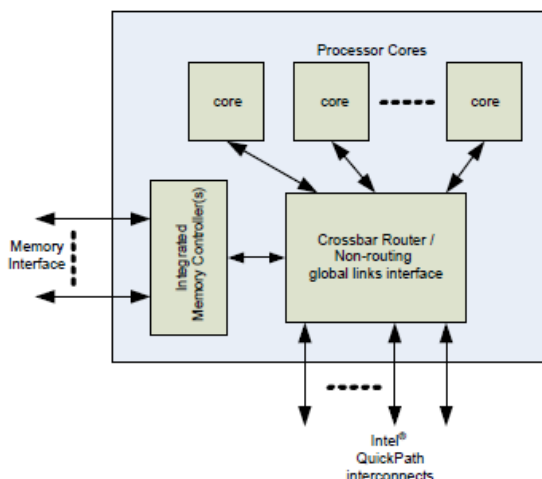
Figure 6. Intel® QuickPath Interconnect



125. As another example, Broadwell Server processors include shared and separate caches. *See, e.g.*, QPI at 8.

Figure 7 shows a schematic of a processor with external Intel® QuickPath Interconnects. The processor may have one or more cores. When multiple cores are present, they may share caches or have separate caches. The processor also typically has one or more integrated memory controllers. Based on the level of scalability supported in the processor, it may include an integrated crossbar router and more than one Intel® QuickPath Interconnect port (a port contains a pair of uni-directional links).

Figure 7. Block Diagram of Processor with Intel® QuickPath Interconnects



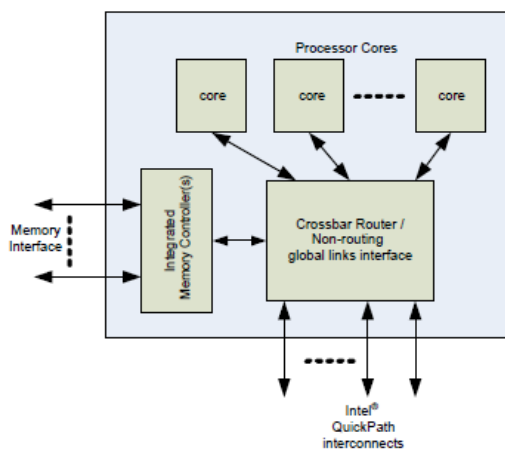
[“performing in parallel at least two series of operations;”]

126. Broadwell Server processors use a method of processing information comprising performing in parallel at least two series of operations.

127. For example, Broadwell Server systems contain multiple cores that perform operations in parallel. *See, e.g.,* QPI at 8.

Figure 7 shows a schematic of a processor with external Intel® QuickPath Interconnects. The processor may have one or more cores. When multiple cores are present, they may share caches or have separate caches. The processor also typically has one or more integrated memory controllers. Based on the level of scalability supported in the processor, it may include an integrated crossbar router and more than one Intel® QuickPath Interconnect port (a port contains a pair of uni-directional links).

Figure 7. Block Diagram of Processor with Intel® QuickPath Interconnects



[“generating an access request when performing one of the operations when the operation involves accessing the data;”]

128. Broadwell Server processors use a method of processing information comprising generating an access request when performing one of the operations when the operation involves accessing the data.

129. For example, when one core in a Broadwell Server system makes an access to data, such as a read from memory, it may create one or more access requests such as, for example, a message in a Home Message Class or a message in a Non-coherent Standard Message Class. *See, e.g., “Weaving High Performance Multiprocessor Fabric” [hereinafter “Fabric”] at 42, 119-120.*

Request Generation

Whenever an Intel QPI agent is required to read from, or write data to memory or I/O, the agent creates an internal request to the appropriate address space. A request bound for coherent shared memory first checks the local cache for a valid copy of the data. If that check fails to find the data in the agent's local cache, the protocol layer uses the rules of the Intel QPI cache coherence protocol to obtain the most up-to-date copy of data from the system memory, or from another caching agent in the system. The protocol layer then issues a set of requests and associated responses to complete the operations.

Whenever an agent generates a request, either from the protocol layer or directly from the processor core for non-coherent operations, the request needs to undergo a mapping operation in order to determine its destination on the Intel QPI fabric. Agents generating a request specify the destination in terms of the system address space, but destinations on Intel QPI are determined by a Node ID field embedded in the packets being sent. A piece of logic that exists in every requesting agent is known as the *source address decoder* (SAD). The SAD logic takes into account the nature of messages that are to be issued, uses a lookup table to map from system address to Node ID, and then provides that information to be included in the messages sent out onto the Intel QPI fabric.

Home Message Class—HOM

HOM class traffic consists of requests for access to the coherent data space, or for snoop responses related to those requests and to the SNP messages. HOM traffic is always directed to the home agent responsible for that particular address in the coherent memory space. Requests are always generated by a caching agent. Snoop responses are also always generated by the caching agent. HOM messages are also the minimum packet size of 1 or 2 flits.

In current products, the ordering requirement placed upon HOM messages is as follows. For any given caching agent (source) and home agent (destination) combination, messages are kept in order on a per-address basis. Messages from different sources or being sent to different destinations have no ordering requirements respective to each other. Likewise messages to different addresses have no ordering requirements relative to each other. Current home agent designs rely upon the link layer maintaining this HOM channel ordering to ensure proper operation of the coherent protocol.

Non-Coherent Standard—NCS

NCS messages are requests dealing with access to system resources that are generally outside of the coherent memory space. This includes things such as reads to non-coherent memory regions, read and write configuration and I/O space transactions, and interrupt acknowledgements. Certain NCS transactions can also access the coherent memory space, but it is not required that coherency be maintained by the Intel QPI hardware in this case. Such read requests would be performed without the addition of snoop requests and responses. Small packets are used, ranging from the minimum packet size, up to three flit packet sizes for some message types. NCS packets do not contain any data flits. However, some messages in NCS support up to an 8 byte payload embedded into the packet's header flits.

["including in each access request an indication of whether or not this occurrence of the access request has a specified order among other occurrences of the access request; including an indication of the specified order in those occurrence[s] of the access request that ar[e] ordered; receiving the access requests from each of the series of operations;"]

130. Broadwell Server processors use a method of processing information comprising including in each access request an indication of whether or not this occurrence of the access request has a specified order among other occurrences of the access request, including an indication of the specified order in those occurrences of the access request that are ordered; receiving the access requests from each of the series of operations.

131. For example, each request sent over the QuickPath Interconnect in connection with Broadwell Server processors contains indications including, *e.g.*, Requester Transaction ID, address, and type that indicates that certain types of requests must be ordered, and if they are to be

ordered, which order they are to be performed in. These ordering requirements are based, *e.g.*, on “stall” rules that require that certain messages be sent before others. *See, e.g.*, Datasheet at 81, 84.

Outgoing Request Buffer (ORB)

When a transaction is issued onto the Intel QuickPath Interconnect, an ORB entry is allocated. This list keeps all pertinent information about the transaction header needed to complete the request. It also stores the cacheline address for coherent transactions to allow conflict checking with snoops and other local requests. See [Section 4.10](#) for details. An Intel QuickPath Interconnect response may come as multiple phases. The ORB tracks each completion phase (that is, Data* and Cmp) and any conflicts (that is, snoops, FrcAckCnflt Response) in the ORB.

When a request is issued, a RTID (Requestor Transaction ID) is assigned based on Home NodeID. Limitations are placed on the RTID based on how many transactions each home agent can support. The RTID allocation limitation binds the tag to a specific range; this range is referred to as *MaxRequests*, where the range is 0 to *MaxRequests-1*. This value is programmable in the IOH configuration registers.

The Home NodeID and RTID are returned in the responses. The ORB must provide a way to associate the response's RTID and Home NodeID with a ORB entry. This can be done through a reverse lookup table based on RTID and Home NodeID or by searching the table for matching RTID and NodeID. Simplification of the reverse lookup function is possible by limiting the scope of the RTID allocation as described in [Section 4.9.2](#).

ORB Depth

The ORB depth is based on the number of requests the IOH needs to be pending on Intel QuickPath Interconnect to achieve full bandwidth, given worst-case average latency for a 4S system. This latency is calculated from allocation and de-allocation of tag. The ORB depth is 256 entries. The IOH will further partition the 256 entries into 128 entries per port. For evenly distributed traffic, this will provide the equivalent of 256 entries.

