congatec Application Note #40

Affected Products	All Intel x86 products
Subject	Real-Time Applications: BIOS Setup Settings and Advice
Confidential/Public	Public
Author	CJR

Revision History

Revision	Date (yyyy-mm-dd)	Author	Changes
1.0	2020-01-20	CJR	Initial release of document
1.1	2021-01-11	CJR	Changed recommended Legacy IO Low Latency setting on page 7

Preface

This application note provides a list of recommended BIOS setup settings for real-time applications, describes each setting and informs about some additional topics related to real-time operating systems (RTOS).

Disclaimer

The information contained within this Application Note, including but not limited to any product specification, is subject to change without notice.

congatec AG provides no warranty with regard to this Application Note or any other information contained herein and hereby expressly disclaims any implied warranties of merchantability or fitness for any particular purpose with regard to any of the foregoing. congatec AG assumes no liability for any damages incurred directly or indirectly from any technical or typographical errors or omissions contained herein or for discrepancies between the product and the Application Note. In no event shall congatec AG be liable for any incidental, consequential, special, or exemplary damages, whether based on tort, contract or otherwise, arising out of or in connection with this Application Note or any other information contained herein or the use thereof.

Intended Audience

This Application Note is intended for technically qualified personnel. It is not intended for general audiences.

Electrostatic Sensitive Device

All congatec AG products are electrostatic sensitive devices and are packaged accordingly. Do not open or handle a congatec AG product except at an electrostatic-free workstation. Additionally, do not ship or store congatec AG products near strong electrostatic, electromagnetic, magnetic, or radioactive fields unless the device is contained within its original manufacturer's packaging. Be aware that failure to comply with these guidelines will void the congatec AG Limited Warranty.

Technical Support

congatec AG technicians and engineers are committed to providing the best possible technical support for our customers so that our products can be easily used and implemented. We request that you first visit our website at www.congatec.com for the latest documentation, utilities and drivers, which have been made available to assist you. If you still require assistance after visiting our website then contact our technical support department by email at support@congatec.com

Symbols

The following are symbols used in this application note.

Note

Notes call attention to important information that should be observed.

Caution

Cautions warn the user about how to prevent damage to hardware or loss of data.

Warning

Warnings indicate that personal injury can occur if the information is not observed.

Copyright Notice

Trademarks

Terminology

Term	Description
CPU	Central Processing Unit
GPU	Graphics Processing Unit
IMC	Integrated Memory Controller
IGD	Integrated Graphics Device
PCH	Platform Controller Hub
SKU	Stock Keeping Unit
DMI	Direct Media Interface
ASPM	Active State Power Management
cTDP	Configurable Thermal Design Power
TCC	Thermal Control Circuit or Time Coordinated Computing
CAT	Cache Allocation Technology
TSN	Time Sensitve Networking
QoS	Quality of Service
RTOS	Real-time Operating System
OS	Operating System

1 Introduction

The purpose of Real-Time Operating Systems (RTOS) is to service interrupts from time critical peripheral devices (typically a field bus device) within the required time limit. The specific time limit may vary depending on the system requirements (hard versus soft real-time) but may not be exceeded at any time during operation.

In contrast to RTOS, generic Operating Systems (OS) like Windows do not have such a strict time limit for servicing interrupts. Power management and optimized performance are typically prioritized. But any feature that saves power or increases short-term performance can increase latency and therefore the limit for servicing interrupts may be exceeded. This is unacceptable for real-time applications because they require predictable and constant performance.

Section 2 of this application note describes the most critical BIOS setup settings for RTOS. To optimize the RTOS for a specific real-time application, it is important to understand the impact of each setting.

Section 3 of this application note gives additional advice on how to optimize RTOS performance.

2 Recommended BIOS Setup Settings

Many BIOS features have a significant impact on the interrupt latency and jitter. Even a single incorrect setting could lead to total system failure.

Note

The available features depend on the specific BIOS. A BIOS may provide relevant features not covered in this application note or not provide features covered in this application note. In general, it is recommended to disable all power-saving and short-term performance enhancing features because the transition between power states introduces latency.

