# **Engineer-to-Engineer Note**

ANALOG Technical notes on using Analog Devices products and development tools Visit our Web resources http://www.analcom/ee-notes and http://www.analog.com/processors or e-mail processor.support@analog.com or processor.tools.support@analog.com for technical support.

# Using ADSP-SC59x/2159x High Performance FIR/IIR Accelerators

Contributed by Sanket Nayak and Mitesh Moonat

*Rev 2 – September 23, 2022* 

# Introduction

ADSP-SC59x/2159x processors (referred in this EE Note as ADSP-SC59x processors) incorporate highperformance hardware accelerators dedicated to performing two widely used signal processing operations, Finite Impulse Response (FIR) and Infinite Impulse response (IIR). These accelerators are hereby referred to FIRA and IIRA respectively.

Figure 1 shows an ADSP-SC59x processor that can have one or two SHARC+ cores, each equipped with one FIR and four IIR accelerators. The SHARC+ cores and accelerators together can provide an overall performance of up to 20 Giga Floating Point MAC (Multiply and Accumulate) Operations Per Second when running at 1 GHz.





<u>Figure 2 Processing Capability with and without FIR/IIR Accelerators</u> on page 2, shows these accelerators can run in tandem with the core performing the dedicated fixed-function computations (FIR and IIR). This increases the overall processing capability of the processor.

This application note provides a brief overview of the functional and performance improvements achieved with the ADSP-SC59x FIR and IIR accelerators. For this EE Note, FIRs and IIRs are referred to as *accelerators*, when compared with their predecessors.

Copyright 2023, Analog Devices, Inc. All rights reserved. Analog Devices assumes no responsibility for customer product design or the use or application of customers' products or for any infringements of patents or rights of others which may result from Analog Devices assistance. All trademarks and logos are property of their respective holders. Information furnished by Analog Devices applications and development tools engineers is believed to be accurate and reliable, however no responsibility is assumed by Analog Devices regarding technical accuracy and topicality of the content provided in Analog Devices Engineer-to-Engineer Notes.





Figure 2 Processing Capability with and without FIR/IIR Accelerators

This Engineer-to-Engineer Note discusses the following topics:

- Accelerator performance under different configurations and contributing factors
- Accelerator drivers from the CrossCore® Embedded Studio (CCES) software, that include usage directions, examples, and benchmark details
- Accelerator usage models for optimum performance in different application scenarios

For more details about the accelerator architecture and programming model, refer to the *ADSP-SC59x* SHARC+ Processor Hardware Reference <sup>[1]</sup>.



A majority of the content of this application note is the same as in *Using ADSP-2156x High Performance FIR/IIR Accelerators* (EE-408)<sup>[5]</sup> EE Note. This application note includes the feature enhancements, figures, tables, and data specific to the ADSP-SC59x/2159x family of processors.

The content of this application note supports both ADSP-SC594 and ADSP-SC598 family of processors. In the following sections, ADSP-SC59x/2159x refers to the *ADSP-SC594 family*. The details applicable to the ADSP-SC598 family are explicitly mentioned.

# ADSP-SC59x Accelerator Enhancements

Accelerator enhancements include both performance and functional improvements.

## **Performance Improvements**

More FIRA and IIRA instances exist in the ADSP-SC59x processors, compared to previous processors. Each SC59x SHARC+ core has one FIRA and four IIRA instances, all running at CCLK frequency. Both SHARC+ cores can access all the FIRA and IIRA instances, and all the instances can run in parallel. Eight IIRA instances can support a total of 2304 biquads <sup>[6]</sup> (288\*8) and 256 channels (32\*8). The ADSP-2156x processors can only support 288 biquads and 32 channels.





When a SHARC+ core accesses the accelerator instance associated with another SHARC+ core, there is additional access latency, compared to accessing its own accelerator instance.