Table 4-14. Local-Local Conflict Actions

Local Request	ORB or Write Cache State	Action
Rd or NonSnpRd or NonSnpWr*	Rd or NonSnpRd or NonSnpWr*	<ul style="list-style-type: none"> Stall until completed.
	RFO	<ul style="list-style-type: none"> Stall until EWB completed. Force EWB on Promotion.
	EWB	<ul style="list-style-type: none"> Stall until EWB completed.
	E- or MG-state	<ul style="list-style-type: none"> Stall until EWB completed. Force EWB on Promotion.
	M-state	<ul style="list-style-type: none"> Stall until EWB completed. Force EWB.
	Non-Coh VT-d table data	<ul style="list-style-type: none"> Data may be snarfed locally for other Intel VT-d Reads.
RFO	Rd or NonSnpRd or NonSnpWr*	<ul style="list-style-type: none"> Stall until completed.
	RFO	<ul style="list-style-type: none"> Stall until Promotion. If promotion does not evict the line then Complete RFO. If promotion causes EWB, then send RFO after EWB completed.
	EWB	<ul style="list-style-type: none"> Stall until EWB completed.
	E- or MG-state	<ul style="list-style-type: none"> Stall until Promotion. If promotion does not evict the line then Complete RFO. If promotion causes EWB, then send RFO after EWB completed.
	M-state	<ul style="list-style-type: none"> Complete immediately, no stall condition.
M-Promote	Rd or NonSnpRd or NonSnpWr* or RFO or EQB or M-state	<ul style="list-style-type: none"> Impossible. EWB can only be received in E or MG state.
	E- or MG-state	<ul style="list-style-type: none"> Normal flow to M-state and Eviction. If M-state does not cause eviction under normal rules then Conflict queue is checked to see if EWB required or RFO completion required. EWB: on EWB completion, next conflict is cleared. If multiple ownerships are granted simultaneously, then state will change to MG state.

132. See also, e.g., Fabric at 118-120.

Table 3.1 Message Classes

Message Class	Abbreviation	Ordering Requirements	Data Content?
Snoop	SNP	None	No
Home	HOM	Limited	No
Data Response	DRS	None	Yes
Non-Data Response	NDR	None	No
Non-coherent Standard	NCS	None	No ¹
Non-coherent Bypass	NCB	None	Yes
Special (Link Management)	SPC	None	No

¹ See section on NCS for more information

Home Message Class—HOM

HOM class traffic consists of requests for access to the coherent data space, or for snoop responses related to those requests and to the SNP messages. HOM traffic is always directed to the home agent responsible for that particular address in the coherent memory space. Requests are always generated by a caching agent. Snoop responses are also always generated by the caching agent. HOM messages are also the minimum packet size of 1 or 2 flits.

In current products, the ordering requirement placed upon HOM messages is as follows. For any given caching agent (source) and home agent (destination) combination, messages are kept in order on a per-address basis. Messages from different sources or being sent to different destinations have no ordering requirements respective to each other. Likewise messages to different addresses have no ordering requirements relative to each other. Current home agent designs rely upon the link layer maintaining this HOM channel ordering to ensure proper operation of the coherent protocol.

Non-Coherent Standard—NCS

NCS messages are requests dealing with access to system resources that are generally outside of the coherent memory space. This includes things such as reads to non-coherent memory regions, read and write configuration and I/O space transactions, and interrupt acknowledgements. Certain NCS transactions can also access the coherent memory space, but it is not required that coherency be maintained by the Intel QPI hardware in this case. Such read requests would be performed without the addition of snoop requests and responses. Small packets are used, ranging from the minimum packet size, up to three flit packet sizes for some message types. NCS packets do not contain any data flits. However, some messages in NCS support up to an 8 byte payload embedded into the packet's header flits.

["determining a performance order for the access requests, wherein if the access requests are ordered then the performance order conforms to the specified order; and performing the access requests in the performance order."]

133. Broadwell Server processors use a method of processing information comprising determining a performance order for the access requests, wherein if the access requests are ordered then the performance order conforms to the specified order and performing the access requests in the performance order.

134. For example, QPI agents in Broadwell Server processors perform the access requests in the order as determined by the rules discussed above. *See, e.g.*, Datasheet at 81, 84.

Outgoing Request Buffer (ORB)

When a transaction is issued onto the Intel QuickPath Interconnect, an ORB entry is allocated. This list keeps all pertinent information about the transaction header needed to complete the request. It also stores the cacheline address for coherent transactions to allow conflict checking with snoops and other local requests. See [Section 4.10](#) for details. An Intel QuickPath Interconnect response may come as multiple phases. The ORB tracks each completion phase (that is, Data* and Cmp) and any conflicts (that is, snoops, FrcAckCnflt Response) in the ORB.

When a request is issued, a RTID (Requestor Transaction ID) is assigned based on Home NodeID. Limitations are placed on the RTID based on how many transactions each home agent can support. The RTID allocation limitation binds the tag to a specific range; this range is referred to as *MaxRequests*, where the range is 0 to *MaxRequests-1*. This value is programmable in the IOH configuration registers.

The Home NodeID and RTID are returned in the responses. The ORB must provide a way to associate the response's RTID and Home NodeID with a ORB entry. This can be done through a reverse lookup table based on RTID and Home NodeID or by searching the table for matching RTID and NodeID. Simplification of the reverse lookup function is possible by limiting the scope of the RTID allocation as described in [Section 4.9.2](#).

ORB Depth

The ORB depth is based on the number of requests the IOH needs to be pending on Intel QuickPath Interconnect to achieve full bandwidth, given worst-case average latency for a 4S system. This latency is calculated from allocation and de-allocation of tag. The ORB depth is 256 entries. The IOH will further partition the 256 entries into 128 entries per port. For evenly distributed traffic, this will provide the equivalent of 256 entries.

Table 4-14. Local-Local Conflict Actions

Local Request	ORB or Write Cache State	Action
Rd or NonSnpRd or NonSnpWr*	Rd or NonSnpRd or NonSnpWr*	<ul style="list-style-type: none"> Stall until completed.
	RFO	<ul style="list-style-type: none"> Stall until EWB completed. Force EWB on Promotion.
	EWB	<ul style="list-style-type: none"> Stall until EWB completed.
	E- or MG-state	<ul style="list-style-type: none"> Stall until EWB completed. Force EWB on Promotion.
	M-state	<ul style="list-style-type: none"> Stall until EWB completed. Force EWB.
	Non-Coh VT-d table data	<ul style="list-style-type: none"> Data may be snarfed locally for other Intel VT-d Reads.
RFO	Rd or NonSnpRd or NonSnpWr*	<ul style="list-style-type: none"> Stall until completed.
	RFO	<ul style="list-style-type: none"> Stall until Promotion. If promotion does not evict the line then Complete RFO. If promotion causes EWB, then send RFO after EWB completed.
	EWB	<ul style="list-style-type: none"> Stall until EWB completed.
	E- or MG-state	<ul style="list-style-type: none"> Stall until Promotion. If promotion does not evict the line then Complete RFO. If promotion causes EWB, then send RFO after EWB completed.
	M-state	<ul style="list-style-type: none"> Complete immediately, no stall condition.
M-Promote	Rd or NonSnpRd or NonSnpWr* or RFO or EQB or M-state	<ul style="list-style-type: none"> Impossible. EWB can only be received in E or MG state.
	E- or MG-state	<ul style="list-style-type: none"> Normal flow to M-state and Eviction. If M-state does not cause eviction under normal rules then Conflict queue is checked to see if EWB required or RFO completion required. EWB: on EWB completion, next conflict is cleared. If multiple ownerships are granted simultaneously, then state will change to MG state.

135. See also, e.g., QPI at 109.

Table 6-1. Ordering Term Definitions

Term	Definition
Intel QuickPath Interconnect Ordering Domain	The Intel QuickPath Interconnect has a relaxed ordering model allowing reads, writes and completions to flow independent of each other. Intel QuickPath Interconnect implements this through the use of multiple, independent virtual channels. In general, the Intel QuickPath Interconnect ordering domain is considered unordered.
PCI Express Ordering Domain	PCI Express (and all other prior PCI generations) have specific ordering rules to enable low cost components to support the Producer-Consumer model. For example, no transaction can pass a write flowing in the same direction. In addition, PCI implements ordering relaxations to avoid deadlocks (for example, completions must pass non-posted requests). The set of these rules are described in <i>PCI Express Base Specification, Revision 1.0a</i> .

136. See also, e.g., Fabric at 118-120.

Table 3.1 Message Classes

Message Class	Abbreviation	Ordering Requirements	Data Content?
Snoop	SNP	None	No
Home	HOM	Limited	No
Data Response	DRS	None	Yes
Non-Data Response	NDR	None	No
Non-coherent Standard	NCS	None	No ¹
Non-coherent Bypass	NCB	None	Yes
Special (Link Management)	SPC	None	No

1 See section on NCS for more information

Home Message Class—HOM

HOM class traffic consists of requests for access to the coherent data space, or for snoop responses related to those requests and to the SNP messages. HOM traffic is always directed to the home agent responsible for that particular address in the coherent memory space. Requests are always generated by a caching agent. Snoop responses are also always generated by the caching agent. HOM messages are also the minimum packet size of 1 or 2 flits.

In current products, the ordering requirement placed upon HOM messages is as follows. For any given caching agent (source) and home agent (destination) combination, messages are kept in order on a per-address basis. Messages from different sources or being sent to different destinations have no ordering requirements respective to each other. Likewise messages to different addresses have no ordering requirements relative to each other. Current home agent designs rely upon the link layer maintaining this HOM channel ordering to ensure proper operation of the coherent protocol.

