

SigStream Operating Guide

The logo for Red Rapids, featuring the text "Red Rapids" in white, bold, sans-serif font inside a black rounded rectangle with a red border.

797 North Grove Rd, Suite 101
Richardson, TX 75081
Phone: (972) 671-9570
www.redrapids.com

Red Rapids reserves the right to alter product specifications or discontinue any product without notice. All products are sold subject to the terms and conditions of sale supplied at the time of order acknowledgment. This product is not designed, authorized, or warranted for use in a life-support system or other critical application.

All trademark and registered trademarks are the property of their respective owners.

Copyright © 2020, Red Rapids, Inc. All rights reserved.

Table of Contents

1.0	Introduction.....	1
1.1	Conventions.....	1
1.2	Revision History.....	2
2.0	Functional Description.....	3
2.1	Trigger Select.....	3
2.2	SigStream Receiver Datapath.....	4
2.2.1	Input Select.....	4
2.2.2	Real or Complex Representation (Model Specific).....	5
2.2.3	Resampler.....	5
2.2.1	DDC-1 Digital Down Converter (Model Specific).....	5
2.2.2	DDC-2 Digital Down Converter (Model Specific).....	6
2.2.1	Output Formatter.....	7
2.3	SigStream Transmitter Datapath.....	7
2.3.1	Output Select.....	7
2.3.2	Sample Probe.....	8
2.4	Operating Modes.....	8
2.4.1	Continuous Mode.....	8
2.4.2	Snapshot Mode.....	8
2.4.3	Periodic Mode.....	8
3.0	VITA Radio Transport Standard (ANSI/VITA 49.0).....	10
3.1	Header.....	10
3.2	Stream and Class Identifier.....	11
3.3	Timestamps.....	11
3.4	Data Payload.....	12
3.5	Trailer.....	13
4.0	Hardware Configuration.....	15
4.1	Time of Day (TOD) Clock Settings.....	15
4.2	LED Source Selection.....	16
4.3	GPIO Settings.....	17
4.4	Device Configuration Settings.....	17
5.0	Operating Configuration.....	20
5.1	Channel Settings.....	20
5.2	Switch Settings.....	20
5.3	Payload Format Settings.....	23
5.4	Datapath Control Settings.....	24
5.5	Synchronizer.....	28
5.6	Time of Day Triggers.....	30
5.7	Event Duration.....	30
5.8	IF Data Packet Format.....	30
5.9	DMA Buffer.....	32

List of Figures

Figure 2-1 SigStream Datapath and Control Framework	3
Figure 2-2 Receiver Datapath.....	4
Figure 2-3 DDC-1 Digital Down Converter.....	5
Figure 2-4 DDC-2 Digital Down Converter.....	6
Figure 2-5 Transmitter Datapath.....	7
Figure 3-1 VRT (VITA 49.0) IF Data Packet Format	10
Figure 3-2 Data Packet Header Format.....	10
Figure 3-3 Class Identifier Contents	11
Figure 3-4 Item Packing Field.....	12
Figure 3-5 Processing Efficient Payload Organization	12
Figure 3-6 Link Efficient Payload Organization	13
Figure 3-7 Trailer.....	13
Figure 4-1 Analog Devices AD9512 Configuration Registers.....	18
Figure 5-1 RX Switch Select Settings	21
Figure 5-2 TX Switch Select Settings	22
Figure 5-3 TX/RX Data Flow.....	26

List of Tables

Table 3-1 TSI/TSF Code Meaning.....	11
Table 3-2 State and Event Indicators.....	13
Table 4-1 LED A Trigger Monitor Assignment	16
Table 4-2 LED B Channel Monitor Assignment.....	17
Table 4-3 GPIO Setting Assignment.....	17
Table 5-1 Event Tag Mask Settings.....	24
Table 5-2 Transform Array Definition.....	25
Table 5-3 Residue Array Definition.....	26
Table 5-4 Residue Variable Settings	26
Table 5-5 SynclnitMask Array Definition	28
Table 5-6 StartEvent and StopEvent Variable Settings.....	29
Table 5-7 CycleSelect Variable Settings.....	29
Table 5-8 CycleStop Variable Settings	30

1.0 Introduction

The SigStream product family addresses a very broad range of applications. The software configurable datapath can be tailored to process a variety of signal types from radio waveforms to instrumentation.