2.1 Overview

The table below lists the recommended BIOS setup setting for real-time applications and the default congatec BIOS setup setting:

	BIOS Setup Setting	Recommended	Default
CPU	Enhanced Intel Speed Step Technology (EIST)	Disabled	Enabled
	Intel Speed Shift Technology (SST)	Disabled	Enabled
	Intel Turbo Boost Technology	Disabled	Enabled
	C-States	Disabled	Disabled
	Thermal Control Circuit (TCC) Activation Offset	0	0
	Intel Hyper-Threading Technology (HTT)	Disabled	Enabled
RAM	System Agent Enhanced Intel SpeedStep (SA GV)	Fixed High	Fixed High
IGD	GT PM Support	Disabled	Disabled
	Render Standby (RC6)	Disabled	Enabled
LPM	DMI Link ASPM Control	Disabled	Disabled
	PCI Express Link ASPM Control	Disabled	Disabled
	Aggressive LPM Support	Disabled	Disabled
	Legacy IO Low Latency	Enabled	Disabled

Note

The BIOS setup settings listed above are only recommendations and may not be ideal for your real-time application. Read the sub-sections below for a description of each setting and additional notes.

2.2 CPU

The CPU performance can be influenced by various power-saving and short-term performance-enhancing features. Such features should be disabled because they introduce latency (e.g. increase response time to a field bus device interrupt).

The following settings can usually be found in the BIOS setup program under submenu “CPU Configuration”, “CPU Thermals and Performance”, or similar.

2.2.1 Enhanced Intel Speed Step Technology (EIST)

If enabled, the OS can change the CPU core frequency and voltage depending on the workload. The CPU cores are not executing instructions during the frequency switching. This feature introduces latency.

2.2.2 Intel Speed Shift Technology (SST)

If enabled, the CPU can change its core frequency and voltage depending on the workload. The CPU cores are not executing instructions during the frequency switching. This feature introduces latency.

2.2.3 Intel Turbo Boost Technology

If enabled, the CPU can dynamically increase its frequency above the rated operating frequency. The CPU cores are not executing instructions during the frequency switching. This feature introduces latency.

Note

The Intel Turbo Boost Technology can also challenge the thermal design, possibly triggering the Thermal Control Circuit (TCC) to throttle the CPU which results in alternating CPU performance.

2.2.4 C-States

Idle States (C-states) are used to save power when the CPU is idle. C0 is the operational state, meaning that the CPU is doing useful work. C1 is the first idle state, C2 the second, and so on, where more power saving actions are taken for numerically higher C-states. These idle states dramatically increase the CPU response time because it takes quiet a long time to resume the CPU from these low power states.

2.2.5 Thermal Control Circuit (TCC) Activation Offset

The TTC Activation Offset defines the CPU die temperature at which the CPU throttling mechanism is started. For optimized real-time behaviour, CPU throttling should not be started at temperatures lower than the maximum specified die temperature.

The set value is subtracted from the max. specified die temperature (100°C). In order to prevent CPU throttling below 100°C, the default value of 0 should not be changed. Proper

CPU and system cooling is also beneficial for good real-time behaviour because overheating always results in some sort of performance reduction.

Note

It is not possible to disable CPU thermal throttling. It is only possible to prevent it via proper thermal design. For more information, refer to section 2.1 "Thermal Design".

2.2.6 Intel Hyper-Threading Technology (HTT)

Delivers two processing threads per physical core. Highly threaded applications can get more work done in parallel, completing tasks sooner. But executing real-time tasks on logical CPU cores might have negative impacts on the real-time performance of the system.

2.3 RAM

A constant high memory bandwidth is also beneficial for an RTOS. Power management features can negatively impact the memory performance.

The following setting can usually be found in the BIOS setup program under submenu Chipset/Processor (Integrated components)/Memory Configuration.

2.3.1 System Agent Geyserville (SA GV)

If enabled, the Intel Memory Controller (IMC) of Intel Core U processors changes the memory clock speed between two operating points depending on memory utilization.

Note

Dual-channel (interleaved) mode improves real-time behavior. Therefore, congatec highly recommends a dual-channel configuration with two identical DRAM modules.

2.4 Integrated Graphics Device (IGD)

In case the real-time OS uses the integrated GPU for video output it is also important to disable certain power management functions for the graphics.