Non-Coherent Standard—NCS

NCS messages are requests dealing with access to system resources that are generally outside of the coherent memory space. This includes things such as reads to non-coherent memory regions, read and write configuration and I/O space transactions, and interrupt acknowledgements. Certain NCS transactions can also access the coherent memory space, but it is not required that coherency be maintained by the Intel QPI hardware in this case. Such read requests would be performed without the addition of snoop requests and responses. Small packets are used, ranging from the minimum packet size, up to three flit packet sizes for some message types. NCS packets do not contain any data flits. However, some messages in NCS support up to an 8 byte payload embedded into the packet's header flits.

137. Intel has had knowledge of the '983 Patent at least since the filing of this Complaint, and if it did not have actual knowledge prior to that time, it was willfully blind to the existence of the '983 Patent based on, for example, its publicly-known corporate policy forbidding its employees from reading patents held by outside companies or individuals, as already described above. Intel should have known that its conduct was infringing both prior to and following the filing of this case.

138. VLSI is informed and believes, and thereon alleges, that Intel actively, knowingly, and intentionally has induced infringement of the '983 Patent by, for example, controlling the design and manufacture of, offering for sale, selling, supplying, and otherwise providing instruction and guidance regarding the above-described products with the knowledge and specific intent to encourage and facilitate infringing uses of such products by its customers both inside and outside the United States. For example, Intel publicly provides documentation, including datasheets available through Intel's publicly accessible ARK service and software developer's manuals, instructing customers on uses of Intel's products that infringe the methods of the '983

Patent. *See, e.g.,* <http://ark.intel.com>. On information and belief, Intel's customers directly infringe the '983 Patent by, for example, making, using, offering to sell, and selling within the United States, and importing into the United States, without authority or license, products containing the above-described Intel products.

139. VLSI is informed and believes, and thereon alleges, that Intel has contributed to the infringement by its customers of the '983 Patent by, without authority, importing, selling and offering to sell within the United States materials and apparatuses for practicing the claimed invention of the '983 Patent both inside and outside the United States. For example, the above-described products constitute a material part of the inventions of the '983 Patent and are not staple articles or commodities of commerce suitable for substantial noninfringing use. On information and belief, Intel knows that the above-described products constitute a material part of the inventions of the '983 Patent and are not staple articles or commodities of commerce suitable for substantial noninfringing use. On information and belief, Intel's customers directly infringe the '983 Patent by, for example, making, using, offering to sell, and selling within the United States, and importing into the United States, without authority or license, products containing the above-described Intel products.

140. As a result of Intel's infringement of the '983 Patent, VLSI has been damaged. VLSI is entitled to recover for damages sustained as a result of Intel's wrongful acts in an amount subject to proof at trial.

141. To the extent 35 U.S.C. § 287 is determined to be applicable, on information and belief its requirements have been satisfied with respect to the '983 Patent.

142. In addition, Intel's infringing acts and practices have caused and are causing immediate and irreparable harm to VLSI.

143. VLSI is informed and believes, and thereon alleges, that Intel's infringement of the '983 Patent has been and continues to be willful. As noted above, Intel has had knowledge of the '983 Patent, or, at minimum, has been willfully blind to the existence of the '983 Patent. Moreover, after service of this Complaint, Intel undisputedly had knowledge of its infringement of the '983 Patent. Intel has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for VLSI's patent rights. Thus, Intel's infringing actions have been and continue to be consciously wrongful.

FIFTH CLAIM

(Infringement of U.S. Patent No. 7,725,759)

144. VLSI re-alleges and incorporates herein by reference Paragraphs 1-143 of its Complaint.

145. The '759 Patent, entitled "System and method of managing clock speed in an electronic device," was duly and lawfully issued May 25, 2010. A true and correct copy of the '759 Patent is attached hereto as Exhibit 5.

146. The '759 Patent names Matthew Henson as inventor.

147. The '759 Patent has been in full force and effect since its issuance. VLSI owns by assignment the entire right, title, and interest in and to the '759 Patent, including the right to seek damages for past, current, and future infringement thereof.

148. The '759 Patent states that it "relates to electronic devices and to managing clock speeds within electronic devices." Ex. 5 at 1:6-7.

149. The '759 Patent explains, for instance, "there is a need for an improved system and method of controlling a clock frequency in an electronic device in order to selectively deliver faster clock speeds." *Id.* at 1:22-24.

150. The ‘759 Patent discloses, among other things, “[a] method of controlling a clock frequency . . . [that] includes monitoring a plurality of master devices that are coupled to a bus within a system.” *Id.* at 1:46-48.

151. VLSI is informed and believes, and thereon alleges, that Intel has infringed and unless enjoined will continue to infringe one or more claims of the ‘759 Patent, in violation of 35 U.S.C. § 271, by, among other things, making, using, offering to sell, and selling within the United States, supplying or causing to be supplied in or from the United States, and importing into the United States, without authority or license, Intel products that use infringing Hardware-Controlled Performance States (“HWP” or “Speed Shift”) technology.

152. For example, the ‘759 accused products embody every limitation of at least claim 1 of the ‘759 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

["1. A method comprising:"]

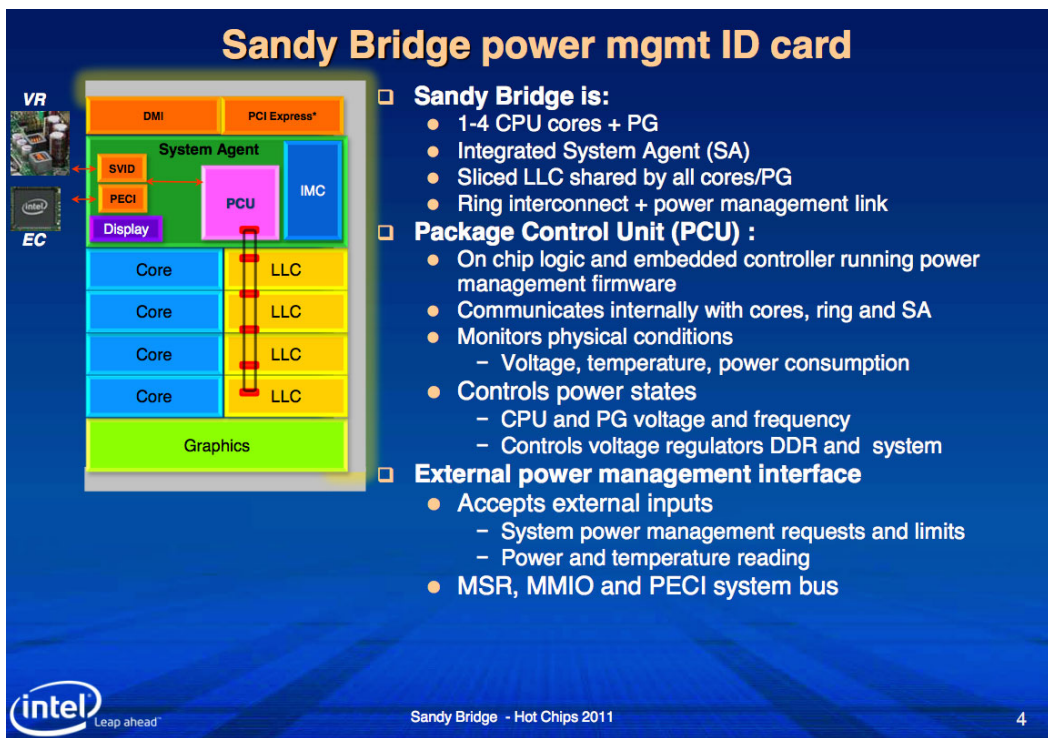
153. Intel Skylake processors, which include Hardware-Controlled Performance States (HWP or “Speed Shift”), are operated using a method comprising the elements listed below. *See, e.g.,* <https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/6th-gen-core-family-mobile-brief.pdf> at 8.

6TH GEN INTEL® CORE PROCESSOR™ FEATURES AT A GLANCE	
FEATURES ¹	BENEFITS
Intel® Active Management Technology (Intel® AMT) ¹⁵	Using built-in platform capabilities and popular third-party management and security applications, Intel AMT allows IT to discover, heal, and protect computing assets on wired and wireless networks. Intel AMT is supported on platforms that have Intel® vPro™.
Intel® Rapid Storage Technology (Intel® RST) ¹⁷	Offers excellent levels of performance, responsiveness, and expandability. Take advantage of the enhanced performance and lower power consumption available with Intel® RST with one or more SATA or PCIe storage drives. With additional SATA drives, Intel® RST provides quicker access to digital photo, video, and data files with RAID 0, 5, and 10, and greater data protection against a storage disk drive failure with RAID 1, 5, and 10. Dynamic Storage Accelerator unleashes the maximum performance of Solid State Drives (SSD) when multitasking. ¹⁷
Intel® Speed Shift Technology	Delivers dramatically quicker responsiveness with single-threaded, transient (short duration) workloads, such as web browsing, by allowing the processor to more quickly select its best operating frequency and voltage for optimal performance and power efficiency.