- Finite Impulse Response Accelerator (FIRA) Direct Memory Access (DMA) overhead optimization.
  - Maximum Burst Transfer size is increased from INCR8 to INCR16 to reduce the Transfer Control Block (TCB) load time. This feature can be disabled by setting FIRCTL1.BURST16\_DIS bit. However, a burst length of 16 requires the end address of the Input Buffers must be 4-KB aligned. (Refer to Anomaly 20000114 in the ADSP-21591-21593-21594-SC591-SC592-SC594 Anomaly List<sup>[4]</sup>.)
  - The Data and Coefficient load now happens in parallel as compared to sequential loading in previous processors thereby reducing the DMA load time. This feature can be disabled by setting the FIR CTL1.DCP DIS bit.
  - The use of Pre-Fetch Buffer reduces the access latency in first data arrival of a Burst Transfer. This feature can be set by the FIR\_CTL1.PFB\_EN bit. Two target SHARC+ Core ports (S1 and S2) can be used to balance the DMA load. One constraint is that this feature can only be used when the Output Buffer and Input/Coefficient Buffers of the respective channels are separated by a minimum size of 12 bytes. (Refer to Anomaly 20000118 in *ADSP-21591-21593-21594-SC591-SC592-SC594 Anomaly List*<sup>[4]</sup>.)

Pre-Fetch Buffer can be enabled only when the Maximum Burst Transfer size of 16 is enabled in the ADSP-SC59x/2159x.

In the ADSP-SC598, the Pre-Fetch Buffer and the Maximum Burst Transfer size of 16 can be enabled or disabled, independent of each other.

# **Functional Improvements**

The functional improvements of the ADSP-SC59x accelerator include:

- Legacy Mode Index Update Scheme in Auto Configuration Mode-In Legacy Mode, after processing a window of data, the FIRA and IIRA updates the TCB Input and Output Buffer Index using a Circular Buffer Scheme. This feature can be enabled in the Auto Configuration Mode (ACM) by setting the FIR\_CTL1.SMQ\_LIUPS\_EN bit. When disabled, the Input Buffer Index updates to 0x0000000 and the Output Buffer Index updates to 0xFFFFFFFF.
- *IIRA Selective Interrupt Generation in Legacy Mode*—In Legacy Mode, an interrupt can be generated either after processing all TCBs or after processing each TCB, by the CTL1.CCINTR bit. Interrupts can be customized in the ACM on a per-channel basis by setting the IIR CTL1.L TIMASK EN bit.
- IIRA Trigger Generation and Trigger Wait in Legacy Mode-In ACM, each channel can generate a controller trigger after its completion. This is then routed to a trigger target by the Trigger Routing Unit (TRU). Each channel can be configured to wait for a controller trigger after loading the TCB. The above trigger mechanism is now available in Legacy Mode and can be enabled using the IIR\_CTL1.L\_TIMASK\_EN and IIR\_CTL1.L\_TWAIT\_EN bits.



## **Accelerator Performance**

Accelerator processing cycles consist of two components, DMA cycles, plus Multiply and Compute (MAC) cycles.

# **DMA Cycles**

Two factors significantly impact the DMA cycles, Filter Parameters and Buffer Placement.

*Filter Parameters*–DMA cycles consist of two operations:

- Initial preload of the coefficients and delay line (FIRA) or state variables (IIRA)-for IIRA, the coefficient load and state variable load can be skipped in Legacy Mode to save these cycles. The initial preload data size depends on the **Tap Length** (FIRA) or the **Number of Biquads** (IIRA).
- Fetch and Store of the Input/Output buffers-these operations are influenced by **Window Size** and take place in parallel with the next compute operation. As such, they do not add any overhead for a larger tap length or number of biquads. For a smaller tap length and biquad values, these cycles *can dominate* the total processing cycles. The crossover point (N) in terms of tap length (FIRA) and biquads (IIRA) depends on the memory access latency. The value of the crossover point increases as buffers are moved from L1 to L2 memory, and further to L3 memory.

