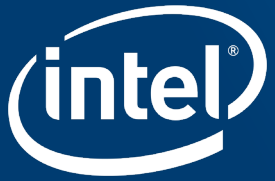

Comet Lake Platform Intel® Turbo Boost Max Technology 3.0

Technical Advisory

WW14, April 2020

Document Number: 621487

Legal Disclaimer

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Intel technologies may require enabled hardware, specific software, or services activation. Check with your system manufacturer or retailer.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or visit www.intel.com/design/literature.htm.

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

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

Intel® Turbo Boost Maximum Technology 3.0

- Intel® Turbo Boost Maximum Technology 3.0. also known as ITBMT 3.0, allows the Operating System (Windows* 10 RS1 or newer) to take advantage of the maximum turbo frequency of each individual core of the processor. In this state, a specific processor core running a workload runs at a higher frequency than the marked maximum turbo frequency of the processor package.
- Individual asymmetric core maximum frequency information is exposed to the operating system via ACPI 5.1 (or newer) Collaborative Processor Performance Control (CPPC2) interface.
- The general value proposition for ITBMT 3.0 is the performance benefit delivered by exposing all cores' native frequency capability instead of limiting the top-end performance of all cores to that of the lowest performance core.
- ITBMT 3.0 performance benefit is most visible in low (1-2) thread count workloads. Highly threaded workloads do not benefit from ITBMT 3.0.

Notes:

- The Logical Core IDs of the faster cores are sample specific.
- All samples within a SKU (SSPEC/QDF) have the same per-core frequencies (in different Cores).
- This document does not list performance numbers.

Power and Thermal Validation

- ITBMT 3.0 enables 1-2 cores to run at higher frequency than other cores. The higher frequency is achievable only if the other cores are not in active (Core C6 state).
- The faster cores are chosen during manufacturing so that they do not require more voltage or power than the slower cores while operating at higher frequency.
- In other words, ITBMT 3.0 does not require any changes to power delivery, cooling solution, etc.,
- No additional checks are required with respect to power and thermal validation.

Note: The Core ID of the faster cores is not fixed within a SKU. The faster cores' IDs are sample specific.

General BIOS Requirements – Enumeration

Enumeration of the feature and per core maximum ratios is accomplished via communication with the Overclocking mailbox processor interface, which is described in the Performance Tuning Guide.

Note: Read access to this interface is available on processors that do not support overclocking (non-K SKUs).

Native OS support for diverse maximum performance cores is indicated by a new ACPI platform-wide OSC bit (12) to defined in ACPI 6.2.

12	**Diverse CPPC Highest Support**	**This bit is set if OSPM supports optimizations for processors with diverse CPPC Highest performance capabilities and can process associated processor device change notifications.**

The means of exposing the diverse core capabilities to the OS is via the _CPC object's Highest Performance element.

Note: Different encoding is required for overclocking vs. non-overclocking platforms.

Non Overclocking Platform Requirements

BIOS reads the per-core max ratio from the OC Mailbox command 1Ch and populate the _CPC Highest Performance element as an **Integer** appropriately (as opposed to a register) for each logical processor.

- Integer Value = core specific max ratio (from OC Mailbox 1Ch – current maximum core ratio field)

Note: Legacy behavior can be achieved via BIOS setup option that uses the lowest maximum ratio of all cores in the population of the _CPC Highest Performance field.

_PSS P0 population can be core specific or a high value can be used (Example: 0xFF Recommended) as the processor internally clips OS requests to the maximum supported value.

Overclocking Platform Requirements

BIOS may implement the same requirements as non-overclocking systems if desired. However to support on-the-fly overclocking (i.e. post boot OC changes), a different approach must be taken in accordance with whether the OS does support optimizations for diverse cores natively as follows:

- If the OS does not support diverse core optimizations (platform-wide _OSC bit 12 is clear), then the BIOS should populate high values (FFh) in _PSS P0 and the _CPC Highest Performance element for each logical processor so that the OS will write values that will be clipped by the HW to the current per core max ratio.
- If the OS supports diverse core optimizations (platform-wide _OSC bit 12 is set), then the diversity must be exposed to the OS and the OS will re-read _CPC Highest Performance value upon receiving an ACPI Notify (0x85) on the processor device. However, the OS will only perform this re-read when _CPC Highest Performance is encoded as a register, i.e., the OS will not re-evaluate the _CPC object itself.