[“monitoring a plurality of master devices coupled to a bus;”]

154. Intel Skylake processors are operated using a method that comprises monitoring a plurality of master devices coupled to a bus.

155. For example, Intel Skylake processors include a “Package Control Unit” that monitors the operations of a plurality of cores. These cores are connected by a “ring interconnect” bus. *See, e.g.*, “Power management architecture of the 2nd generation Intel Core microarchitecture, formally codenamed Sandy Bridge” [hereinafter “PCU”] at 4.



156. On information and belief, these and the other pertinent portions of the Sandy Bridge family of processors were carried over to Skylake processors in a manner that is materially the same with respect to the infringement analysis presented in this example.

[“receiving a request, from a first master device of the plurality of master devices, to change a clock frequency of a high-speed clock, the request sent from the first master device in response to a predefined change in performance of the first master device,”]

157. Intel Skylake processors are operated using a method that comprises receiving a request, from a first master device of the plurality of master devices, to change a clock frequency of a high-speed clock, the request sent from the first master device in response to a predefined change in performance of the first master device.

158. For example, each core monitors its own “applied workload” for a change in performance in that device. If a first core detects a predefined change in performance, it will

request that the system change P-state, which corresponds to the high-speed frequency of operation of the cores on the device. This detection is signaled as a request to the PCU, which is responsible for setting frequencies on the microprocessor. *See, e.g.*, SDM at 3158-3159.

14.4 HARDWARE-CONTROLLED PERFORMANCE STATES (HWP)

Intel processors may contain support for Hardware-Controlled Performance States (HWP), which autonomously selects performance states while utilizing OS supplied performance guidance hints. The Enhanced Intel Speed-Step® Technology provides a means for the OS to control and monitor discrete frequency-based operating points via the IA32_PERF_CTL and IA32_PERF_STATUS MSRs.

In contrast, HWP is an implementation of the ACPI-defined Collaborative Processor Performance Control (CPPC), which specifies that the platform enumerates a continuous, abstract unit-less, performance value scale that is not tied to a specific performance state / frequency by definition. While the enumerated scale is roughly linear in terms of a delivered integer workload performance result, the OS is required to characterize the performance value range to comprehend the delivered performance for an applied workload.

When HWP is enabled, the processor autonomously selects performance states as deemed appropriate for the applied workload and with consideration of constraining hints that are programmed by the OS. These OS-provided hints include minimum and maximum performance limits, preference towards energy efficiency or performance, and the specification of a relevant workload history observation time window. The means for the OS to override HWP's autonomous selection of performance state with a specific desired performance target is also provided, however, the effective frequency delivered is subject to the result of energy efficiency and performance optimizations.

["wherein the predefined change in performance is due to loading of the first master device as measured within a predefined time interval; and"]

159. Intel Skylake processors are operated using a method wherein the predefined change in performance is due to loading of the first master device as measured within a predefined time interval.

160. For example, the change in performance measured is due to the loading of the processor in a predefined time interval, the “moving workload history observation window.” This window can either be specified by an operating system or determined by the processor. *See, e.g.*, SDM at 3161-3162.

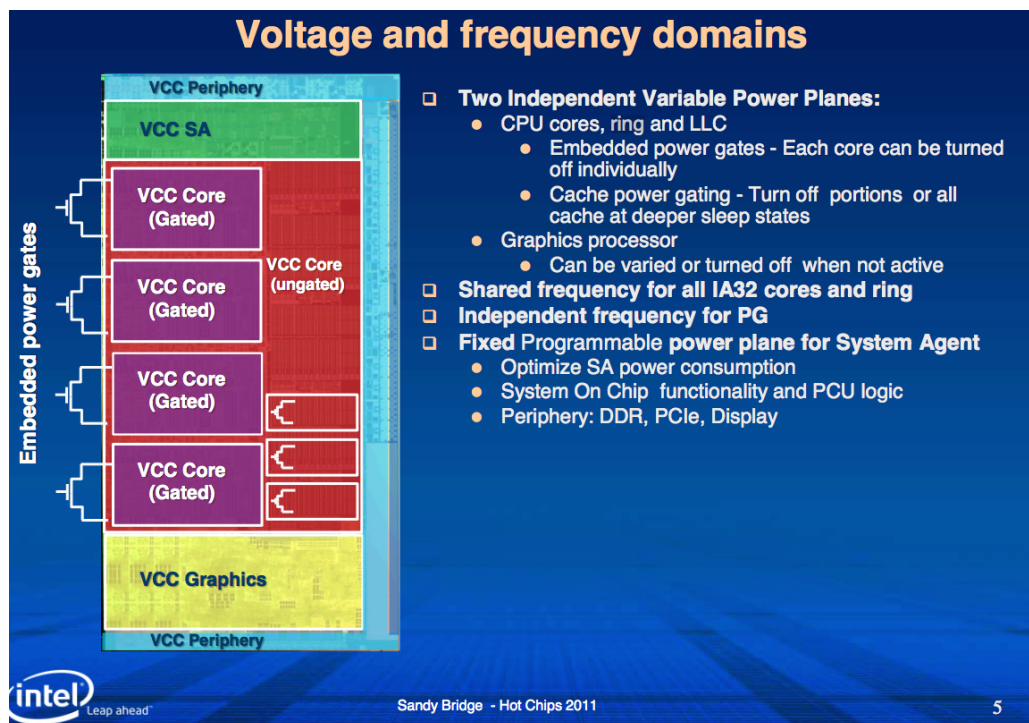
Activity_Window (bits 41:32, RW) — Conveys a hint to the HWP hardware specifying a moving workload history observation window for performance/frequency optimizations. If 0, the hardware will determine the appropriate window size. When writing a non-zero value to this field, this field is encoded in the format of bits 38:32 as a 7-bit mantissa and bits 41:39 as a 3-bit exponent value in powers of 10. The resultant value is in microseconds. Thus, the minimal/maximum activity window size is 1 microsecond/1270 seconds. Combined with the Energy_Performance_Preference input, Activity_Window influences the rate of performance increase

/ decrease. This non-zero hint only has meaning when Desired_Performance = 0. The default value of this field is 0.

["in response to receiving the request from the first master device: providing the clock frequency of the high-speed clock as an output to control a clock frequency of a second master device coupled to the bus; and"]

161. Intel Skylake processors are operated using a method that comprises, in response to receiving the request from the first master device, providing the clock frequency of the high-speed clock as an output to control a clock frequency of a second master device coupled to the bus.

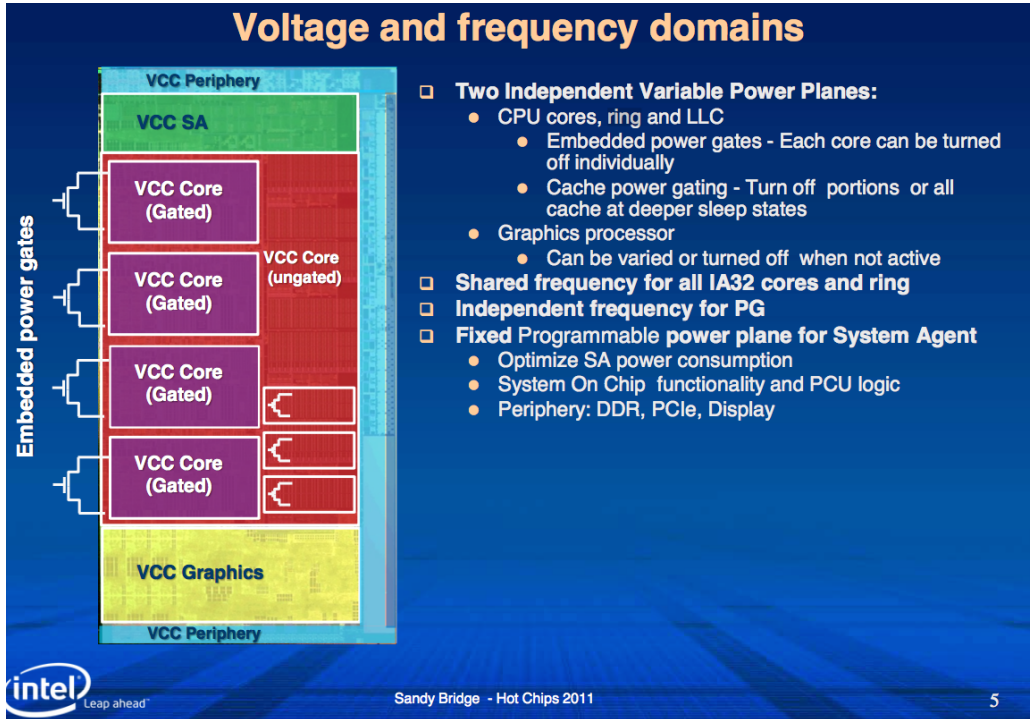
162. For example, when the PCU receives the request from the first core to change the clock frequency, the same frequency is also supplied to each other core in the platform, at least one of which is a second master device coupled to the ring interconnect bus. Each core shares a common frequency. *See, e.g.*, PCU at 5.