*Buffer Placement*–The DMA cycles depend on the location of the input, output, and coefficient buffers in DSP memory, as each memory type, L1 or L2, or L3, has different access latencies.



For optimal performance in FIRA, Analog Devices recommends that the Input/Coefficient Buffers are 4-KB aligned and the TCB is 64-Byte aligned.

# **Compute (MAC) Cycles**

Unlike previous SHARC processors, the MAC units in the ADSP-SC59x/ADSP-2156x accelerators are derived from the SHARC+ core MAC units, where a multiplication operation takes place in two stages. This design implies that a minimum of N+1 number of cycles are required to perform N MAC operations. For the FIRA, the tap length size is generally significantly higher; consequently, the theoretical number of cycles per sample per tap is approximately 0.25 (1/4). For the IIRA, processing each biquad requires 6 cycles (for 5 MAC operations).

In an ideal scenario, to make maximum use of the accelerator computing power, the percent contribution of the compute and MAC cycles, called the **Compute Efficiency (CE)**, must be 100%.

CE can be measured using the expression:

CE = (M/N)\*100 where, M = Theoretical Cycles per Tap (FIRA = 0.25) or per biquad (IIRA = 6) N = Measured Cycles per Tap (FIRA) or per biquad (IIRA)



## **FIR/IIR Accelerators Performance**

This section illustrates FIR and IIR accelerator performance with different filter parameters and in different processors:

- ADSP-SC59x FIRA/IIRA at maximum ACLK (=CCLK) = 1 GHz
- ADSP-2156x FIRA/IIRA at maximum ACLK (=CCLK) = 1 GHz
- ADSP--SC57x FIRA/IIRA at maximum ACLK (=SCLK0) = 125 MHz and at maximum CCLK = 500 MHz
- ADSP-214xx FIRA/IIRA at maximum ACLK(=PCLK=CCLK/2) = 225 MHz and at maximum CCLK = 450 MHz
- ADSP-2156x (SHARC+) core at maximum CCLK = 1 GHz



The ADSP-SC58x accelerator performance is the same as the ADSP-SC57x accelerator, except for a slight performance loss due to the lack of a DMA burst mode.

All the buffers are placed in L1 memory during the measurements. Core FIR and IIR cycles are measured by the CCES library functions.

Additional measurements and analysis discuss how the various factors affect the accelerator MAC Utilization Efficiency (MUE).

#### FIR Performance

This section contains the following information:

- Measured performance of the ADSP-SC59x FIRA for different filter parameters and comparison with other cases
- Measured FIRA Compute Efficiency (CE) and contributing factors
- The effect of buffer placement (L1, L2, and L3 memory) on FIRA performance

#### Comparison with ADSP-2156x/SC57x/214xx Accelerators and SHARC+(measured on ADSP-2156x) Core

Figure 3 FIR Performance in Core Cycles Across Window Size and Tap Length on page 6 and Figure 4 FIR Performance in Time (µs) Across Window Size and Tap Length on page 7 show the measured FIR performance in core cycles and time units (µs) for the six cases. Figure 5 ADSP-SC59x FIR Performance Improvement Factor for 512 Taps on page 8 shows the performance improvement factor of one ADSP-SC598 FIR accelerator instance at different window sizes and a tap length of 512, in comparison with other cases.



The following figures (Figure 3, Figure 4, and Figure 5 on the next 3 pages) compare the performance of a single accelerator instance on different processors.