The following settings can usually be found in the BIOS setup program under submenu Chipset/Processor (Integrated components)/Graphics Configuration and/or Advanced/Power&Performance/GT Power Mangement Control.

2.4.1 GT PM Support

If disabled, power saving features of the IGD are disabled resulting in best response time.

2.4.2 Render Standby (RC6)

If enabled, the IGD core voltage is significantly lowered to reduce power consumption. If GT PM Support is disabled, this feature is automatically disabled.

2.4.3 Maximum GT Frequency

If enabled, the IGD frequency is limited. Since the Thermal Design Power (TDP) budget is shared with the x86 cores, enabling this feature can improve the performance stability of the x86 cores at the cost of IGD performance. Whether this trade-off is acceptable or not, depends on the real-time application. For more information, see section 3.2 "Thermal Design Power (TDP) Constraints".

2.5 LPM

Link Power Management (LPM) provides power saving features for PCI Express (PCIe), SATA, and the Direct Media Interface (DMI) links. Such features should be disabled because they will introduce latency.

Note

The BIOS setup program might provide power-saving features that are not described in the sub-sections below but should also be disabled. For example, clock gating (e.g. for PCIe) and throttling (e.g. PCH Throttling, Thermal Throttling) should be disabled.

2.5.1 DMI Link ASPM Control

If enabled, the Direct Media Interface (DMI) link between the CPU and the PCH can go into low power states. The transition between power states increases latency.

2.5.2 PCI Express Link ASPM Control

If enabled, the PCI Express link can go into active low power states. The transition between these power states increases latency. Each PCIe port (including PEG) has a dedicated BIOS setup option to enable ASPM.

2.5.3 Aggressive LPM Support

If enabled, the SATA controller can go into low power states. The transition between power states can result in delays when reading from or writing to the SATA device.

2.5.4 Legacy IO Low Latency

This setup option disables additional clock gating and power management functions and should always be Enabled if it is available in the BIOS setup.

3 Additional Advice

In addition to the BIOS setup settings described in the previous section, other (not BIOS related) optimizations may be required to further improve the performance of an RTOS.

Note

It is outside the scope of this application note to cover all the topics that might be relevant for an RTOS. Contact your RTOS or Hypervisor vendor for further support.

3.1 Thermal Design

CPU thermal throttling results in unstable CPU performance and can lead to total system failure. The thermal design must ensure adequate cooling of the CPU at all times to prevent the Thermal Control Circuit (TCC) in the CPU from triggering thermal throttling. Furthermore, the thermal design must also ensure adequate cooling of other components (e.g. memory) at all times to prevent unstable performance.

Most Intel CPUs are rated for a maximum die temperature of 100°C. The TCC triggers CPU thermal throttling at this temperature. This temperature can be lowered (but not raised) via a BIOS setup setting. For more information, refer to section 1.2.5 “Thermal Control Circuit (TCC) Activation Offset”.

3.2 Thermal Design Power (TDP) Constraints

All circuits integrated in the processor package share the TDP budget with the x86 CPU cores. For example, most recent mobile processors made by Intel have an Integrated Graphics Device (IGD) and Integrated Memory Controller (IMC). The TDP is usually also the power limit of all the integrated circuits in the processor package.

Especially the IGD can draw a lot of power under load. If the TDP budget is reached, the CPU performance becomes unstable. Most importantly, the x86 cores may run at a much lower clock speed.

There are several ways to prevent such unstable CPU performance:

- Reduce the IGD power consumption (see section 2.3.3 “Maximum GT Frequency”)
- Increase the power limit (cTDP – Only available on select Intel processors)
- Choose a different processor with a sufficiently high TDP

3.3 QoS Support in Silicon and RTOS/Hypervisor

To run hard real-time applications on top of a hypervisor, ensure the following requirements are met by the hypervisor:

- Strict temporal isolation of guest operating systems
- Use of QoS features found in the silicon to eliminate interference on hardware level (e.g. Cache Allocation Technology)
- Use of real-time execution mode without adding latencies and/or jitter while accessing hardware or processing interrupts

For optimal real-time performance (lowest latencies and jitter), congatec recommends to use the RTS Hypervisor.