["providing the clock frequency of the high-speed clock as an output to control a clock frequency of the bus."]

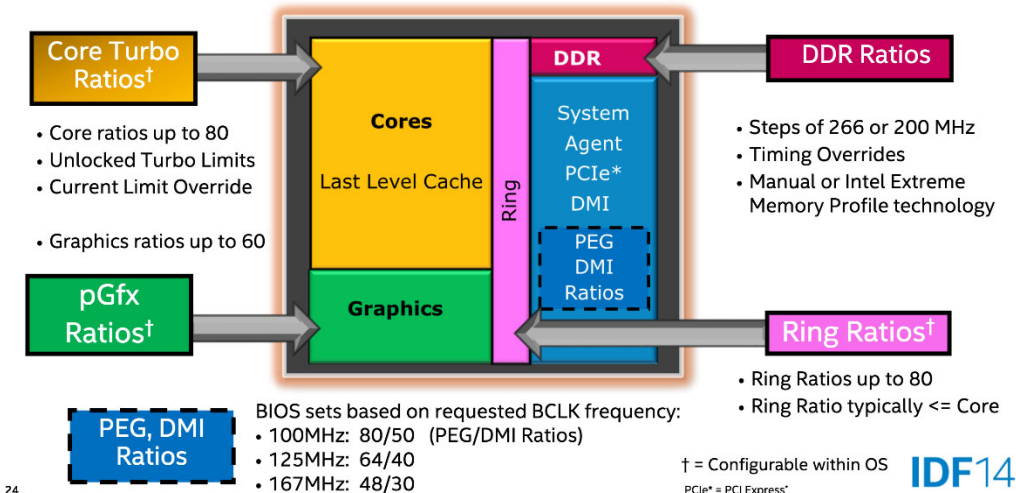
163. Intel Skylake processors are operated using a method that comprises providing the clock frequency of the high-speed clock as an output to control a clock frequency of the bus.

164. For example, when the PCU receives the request from the first core to change the clock frequency, the same frequency is also supplied to the ring interconnect itself, which shares a clock frequency with the cores. *See, e.g.,* PCU at 5.



165. *See also, e.g.,* <http://docplayer.net/38640265-Overclocking-intel-core-processors-taking-overclocking-to-the-next-level.html> at 24.

Ratio Tuning Capability Summary



166. Intel has had knowledge of the ‘759 Patent at least since the filing of this Complaint, and if it did not have actual knowledge prior to that time, it was willfully blind to the existence of the ‘759 Patent based on, for example, its publicly-known corporate policy forbidding its employees from reading patents held by outside companies or individuals, as already described above. Intel should have known that its conduct was infringing both prior to and following the filing of this case.

167. VLSI is informed and believes, and thereon alleges, that Intel actively, knowingly, and intentionally has induced infringement of the ‘759 Patent by, for example, controlling the design and manufacture of, offering for sale, selling, supplying, and otherwise providing instruction and guidance regarding the above-described products with the knowledge and specific intent to encourage and facilitate infringing uses of such products by its customers both inside and outside the United States. For example, Intel publicly provides documentation, including datasheets available through Intel’s publicly accessible ARK service and software developer’s manuals, instructing customers on uses of Intel’s products that infringe the methods of the ‘759 Patent. *See, e.g.,* <http://ark.intel.com>. On information and belief, Intel’s customers directly infringe

the '759 Patent by, for example, making, using, offering to sell, and selling within the United States, and importing into the United States, without authority or license, products containing the above-described Intel products.

168. VLSI is informed and believes, and thereon alleges, that Intel has contributed to the infringement by its customers of the '759 Patent by, without authority, importing, selling and offering to sell within the United States materials and apparatuses for practicing the claimed invention of the '759 Patent both inside and outside the United States. For example, the above-described products constitute a material part of the inventions of the '759 Patent and are not staple articles or commodities of commerce suitable for substantial noninfringing use. On information and belief, Intel knows that the above-described products constitute a material part of the inventions of the '759 Patent and are not staple articles or commodities of commerce suitable for substantial non-infringing use. On information and belief, Intel's customers directly infringe the '759 Patent by, for example, making, using, offering to sell, and selling within the United States, and importing into the United States, without authority or license, products containing the above-described Intel products.

169. As a result of Intel's infringement of the '759 Patent, VLSI has been damaged. VLSI is entitled to recover for damages sustained as a result of Intel's wrongful acts in an amount subject to proof at trial.

170. To the extent 35 U.S.C. § 287 is determined to be applicable, on information and belief its requirements have been satisfied with respect to the '759 Patent.

171. In addition, Intel's infringing acts and practices have caused and are causing immediate and irreparable harm to VLSI.

172. VLSI is informed and believes, and thereon alleges, that Intel's infringement of the '759 Patent has been and continues to be willful. As noted above, Intel has had knowledge of the '759 Patent, or, at minimum, has been willfully blind to the existence of the '759 Patent. Moreover, after service of this Complaint, Intel undisputedly had knowledge of its infringement of the '759 Patent. Intel has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for VLSI's patent rights. Thus, Intel's infringing actions have been and continue to be consciously wrongful.

SIXTH CLAIM

(Infringement of U.S. Patent No. 7,793,025)

173. VLSI re-alleges and incorporates herein by reference Paragraphs 1-172 of its Complaint.

174. The '025 Patent, entitled "Hardware managed context sensitive interrupt priority level control," was duly and lawfully issued September 7, 2010. A true and correct copy of the '025 Patent is attached hereto as Exhibit 6.

175. The '025 Patent names Robert Ehrlich, Brett W. Murdock, and Craig D. Shaw as co-inventors.

176. The '025 Patent has been in full force and effect since its issuance. VLSI owns by assignment the entire right, title, and interest in and to the '025 Patent, including the right to seek damages for past, current, and future infringement thereof.

177. The '025 Patent states that it "relates generally to interrupt controllers." Ex. 6 at 1:8-9.

178. The '025 Patent explains, for instance, that by using a "mode control selector to selectively couple different priority level assignments to a priority encoding module, context

sensitive switching of the priority levels assigned to each interrupt request can be implemented with reduced latency.” *Id.* at Abstract.

179. VLSI is informed and believes, and thereon alleges, that Intel has infringed and unless enjoined will continue to infringe one or more claims of the ‘025 Patent, in violation of 35 U.S.C. § 271, by, among other things, making, using, offering to sell, and selling within the United States, supplying or causing to be supplied in or from the United States, and importing into the United States, without authority or license, Intel products that use infringing interrupt routing technology.

180. For example, the ‘025 accused products embody every limitation of at least claim 1 of the ‘025 Patent, literally or under the doctrine of equivalents, as set forth below. The further descriptions below, which are based on publicly available information, are preliminary examples and are non-limiting.

[“1. A method for implementing interrupts in a data processing system, comprising:”]

181. Intel Ivy Bridge processors operate using a method for implementing interrupts in a data processing system.

182. For example, Ivy Bridge processors handle interrupts using the “Power Aware Interrupt Routing” system. *See, e.g.*, Intel Software Developer’s Manual [hereinafter “SDM”] at 2869.

CHAPTER 6 INTERRUPT AND EXCEPTION HANDLING

This chapter describes the interrupt and exception-handling mechanism when operating in protected mode on an Intel 64 or IA-32 processor. Most of the information provided here also applies to interrupt and exception mechanisms used in real-address, virtual-8086 mode, and 64-bit mode.

Chapter 20, "8086 Emulation," describes information specific to interrupt and exception mechanisms in real-address and virtual-8086 mode. Section 6.14, "Exception and Interrupt Handling in 64-bit Mode," describes information specific to interrupt and exception mechanisms in IA-32e mode and 64-bit sub-mode.

6.1 INTERRUPT AND EXCEPTION OVERVIEW

Interrupts and exceptions are events that indicate that a condition exists somewhere in the system, the processor, or within the currently executing program or task that requires the attention of a processor. They typically result in a forced transfer of execution from the currently running program or task to a special software routine or task called an interrupt handler or an exception handler. The action taken by a processor in response to an interrupt or exception is referred to as servicing or handling the interrupt or exception.

183. *See also, e.g.,* Power Management of the Third Generation Intel Core Micro Architecture formerly codenamed Ivy Bridge [hereinafter "Hot Chips"] at 9.