#### Figure 3 FIR Performance in Core Cycles Across Window Size and Tap Length















#### Figure 4 FIR Performance in Time (µs) Across Window Size and Tap Length















Figure 5 ADSP-SC59x FIR Performance Improvement Factor for 512 Taps for Different Window Sizes

The following observations are derived from Figure 5:

- The ADSP-SC598 FIRA outperforms the ADSP-SC59x/2159x for small window sizes. The improvement primarily results from the reduced ADSP-SC598 FIRA DMA overhead. The DMA load time accounts for a large portion of the total processing time in the small window cases, whereas the DMA load time is insignificant in the large window cases.
- The ADSP-SC598 FIRA outperforms the ADSP-2156x for small window sizes. The improvement primarily results from the reduced ADSP-SC598 FIRA DMA overhead. The DMA load time accounts for a large portion of the total processing time in the small window cases whereas the DMA load time is insignificant in the large window cases.
- The ADSP-SC598 FIRA performance is approximately 8 times that of ADSP-SC57x FIRA. The improvement is because the ADSP-SC598 FIRA runs at CCLK (1 GHz), whereas the ADSP-SC57x FIRA runs at SCLK, which is a maximum of CCLK/4 (125 MHz). For small window sizes, the performance factor increases due to the enhancements to reduce the DMA load time on ADSP-SC598 FIRA.
- The ADSP-SC598 FIRA performance is approximately 4.46 times that of the ADSP-214xx FIRA. The improvement is because the ADSP-SC598 FIRA runs at CCLK (1 GHz), whereas the ADSP-214xx FIRA runs at PCLK, which is a maximum of CCLK/2 (225 MHz). For small window sizes, the performance factor increases because of the enhancements that reduce the DMA load time for the ADSP-SC598 FIRA.
- The ADSP-SC598 FIRA performance is approximately two times that of the ADSP-2156x SHARC+ core's FIR code execution. The improvement is because the ADSP-SC598 FIRA has four



MAC units, whereas the ADSP-2156x SHARC+ core has two MAC units. An exception to this conclusion applies to smaller block sizes, where the DMA overhead dominates the compute cycles.

## Compute Efficiency (CE)

Figure 6 shows how the Compute Efficiency (CE) percentage for the FIRA varies with window size and tap length. The CE approximates 100% for large (>128) window sizes. The CE significantly decreases for smaller window sizes. Compared to previous processors, the CE for smaller window sizes are significantly better because of the DMA load time improvement in ADSP-SC59x FIRA. For a constant window size, the CE percentage does not vary significantly for different tap lengths.



Figure 6 FIRA Compute Efficiency (CE) for Different Tap Lengths and Window Sizes

## **Buffer Placement**

Figure 7 ADSP-SC59x FIRA Performance Comparison for Buffers in L1, L2, and L3 Memory on page 10 shows how the core cycles change with increased tap lengths when the input/output/coefficient buffers are placed in L1/L2/L3 memory for a window size of 1024. It also shows a plot of the performance factor when the buffers are in L1 memory, versus in L2 and L3 memory.





Figure 7 ADSP-SC59x FIRA Performance Comparison for Buffers in L1, L2, and L3 Memory

The following observations are derived from the Figure 7 performance factor graphs:

- As the tap length increases for a constant window size, the core cycles remain constant until a certain point (N), then increase proportionally. For small tap lengths, the input/output DMA cycles dominate the compute cycles.
- As the memory access latency increases from L1 to L2 memory and then further to L3 memory, the value of certain point (N) increases, and the performance factor approaches 1.0. Placing buffers in L2 and L3 memory instead of L1 memory might be expensive for tap lengths less than certain point (N) and comparable for tap lengths greater than or equal to certain point(N).

## **IIR Performance**

This section contains the following topics:

- Measured performance of the ADSP-SC59x IIRA for different filter parameters and comparison with other cases
- Measured IIRA Compute Efficiency (CE) and contributing factors
- The effect of buffer placement (L1, L2, and L3 memory) on IIRA performance



#### Comparison with ADSP-2156x/SC57x/214xx Accelerators and SHARC+ (measured on ADSP-2156x) Core

<u>Figure 8</u> and <u>Figure 9</u> show the measured IIR performance of one IIRA instance in core cycles and time units ( $\mu$ s), respectively, for all the cases. <u>Figure 10 ADSP-SC59x IIR Performance Improvement Factor</u> for 6 Biquads Considering Time on page 13 shows the performance factor of the ADSP-SC9x IIRA, as compared to the other cases for different window sizes and a constant value of six biquads.