The SigStream software application demonstrates a complete signal acquisition and generation thread. The behavior of the application is controlled through several structures that are grouped into two categories; hardware configuration and operating configuration.

The hardware configuration structures tend to be channel independent and frequently apply to specific chips that are not present on every product. The variables assigned to these structures are used to perform the following operations:

- Set the time of day (TOD) clock.
- Select the action that will illuminate status LEDs.
- Set the direction and logic state of GPIO or USER ports.
- Load internal registers of all chips that have a serial bus interface (SPI, I2C, etc).

The operating configuration structures target the signal conditioning and data manipulation features available to each channel. The variables assigned to these structures are used to perform the following operations:

- Select the source or sink attached to each datapath.
- Format the data payload.
- Define the behavior of the channel.
- Format the IF data packets.
- Set DMA operating parameters.
- Set interrupt processing options.

The SigStream application automatically scales to any number of transmitter or receiver channels supported by the target product. It obtains this product specific information by reading a feature identification register from the Red Rapids device.

The latest product documentation and software is available for download from the Red Rapids web site (www.redrapids.com).

1.1 Conventions

This manual uses the following conventions:

- Hexadecimal numbers are prefixed by “0x” (e.g. 0x00058C).
- *Italic* font is used for names of registers.
- **Blue** font is used for names of directories, files and OS commands.
- **Green** font is used to designate source code.
- Active low signals are followed by ‘#’, For example, TRST#.



Text in this format highlights useful or important information.



Text shown in this format is a warning. It describes a situation that could potentially damage your equipment. Please read each warning carefully.

The following are some of the acronyms used in this manual.

- **ADC** Analog to Digital Converter
- **DAC** Digital to Analog Converter
- **DDC** Digital Down Converter
- **F_s** Sample Clock Frequency
- **HBF** Halfband Filter
- **PPS** Pulse per Second
- **RX** Receiver
- **TOD** Time of Day
- **TSF** TimeStamp-Fractional (VITA 49)
- **TSI** TimeStamp-Integer (VITA 49)
- **TX** Transmitter
- **VITA** VME International Trade Association
- **VRT** VITA Radio Transport

1.2 Revision History

Version	Date	Description
R04	04/07/2020	Updated transmitter datapath description for the Model 272.
R03	07/20/2017	Added item and packet size limits imposed on the Model 276.
R02	03/24/2017	Updated DDC descriptions.
R01	08/04/2016	Added architecture description.
R00	01/12/2016	Initial release.

2.0 Functional Description

All SigStream products share a common datapath and control framework as shown in Figure 2-1. The hardware architecture is designed to accommodate up to eight receiver (ADC) or transmitter (DAC) channels in any combination. The number of channels varies by model number along with the sample rate and resolution of the ADC or DAC. Some SigStream products have a single SMA connector that functions as either the Coaxial Trigger or REF/CLK input. The selection is made as a build option when the product is ordered.