Power efficiency via interrupt routing

- ❑ **PAIR algorithm lowers power or performance impact of re-routable interrupts**
 - Compares power-state of all cores eligible to service interrupt
 - Chooses "best core" based on optimization mode (Power vs. Performance)
 - "Best Core" based on the following
 - Core C-states
 - P-state request (turbo vs. non-turbo)
- ❑ **Example: 1 core in C6 & 1 in C0**
 - Power bias will direct the interrupt to core in C0
 - Performance bias will wake the C6 core

intel Leap ahead Ivy Bridge - Hot Chips 2012 9

["receiving one or more interrupt requests from one or more interrupt sources;"]

184. Intel Ivy Bridge processors operate using a method that includes receiving one or more interrupt requests from one or more interrupt sources.

185. For example, Ivy Bridge processors can receive one or more interrupt requests by means of, *e.g.*, the INTR pin or through an APIC (Advanced Programmable Interrupt Controller). *See, e.g.*, SDM at 2871.

6.3.2 Maskable Hardware Interrupts

Any external interrupt that is delivered to the processor by means of the INTR pin or through the local APIC is called a maskable hardware interrupt. Maskable hardware interrupts that can be delivered through the INTR pin include all IA-32 architecture defined interrupt vectors from 0 through 255; those that can be delivered through the local APIC include interrupt vectors 16 through 255.

[“selectively masking the one or more interrupt requests to generate one or more pending interrupt signals;”]

186. Intel Ivy Bridge processors operate using a method that includes selectively masking the one or more interrupt requests to generate one or more pending interrupt signals.

187. For example, Intel Ivy Bridge processors can mask the one or more maskable interrupts to generate pending interrupts to the cores. *See, e.g.*, SDM at 2874.

6.8 ENABLING AND DISABLING INTERRUPTS

The processor inhibits the generation of some interrupts, depending on the state of the processor and of the IF and RF flags in the EFLAGS register, as described in the following sections.

6.8.1 Masking Maskable Hardware Interrupts

The IF flag can disable the servicing of maskable hardware interrupts received on the processor's INTR pin or through the local APIC (see Section 6.3.2, “Maskable Hardware Interrupts”). When the IF flag is clear, the processor inhibits interrupts delivered to the INTR pin or through the local APIC from generating an internal interrupt request; when the IF flag is set, interrupts delivered to the INTR or through the local APIC pin are processed as normal external interrupts.

188. The selectively masked interrupts are stored as pending interrupt signals to be processed. *See, e.g.*, SDM at 2876.

6.9 PRIORITY AMONG SIMULTANEOUS EXCEPTIONS AND INTERRUPTS

If more than one exception or interrupt is pending at an instruction boundary, the processor services them in a predictable order. Table 6-2 shows the priority among classes of exception and interrupt sources.

[“providing a plurality of interrupt priority storage devices comprising a first interrupt priority storage device for storing priority level information associated with a first system mode, and a second interrupt priority storage device for storing priority level information associated with a second system mode; and”]

189. Intel Ivy Bridge processors operate using a method that includes providing a plurality of interrupt priority storage devices comprising a first interrupt priority storage device for storing priority level information associated with a first system mode, and a second interrupt priority storage device for storing priority level information associated with a second system mode.

190. For example, each of the plurality of cores in Intel Ivy Bridge processors include a “local APIC” (Advanced Programmable Interrupt Control) that stores interrupts. *See, e.g.*, SDM at 3047.

10.1 LOCAL AND I/O APIC OVERVIEW

Each local APIC consists of a set of APIC registers (see Table 10-1) and associated hardware that control the delivery of interrupts to the processor core and the generation of IPI messages. The APIC registers are memory mapped and can be read and written to using the MOV instruction.

Local APICs can receive interrupts from the following sources:

- **Locally connected I/O devices** — These interrupts originate as an edge or level asserted by an I/O device that is connected directly to the processor’s local interrupt pins (LINT0 and LINT1). The I/O devices may also be connected to an 8259-type interrupt controller that is in turn connected to the processor through one of the local interrupt pins.
- **Externally connected I/O devices** — These interrupts originate as an edge or level asserted by an I/O device that is connected to the interrupt input pins of an I/O APIC. Interrupts are sent as I/O interrupt messages from the I/O APIC to one or more of the processors in the system.
- **Inter-processor interrupts (IPIs)** — An Intel 64 or IA-32 processor can use the IPI mechanism to interrupt another processor or group of processors on the system bus. IPIs are used for software self-interrupts, interrupt forwarding, or preemptive scheduling.
- **APIC timer generated interrupts** — The local APIC timer can be programmed to send a local interrupt to its associated processor when a programmed count is reached (see Section 10.5.4, “APIC Timer”).
- **Performance monitoring counter interrupts** — P6 family, Pentium 4, and Intel Xeon processors provide the ability to send an interrupt to its associated processor when a performance-monitoring counter overflows (see Section 18.6.3.5.8, “Generating an Interrupt on Overflow”).
- **Thermal Sensor interrupts** — Pentium 4 and Intel Xeon processors provide the ability to send an interrupt to themselves when the internal thermal sensor has been tripped (see Section 14.7.2, “Thermal Monitor”).

191. Each local APIC also masks and stores priority information for interrupts. Therefore, each multicore Ivy Bridge processor contains a plurality of interrupt priority storage devices. See, e.g., SDM at 3051.