IIR performance for window size = 64 14000 **CCLK Cycles** SC59x\_ACC\_W64 2156x\_ACC\_W64 SC57x\_ACC\_W64 214xx\_ACC\_W64 2156x Core ACC W64 Biguads







#### Figure 9 IIR Performance in Time (µs) Across Window Sizes and Number of Biquads

 $(\mathbf{i})$ 

The <u>Figure 9</u> values compare the performance of a single accelerator instance on different processors. The performance of ADSP-SC598 IIRA is similar to the ADSP-SC59x/2159x IIRA performance. Hence, the ADSP-SC59x/2159x IIRA values shown are also applicable to the ADSP-SC598 IIRA.





#### Figure 10 ADSP-SC59x IIR Performance Improvement Factor for 6 Biquads Considering Time

The following observations are derived from the IIR performance values shown in Figure 10:

- The ADSP-SC59x IIRA performance is similar to the ADSP-2156x IIRA performance. Both ADSP-SC59x IIRA and ADSP-2156x IIRA run at CCLK (1 GHz).
- The ADSP-SC59x IIRA performance is approximately 6.7 times the ADSP-SC57x IIRA performance. The ADSP-SC59x IIRA runs at CCLK (1 GHz), whereas the ADSP-SC57x IIRA runs at SCLK, which is a maximum of CCLK/4 (125 MHz). The reduction in the improvement factor from 8 to 6.7, compared to the FIRA, is because the ADSP-SC59x IIRA compute unit takes 6 cycles per biquad. The ADSP-SC57x IIRA takes 5 cycles per biquad.
- The ADSP-SC59x IIRA performance is approximately 3.7 times the ADSP-214xx IIRA performance. The ADSP-2156x IIRA runs at CCLK (1 GHz), whereas the ADSP-214xx IIRA runs at PCLK, which is a maximum of CCLK/2 (225 MHz). The reduction in the improvement factor from 4.44 to 3.7, compared to the FIRA, is because the ADSP-2156x IIRA compute unit takes 6 cycles per biquad. The ADSP-214xx IIRA takes 5 cycles per biquad.
- The ADSP-SC59x IIRA performance is approximately 0.42 times the ADSP-2156x SHARC+ core's execution of IIR code. Each ADSP-SC59x IIRA instance has one MAC unit, whereas each SHARC+ core has two MAC units. Additionally, the IIRA compute unit takes 6 cycles per biquad. The SHARC+ core takes 2.5 cycles.



## Compute Efficiency

Figure 11 shows how the Compute Efficiency (CE) percentage for IIRA varies with the window size and number of biquads. As can be seen in the plot, the CE approximates 100% for larger window sizes, and it significantly reduces for small window sizes. For a constant window size, the CE percentage does not vary significantly with the number of biquads, except when the biquad count is one or two. This is because the input and output DMA cycles dominate the total processing cycles for a low biquad count.





## **Buffer Placement**

Figure 12 ADSP-SC59x IIRA Performance Comparison for Buffers in L1, L2, and L3 Memory on page 15 shows how the performance measured in terms of core cycles change when the input/output/coefficient buffers are placed in L1/L2/L3 memory with a window size equal to 1024 and the number of biquads is increased. It also shows a plot of the performance factor when the buffers are in L1 memory, compared to when they are in L2 or L3 memory.

The following observations are derived from the Figure 12 CE graphs:

- As the number of biquads increases, the core cycles do not increase significantly until a point (N), at which the core cycles start increasing proportionally. This change happens because the input and output DMA cycles dominate the compute cycles when the biquad count is low.
- As the memory access latency increases from L1 to L2 memory, and then further to L3 memory, the N value increases, and the performance factor approximates 1.0, as the number of biquads becomes large. Placing buffers in L2 and L3 memory, instead of L1 memory, might be expensive for biquad counts less than N; and comparable for biquad counts greater than or equal to N.