OC Mailbox Command 1Ch Details

- BIOS identifies the maximum per-core frequency via the Overclocking Mailbox (MSR 150h) command 1Ch (OC_MAILBOX_READ_PER_CORE_RATIO_LIMITS_CAPABILITY_CMD).

  The per-core max ratio is found in bits 7:0 (current setting) of the result.

- Usage: Execute from each logical core to obtain its maximum ratio:
  - Write MSR 150h: 0x8000001C00000000
  - Read MSR 150h:
    - Bits 39:32 (error code) should be zero
    - Bits 7:0 provide the Core Max Ratio

Parameter		Definition
Command		1Ch (read)
Param1		00h
Param2		00h
Data	7:0	Current Core max ratio
	15:8	Fused Max ratio (default P_0 for this core)
	23:16	Favored Core Index in the favored core list (Mailbox ID)
	31:24	Reserved
Completion Codes		0 - success

Intel® Turbo Boost Maximum Validation Guidance

To capture ITBMT 3.0 higher frequencies, run one of the options:

- Run Thermal Analysis Tool (TAT) with log capture.
- Run Cinebench R15/R20 single threaded or any other single threaded workload like SPEC (ST).
- View the log generated by TAT as shown below:

CML-S RVP with Core i9-10900K Example:

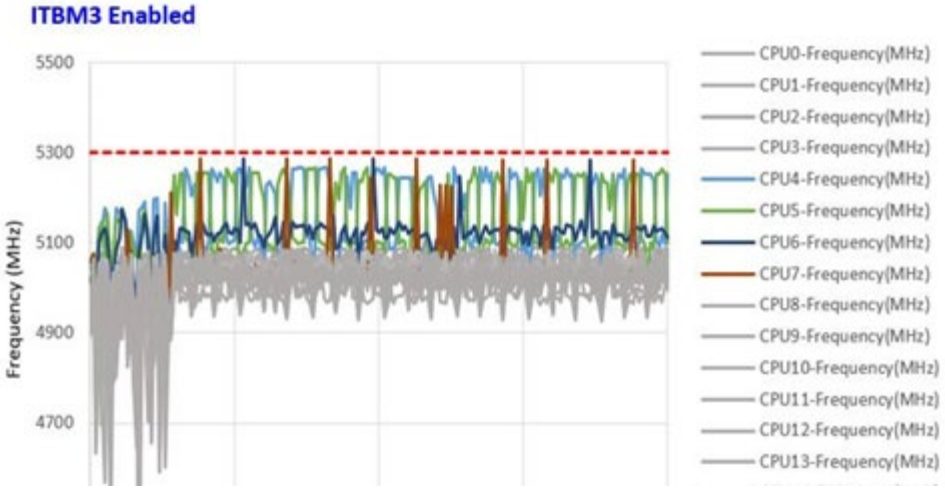

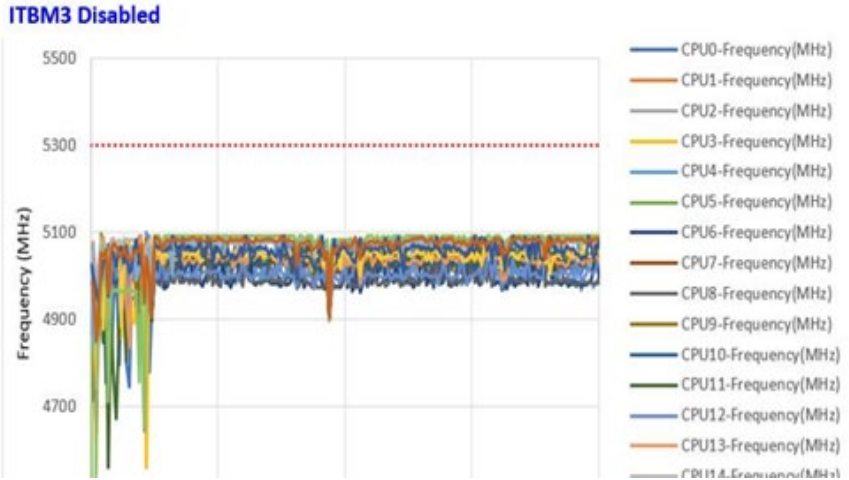

P-State Resolution with ITBMT 3.0

P-States in the processor follow these phases:

- P-State request by the Operating System
- P-State shaping by Power Management FW (pCode)
- P-State Limits by pCode
- All active cores are set to the same resolved frequency (P-state)

Following slides provide more details on these phases.