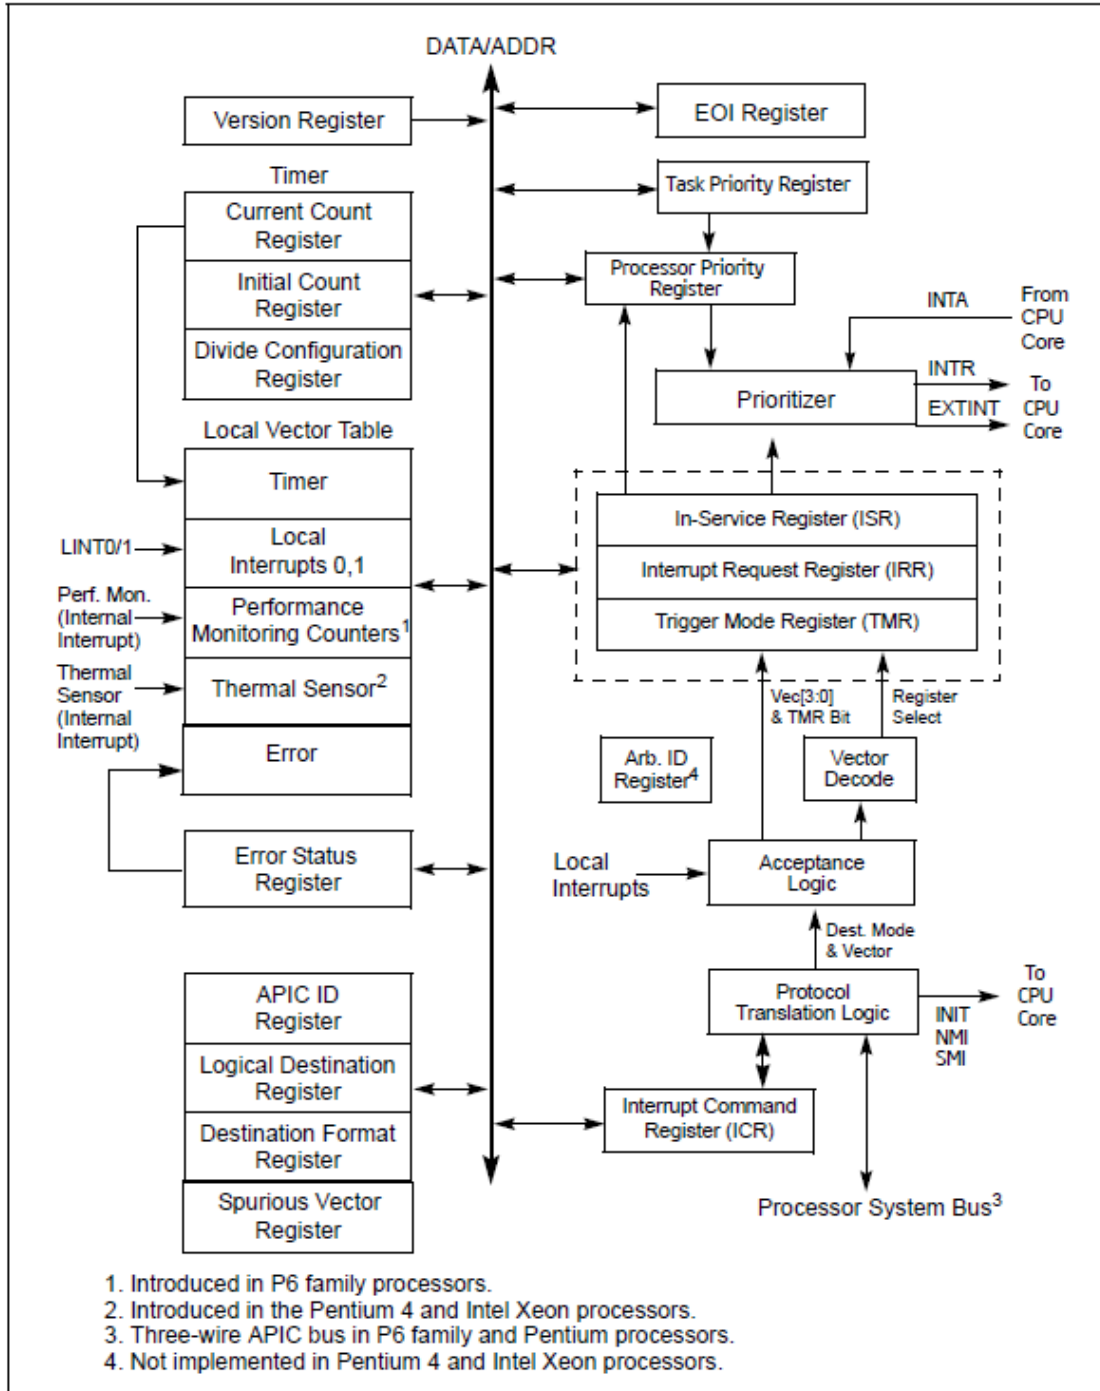


Figure 10-4. Local APIC Structure

192. Additionally, interrupts are routed to a particular core depending on parameters such as loading, power state, and power mode. Each core's APIC is associated with a given system mode depending on those parameters. *See, e.g., Hot Chips at 9.*

Power efficiency via interrupt routing

- PAIR algorithm lowers power or performance impact of re-routable interrupts
 - Compares power-state of all cores eligible to service interrupt
 - Chooses “best core” based on optimization mode (Power vs. Performance)
 - “Best Core” based on the following
 - Core C-states
 - P-state request (turbo vs. non-turbo)
- Example: 1 core in C6 & 1 in C0
 - Power bias will direct the interrupt to core in C0
 - Performance bias will wake the C6 core

intel Leap ahead Ivy Bridge - Hot Chips 2012 9

“providing a plurality of interrupt priority storage devices comprising a first interrupt priority storage device for storing priority level information associated with a first system mode for each of the one or more interrupt requests, and a second interrupt priority storage device for storing priority level information associated with a second system mode for each of the one or more interrupt requests; and”]

193. Intel Ivy Bridge processors operate using a method that includes providing a plurality of interrupt priority storage devices comprising a first interrupt priority storage device for storing priority level information associated with a first system mode for each of the one or more interrupt requests, and a second interrupt priority storage device for storing priority level information associated with a second system mode for each of the one or more interrupt requests.

194. For example, each local APIC can process events from each of the one or more interrupt requests. See the evidence cited above in connection with the previous element.


“selectively coupling, in response to a mode control signal, one of the plurality of interrupt priority storage devices to provide logic circuitry with priority level information corresponding to the mode control signal.”]

195. Intel Ivy Bridge processors operate using a method that includes selectively coupling, in response to a mode control signal, one of the plurality of interrupt priority storage devices to provide logic circuitry with priority level information corresponding to the mode control signal.

196. For example, Intel Ivy Bridge processors selectively couple, in response to power mode and power state control signals, interrupts including priority level information to a given core, and that core’s local APIC. *See, e.g., Hot Chips at 9.* Each core’s APIC is associated with a given system mode depending on parameters such as loading, power mode, and power state. *Id.*

Power efficiency via interrupt routing

- ❑ **PAIR algorithm lowers power or performance impact of re-routable interrupts**
 - Compares power-state of all cores eligible to service interrupt
 - Chooses “best core” based on optimization mode (Power vs. Performance)
 - “Best Core” based on the following
 - Core C-states
 - P-state request (turbo vs. non-turbo)
- ❑ **Example: 1 core in C6 & 1 in C0**
 - Power bias will direct the interrupt to core in C0
 - Performance bias will wake the C6 core

 Leap ahead Ivy Bridge - Hot Chips 2012 9

["where the logic circuitry uses the provided priority level information to prioritize one or more of the pending interrupt signals and"]

197. Intel Ivy Bridge processors operate using a method where the logic circuitry uses the provided priority level information to prioritize one or more of the pending interrupt signals.

198. For example, the local APIC performs interrupt prioritization. *See, e.g.*, SDM at 3051.

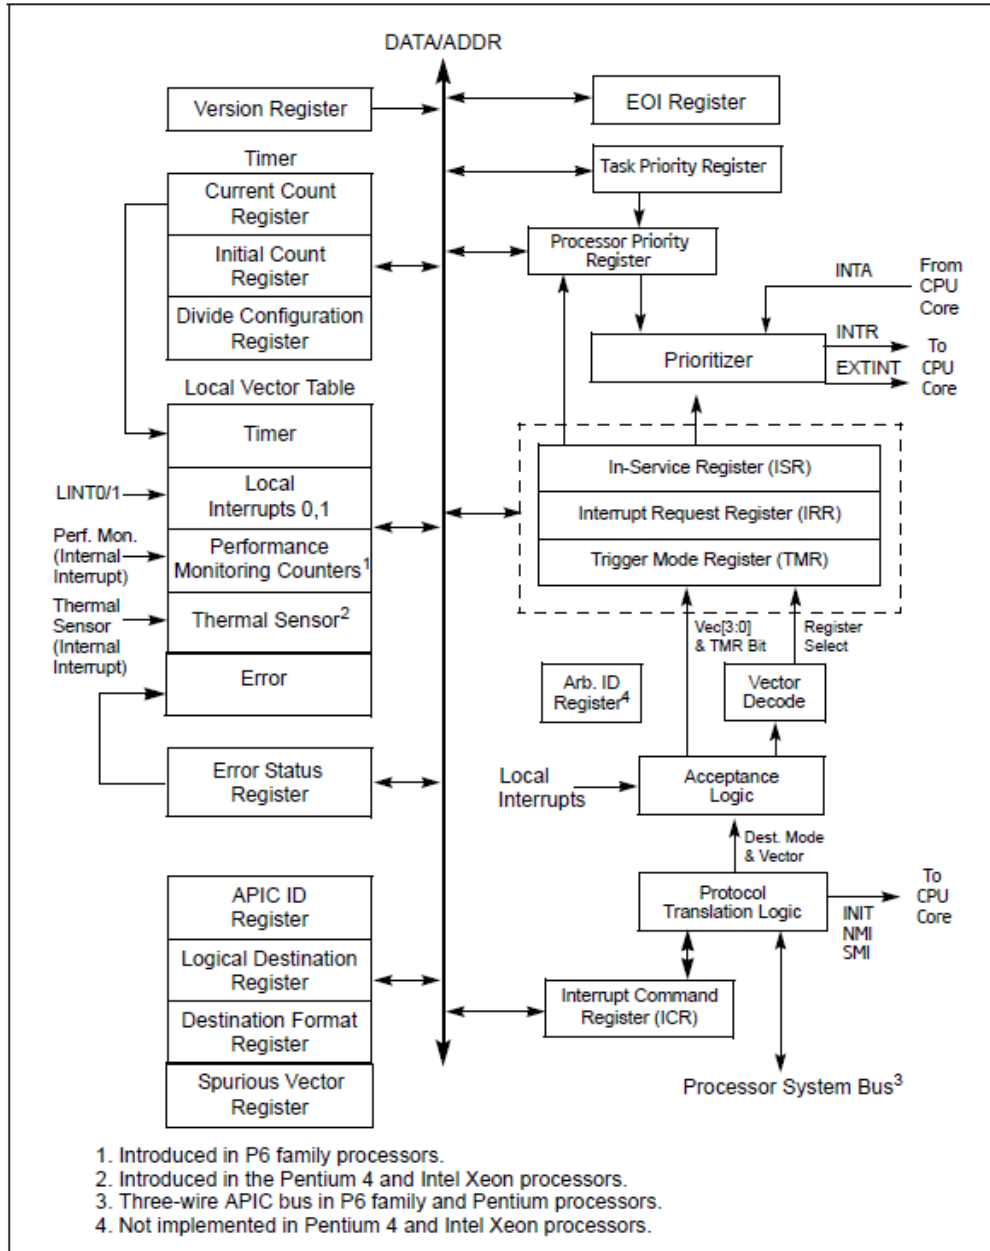


Figure 10-4. Local APIC Structure

199. See also, e.g., SDM at 3074.

10.8.3 Interrupt, Task, and Processor Priority

Each interrupt delivered to the processor through the local APIC has a priority based on its vector number. The local APIC uses this priority to determine when to service the interrupt relative to the other activities of the processor, including the servicing of other interrupts.

["also provides an interrupt request signal which will cause an interrupt to occur in the data processing system."]

200. Intel Ivy Bridge processors operate using a method that also provides an interrupt request signal which will cause an interrupt to occur in the data processing system.

201. For example, the local APIC generates an interrupt request signal to a core that causes interrupt handlers to be called in that core. *See, e.g.*, SDM at 2879.