Figure 12 ADSP-SC59x IIRA Performance Comparison for Buffers in L1, L2, and L3 Memory

# **Programming the Accelerators**

The ADSP-SC59x accelerators can be programmed using the device drivers available with the CrossCore® Embedded Studio (CCES) software.





Figure 13 ADSP-SC59x Accelerator Device Driver Structure

Each ADSP SC59x processor has multiple instances of a FIRA(2) and IIRA(8). The SHARC+ cores can access all FIRA/IIRA instances. To minimize the software complexity of accessing those accelerator instances in parallel by the SHARC+ cores, a static configuration scheme can be employed. In this scheme, all accelerator instances are assigned to one SHARC+ core or divided into two groups for the cores. With this organization, both SHARC+ cores can operate independently in parallel using the assigned accelerator instances for processing.

The programming model is the same for the FIR and IIR accelerators. <u>Figure 13</u> shows the accelerator drivers adhere to a queueing programming model. The model uses abstractions of task and queues. A *task* contains a set of *channels* that are processed sequentially when given to the accelerator, A *queue* is a set of tasks the accelerator processes sequentially and maintains on a first-in, first-processed basis. Applications can create tasks using the adi\_fir(/iir)\_CreateTask() call and queue those tasks using the adi\_fir(/iir)\_QueueTask() call.

The driver maintains a processing queue of the tasks given or queued by the application for the accelerator and sequentially schedules the tasks to the accelerator. The applications can register a user-defined callback with a adi\_fir(/iir)\_RegisterCallback() call. The driver uses the callback to notify the application when the tasks and channels complete their work.

The driver maintains the status of all the application tasks it has queued. The application can query the status of the task (when necessary) using the adi\_fir(/iir)\_GetFir(/Iir)TaskStatus() call. Once the accelerator completes the queued task, it is removed from the processing queue. The driver maintains all task information for the tasks which are completed or created. The application can later reuse the created tasks by re-queuing them with the adi\_fir(/iir)\_QueueTask() call.



The accelerator can be configured using a static configuration to operate in legacy mode or Auto Configuration Mode (ACM). In legacy mode, the driver maintains a software queue (linked list) of the tasks queued by the application and sequentially schedules the tasks to the accelerator. In ACM, the driver software uses the hardware halt feature to implement a hardware queuing mechanism. Additional channel-customized features exist in the ACM, such as interrupts and triggers. For more feature details, see the ADI\_FIR(/IIR)\_CHANNEL\_INFO topic in the ADSP-SC59x SSDD API Reference Manual for SHARC+ Core <sup>[2]</sup> of the CCES. The driver programming model is the same for legacy and ACM operation.

The FIR\_Multi\_Channel\_Processing and IIR\_Multi\_Channel\_Processing source code examples provided with this application note <sup>[3]</sup> (also available with the ADSP-SC594 EZ-Kit *Board Support Package*), can be used as a reference for the ADSP-SC59x accelerator device drivers.

<u>Table 1</u> summarizes the accelerator driver benchmarks captured using the ADSP-SC594 EZ-Kit evaluation board. The source code examples used to obtain these performance measurements, ADSP\_SC59x\_FIRA\_Driver\_Benchmark and ADSP\_SC59x\_IIRA\_Driver\_Benchmark, can be found in the ZIP file associated with the EE436 application note <sup>[3]</sup>. The measured values obtained were done with all buffers in L1 memory.