Note: P-State resolution for GT and the Ring are outside the scope of this document.

P-State Request

The OS makes a P-State request via these registers:

- Intel SpeedShift: IA32_HWP_REQUEST (MSR 774h)
  - Provides an operating range (minimum/maximum frequency).
  - The minimum and maximum frequency is derived from the _CPC table BIOS exposes.
  - Minimum and maximum frequencies are dependent on the OS Power Plan:
    - Maximum is _CPC's Highest Performance.
    - Minimum can be the base frequency or the LFM frequency.
  - Energy Performance Preference (EPP) field tells the processor which is preferred power saving or performance.
- Legacy P-States: IA32_PERF_CTL (MSR 199h) – Not supported with ITBMT3 and not referenced further in this document.

P-State Shaping

The requested P-State range (MSR 774h) is clipped in the following manor:

- Minimum/Maximum is clipped to [Pn,P0] range without error.
- Minimum frequency request may be raised by Race to Halt (RTH).
- Maximum frequency request may be decreased by Energy Efficient Turbo (EET).

Notes:

- Both RTH and EET can be disabled by BIOS and at runtime (both in MSR_POWER_CTL - MSR 1FCh). It is recommended to disable EET in desktops.
- EET is disabled when “High Performance” OS Power plan is selected.

P-State Limits

Frequency is limited (throttled) via various limits.

- Each limit mechanism produces a maximum operating frequency.
- The final frequency is the minimum of all limits.
- There are power, current, voltage, thermal and frequency limits. There are several limits in each category.

  Note: The operation of the various limits is outside the scope of this document.
- Processors with ITBMT 3.0 have the following frequency (P-State) limits
  - Turbo Ratio Limits (MSR 1ADh): Maximum frequency depending how many cores are active (not in C6/C7). It does not matter which cores are active, just how many. Turbo Ratio Limits can have an offset when software is executing AVX instructions.
  - OC Max Ratio (overclocking only): Set via OC Mailbox command 0x11 (in K-SKUs only). The default OC Max Ratio is the same as Turbo Ratio Limits for 1 active core.
  - Per Core maximum frequency (ITBMT3): Different maximum frequency per physical core.
    - Can be read via OC Mailbox command 0x1C (all SKUs).
    - Can be set via OC Mailbox command 0x1D (K SKUs only).
    - Example: Core #1 can have maximum of 50 and Core #6 can have maximum of 52.
    - This P-state limit is the lowest per-core frequency of all active cores.
- The limit reason is reported in IA_PERF_LIMIT_REASON (MSR 64Fh).

Known and Potential Issues

Feature Enumeration

- Due to an Erratum CML132, the Comet Lake Processor does not enumerate ITBMT 3.0 via the CPUID instruction (leaf 06h, EAX[14]).
- BIOS can enumerate ITBMT 3.0 by comparing the values returned by OC Mailbox command 1Ch as described in previous slides of this document.
- Software cannot accurately enumerate ITBMT 3.0 due to the OC Mailbox interface; can be used to modify per-core maximum frequency in Overclocking SKUs.

Unexpected Wakeup from S0ix When ITBMT3 Enabled (sighting ID 22010169198)

- When ITBMT3 was enabled, the S0ix wakeup timing was misconfigured at BIOS for longer time than expected.
- Fixed in BIOS version v9.0.30.50.

Potential issue: Unexpected Wakeup from S0ix When ITBMT3 Enabled with Overclocking Enabled

- Similar to the non-overclocking sighting.
- New issue, no root cause.

References

Document Title	Document Number
Skylake, Kaby Lake, Coffee Lake, Whiskey Lake, and Comet Lake Processor Family Core and Uncore BIOS Specification	550049
Turbo and Thermal Power Management Guide for Intel® Core™ Based Processors User Guide	571040
Comet Lake Processor External Design Specification, Volume 1 of 2	606599
Comet Lake Platform Performance Tuning User Guide	610996
10th Generation Intel® Core™ Processor Family Specification Update - NDA	613849