6.12 EXCEPTION AND INTERRUPT HANDLING

The processor handles calls to exception- and interrupt-handlers similar to the way it handles calls with a CALL instruction to a procedure or a task. When responding to an exception or interrupt, the processor uses the exception or interrupt vector as an index to a descriptor in the IDT. If the index points to an interrupt gate or trap gate, the processor calls the exception or interrupt handler in a manner similar to a CALL to a call gate (see Section 5.8.2, "Gate Descriptors," through Section 5.8.6, "Returning from a Called Procedure"). If index points to a task gate, the processor executes a task switch to the exception- or interrupt-handler task in a manner similar to a CALL to a task gate (see Section 7.3, "Task Switching").

202. *See also, e.g.*, SDM at 3051.

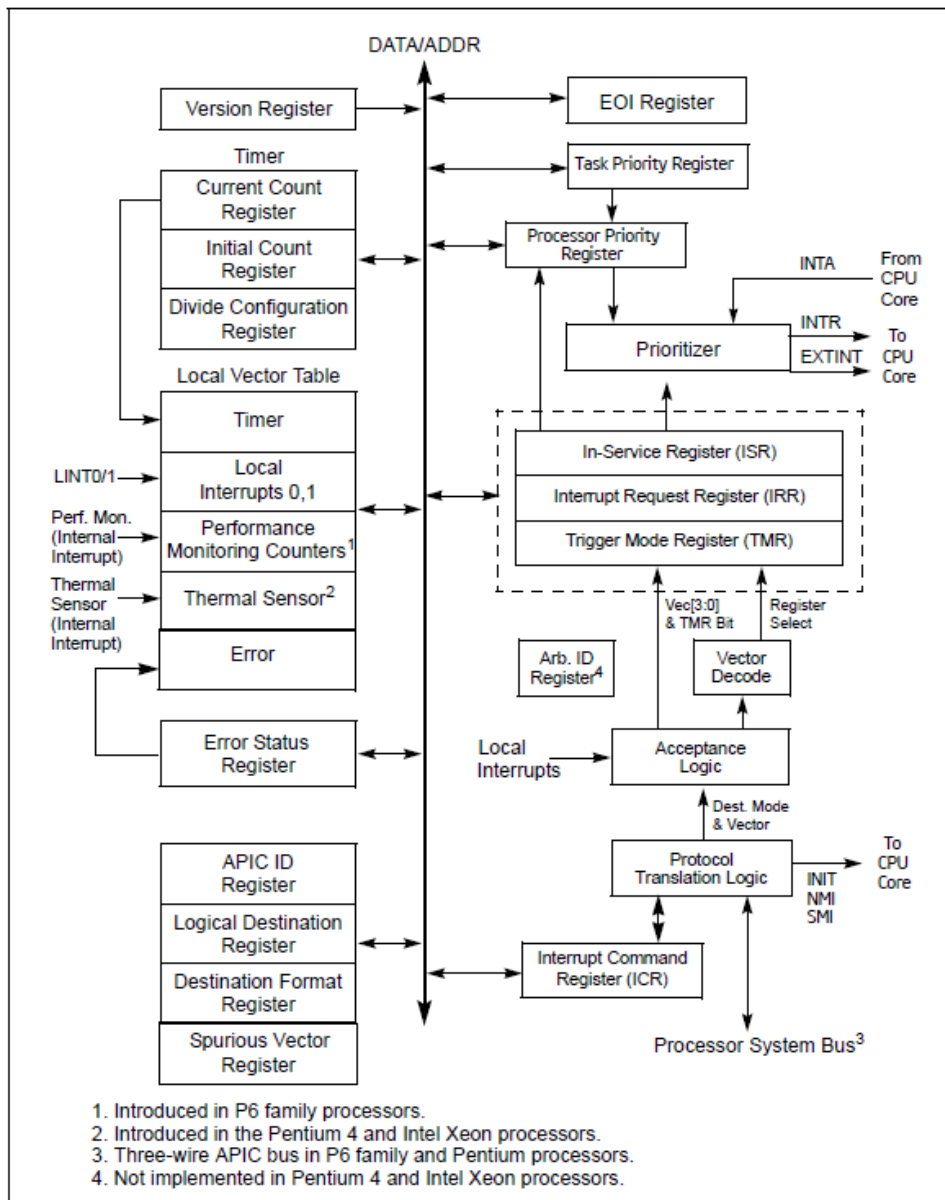


Figure 10-4. Local APIC Structure

203. Intel has had knowledge of the ‘025 Patent at least since the filing of this Complaint, and if it did not have actual knowledge prior to that time, it was willfully blind to the existence of the ‘025 Patent based on, for example, its publicly-known corporate policy forbidding its employees from reading patents held by outside companies or individuals, as already described above. Intel should have known that its conduct was infringing both prior to and following the filing of this case.

204. VLSI is informed and believes, and thereon alleges, that Intel actively, knowingly, and intentionally has induced infringement of the '025 Patent by, for example, controlling the design and manufacture of, offering for sale, selling, supplying, and otherwise providing instruction and guidance regarding the above-described products with the knowledge and specific intent to encourage and facilitate infringing uses of such products by its customers both inside and outside the United States. For example, Intel publicly provides documentation, including datasheets available through Intel's publicly accessible ARK service and software developer's manuals, instructing customers on uses of Intel's products that infringe the methods of the '025 Patent. *See, e.g.,* <http://ark.intel.com>. On information and belief, Intel's customers directly infringe the '025 Patent by, for example, making, using, offering to sell, and selling within the United States, and importing into the United States, without authority or license, products containing the above-described Intel products.

205. VLSI is informed and believes, and thereon alleges, that Intel has contributed to the infringement by its customers of the '025 Patent by, without authority, importing, selling and offering to sell within the United States materials and apparatuses for practicing the claimed invention of the '025 Patent both inside and outside the United States. For example, the above-described products constitute a material part of the inventions of the '025 Patent and are not staple articles or commodities of commerce suitable for substantial noninfringing use. On information and belief, Intel knows that the above-described products constitute a material part of the inventions of the '025 Patent and are not staple articles or commodities of commerce suitable for substantial non-infringing use. On information and belief, Intel's customers directly infringe the '025 Patent by, for example, making, using, offering to sell, and selling within the United States,

and importing into the United States, without authority or license, products containing the above-described Intel products.

206. As a result of Intel's infringement of the '025 Patent, VLSI has been damaged. VLSI is entitled to recover for damages sustained as a result of Intel's wrongful acts in an amount subject to proof at trial.

207. To the extent 35 U.S.C. § 287 is determined to be applicable, on information and belief its requirements have been satisfied with respect to the '025 Patent.

208. In addition, Intel's infringing acts and practices have caused and are causing immediate and irreparable harm to VLSI.

209. VLSI is informed and believes, and thereon alleges, that Intel's infringement of the '025 Patent has been and continues to be willful. As noted above, Intel has had knowledge of the '025 Patent, or, at minimum, has been willfully blind to the existence of the '025 Patent. Moreover, after service of this Complaint, Intel undisputedly had knowledge of its infringement of the '025 Patent. Intel has deliberately continued to infringe in a wanton, malicious, and egregious manner, with reckless disregard for VLSI's patent rights. Thus, Intel's infringing actions have been and continue to be consciously wrongful.

PRAYER FOR RELIEF

WHEREFORE, VLSI prays for judgment against Intel as follows:

- A. That Intel has infringed, and unless enjoined will continue to infringe, each of the Asserted Patents;
- B. That Intel has willfully infringed each of the Asserted Patents;

- C. That Intel pay VLSI damages adequate to compensate VLSI for Intel's infringement of each of the Asserted Patents, together with interest and costs under 35 U.S.C. § 284;
- D. That Intel be ordered to pay prejudgment and post-judgment interest on the damages assessed;
- E. That Intel pay VLSI enhanced damages pursuant to 35 U.S.C. § 284;
- F. That Intel be ordered to pay supplemental damages to VLSI, including interest, with an accounting, as needed;
- G. That Intel be enjoined from infringing the Asserted Patents, or if its infringement is not enjoined, that Intel be ordered to pay ongoing royalties to VLSI for any post-judgment infringement of the Asserted Patents;
- H. That this is an exceptional case under 35 U.S.C. § 285, and that Intel pay VLSI's attorneys' fees and costs in this action; and
- I. That VLSI be awarded such other and further relief, including equitable relief, as this Court deems just and proper.

DEMAND FOR JURY TRIAL

Pursuant to Federal Rule of Civil Procedure 38(b), VLSI hereby demands a trial by jury on all issues triable to a jury.

Dated: March 1, 2019

Respectfully submitted,

FARNAN LLP

/s/ Brian E. Farnan
Brian E. Farnan (Bar No. 4089)
Michael J. Farnan (Bar No. 5165)
919 N. Market Street, 12th Street
Wilmington, Delaware 19801
Telephone: (302) 777-0300
Facsimile: (302) 777-0301

Attorneys for Plaintiff VLSI Technology LLC