Operation	Description	Measured Core CyclesFIRAIIRA			
				A	
		Legacy	ACM	Legacy	ACM
CreateTask()	Constant Overhead	89	149	151	125
	Per Channel Overhead	335	347	332	353
QueueTask()	Pushing a task to the empty queue	350	536	355	537
	Pushing a task to the non-empty queue	194	445	189	436
Interrupt Overhead	For a single task in queue	495	456	500	458
(Round trip cycles include the SEC interrupt dispatcher latency.)	For more than one task in queue	623	646	645	622

Table 1: ADSP-SC59x Accelerator Driver Benchmarks



Additional overhead for cache flushing and invalidation can occur when the buffers are placed in L2 and L3 memories. Analog Devices recommends placing the buffers in L1 memory to avoid cache overhead. The <u>Table 1</u> numbers are also applicable to the ADSP-SC598 FIRA/IIRA accelerators.



# **Accelerator Usage Models**

This section describes models that optimally use the accelerator for different application scenarios. The optimum solution is to offload all FIR and IIR tasks from the SHARC+ cores to the accelerators and allow the cores to do other parallel processing. This solution may not be feasible in some scenarios, particularly when the core needs to use the output from the accelerator for additional processing and has no other independent tasks to finish in parallel. The EE-436 application note discusses the four accelerator usage models (Core Only, Direct Replacement, Split Task, and Data Pipelining) shown in Figure 14.



#### Figure 14 Accelerator Usage Models

## **Direct Replacement**

Direct Replacement is the most straightforward and easy-to-use method. The core FIR and IIR processing is directly replaced by the accelerator. The core simply waits for the accelerator to finish the job. This model is useful only when the FIRA is employed, because it is faster than the core.

## Split Task

Split Task divides the overall FIR and IIR processing between the core and the accelerator. This model is useful when there are multiple channels available for parallel processing. Based on a timing estimation, Split Task divides the total number of channels between the core and the accelerator so that both finish processing at approximately the same time. In some scenarios, this approach yields better results than the Direct Replacement approach. The Split Task model uses the core idle time to finish the overall FIR and IIR processing tasks faster than the accelerator working alone.



## **Data Pipelining**

In this model, data flow between the core and accelerator can be pipelined in such a way that both can work in parallel on different data frames. As shown in <u>Figure 14 Accelerator Usage Models</u> on page 18, the core processes the Nth frame and then initiates the accelerator's processing of this frame. The core continues in parallel to further process the (N-1)<sup>th</sup> frame output produced by the accelerator in the previous frame. This sequence provides for the complete offloading of the FIR/IIR processing task to the accelerator, at the cost of additional output latency. The pipeline stages and consequently the output latency, can increase depending upon the number of such FIR/IIR processing stages in the processing chain.

Table 2 compares the model approaches and comments on the performance gains and constraints.

Approach	Description	Performance Advantage	Performance Constraints
Direct Replacement	Core processing is replaced by the accelerator. Core waits until the accelerator completes the job.	+	Suitable only for FIRA, as it is faster than the core.
Split Task	FIR/IIR processing of multiple channels is split between the core and the accelerator.	++	Cannot be used for single-channel processing.
Data Pipelining	Data between the core and accelerator is pipelined to allow the core and accelerator to run in parallel	+++	Costs additional latency of at least one frame.

Table 2: Accelerator Usage Approach Comparison

# **Real-Time Implementation of an Example FIR/IIR Use Case**

<u>Figure 15 FIR/IIR Real-Time Use Case</u> on page 20 provides an example of FIR/IIR real-time processing of 12-channel audio data.

- Each channel passes through a 1024-tap FIR filter, followed by a five-band IIR equalizer.
- For simplification, the FIR filter is an all-pass filter (the b0 coefficient is one and all other coefficients are zero).
- Each IIR band is an 8<sup>th</sup> order IIR filter, implemented with four-cascaded biquad stages. <u>Table 3</u>: <u>IIR Band Specifications</u> on page 20 describes the IIR band details. The output of each band is multiplied with the respective gain (from 0.0 minimum through 1.0 maximum) and added together to generate the final output. The *IIRGains* flag can be set (bypass) or cleared (enable processing) dynamically using the CCES memory browser.





Figure 15 FIR/IIR Real-Time Use Case

Table 3: IIR Band Specifications

IIR Band Number	Lower Cutoff in Hz	Higher Cutoff in Hz	Filter Type
1	-	250	Low Pass
2	250	500	Band Pass
3	500	2,000	Band Pass
4	2,000	4,000	Band Pass
5	4,000	-	High Pass

- The block size is 256 samples per channel.
- All buffers are placed in L1 memory.

The Direct\_Replacement, Pipelined, and Split\_Task example code is provided with the application note EE436 ZIP file<sup>[3]</sup>, that implements these models on the ADSP-21593 EZ-Kit evaluation board.

A similar set of code examples is provided for the ADSP-21569 EZ-Kit evaluation board. The core and accelerator performance numbers are measured for the *core only* case and for the single accelerator usage model. The resultant core MIPS savings for each model are displayed. <u>Table 4 MIPS Summary for Different Accelerator Usage Models</u> on page 21 shows the MIPS of the core and accelerator for different usage models, plus the Core MIPS savings. *The numbers measured on the ADSP-21569 are mentioned in parentheses for comparison*.

From a performance perspective, the *Data Pipelining* model is the best, resulting in saving approximately 333 core MIPS. Next best is the *Split Task* model (saving 252 core MIPS), and followed the *Direct Replacement* model (saving 239 core MIPS). These results align with the previous discussion about the accelerator usage models.



	MIPS Usage			
Usage Model	Core	FIRA	IIRA	<b>Core MIPS Saving</b>
Core Only	339 (339)	0.00		-
Direct Replacement	100 (241)	82 (161)	12 (75)	239 (98)
Split Task	87 (165)	69 (110)	12 (50)	252 (174)
Data Pipelining	6 (5)	82 (161)	12 (75)	333 (334)

Table 4 MIPS Summary for Different Accelerator Usage Models



The numbers were measured on ADSP-SC59x/2159x and Pre-Fetch Buffer and Maximum Burst Transfer size of 16 were enabled for FIRA.

# Key Take-aways

The following observations derive from the comparison of <u>Table 4</u> MIPS usage for the ADSP-SC59x/2159x processors and the ADSP-2156x:

- The core MIPS saving achieved on the ADSP-SC59x/2159x is significantly better than that on ADSP-2156x. This is because multiple instances of FIRA / IIRA(2x/8x) operate in parallel on the ADSP-SC59x/2159x, compared to a single instance of FIRA/IIRA on ADSP-2156x.
- As we move from the Core Only to Direct Replacement model, we observe a significant amount of core MIPS saving on the ADSP-SC59x/2159x. This means for the customer that with a minimal amount of code changes, that is, replacing the *Core* function calls with *Accelerator* function calls, a significant amount of core MIPS can be saved.



# References

- [1] ADSP-SC59x SHARC+ Processor Hardware Reference, Rev 0.1, October 2018. Analog Devices, Inc.
- [2] ADSP-SC59x SSDD API Reference Manual for SHARC+ Core, Version 2.0, CrossCore® Embedded Studio Help.
- [3] Associated ZIP file for EE-436: Using ADSP-SC59x/2159x High Performance FIR/IIR Accelerators, March 2022. Analog Devices, Inc.
- [4] ADSP-21591-21593-21594-SC591-SC592-SC594 Anomaly List (Rev. D)
- [5] EE-408: Using ADSP-2156x High Performance FIR/IIR Accelerators, August 2019. Analog Devices, Inc.
- [6] **Biquad**—any block that realizes a biquadratic transfer function (that is, the numerator and denominator are quadratic polynomials in s [continuous time] or in z to the -1 [discrete time]). The Biquad Filter block independently filters each channel of the input signal with the specified biquadratic infinite impulse response (IIR) filter.

# **Document History**

Revision	Description
Rev 2– September 23, 2022 by Sanket Nayak and Mitesh Moonat	Added changes to include the details and graphs for ADSP-SC598.
Rev 1– June 14, 2022 by Sanket Nayak and Mitesh Moonat	Initial Release