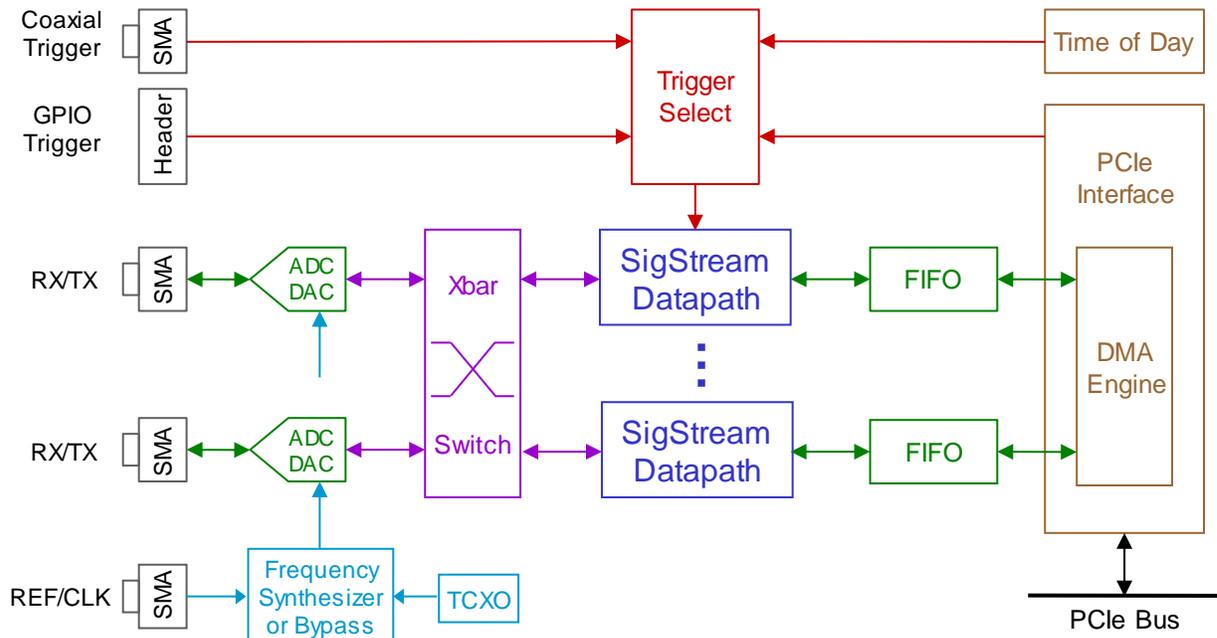


Figure 2-1 SigStream Datapath and Control Framework

The channel inputs are consecutively numbered from one through the total number of channels available. The numbering always begins with receiver channels on units configured as a transceiver. There is one SigStream datapath for every physical channel. However, the crossbar switch allows any physical channel to connect with any SigStream datapath. This makes it possible to connect multiple datapaths to a single receiver input if desired.

The remainder of this section will focus on the trigger select function, SigStream datapath, and operating modes. There are separate manuals that discuss the front-end (ADC/DAC) performance, clock distribution, and DMA engine.

2.1 Trigger Select

It is often necessary to synchronize signal acquisition or generation to an external event, or simply synchronize multiple channels within a system. The SigStream datapath can accept triggers from a variety of sources to initiate or terminate channel processing.

An external hardware trigger can be applied through either the coaxial connector or the GPIO connector. The GPIO port offers four high impedance input triggers and one 50 ohm terminated trigger. It is important to meet the setup and hold requirement of the trigger if accuracy to a single sample clock period is required. It may not be possible to achieve this accuracy across multiple cards if the sample clock period is lower than the propagation delay ambiguity through the trigger input buffer.

The on-board clock can be used to trigger a channel based on a specific time of day selected by software. The time of day trigger has a resolution of one sample clock period, but this does not necessarily translate to accuracy. The time of day clock is usually synchronized to an external source using a 1 PPS strobe on the coaxial trigger input. The seconds counter advances on the rising edge of this strobe and the fractional seconds counter is cleared. Just like an external trigger, this 1 PPS input will have to meet strict setup and hold requirements to obtain accuracy within a single period of the fractional seconds counter. The fractional seconds counter usually runs at the sample clock frequency, but there are physical limits that result in a lower frequency for some of the very high performance products.

Individual channels can also be triggered from application software by simply writing a register value. This technique incurs the greatest timing ambiguity because there is no way to predict the exact length of the PCIe bus arbitration process, but it can be useful in applications that do not require precision timing.

All of the triggers can be internally time delayed by a fixed number of fractional second clock periods. This alignment feature allows the channel to compensate for latencies through other parts of a system.

2.2 SigStream Receiver Datapath

The internal elements of a SigStream receiver datapath are shown in Figure 2-2. Digitized samples enter the datapath from the crossbar switch and exit into a FIFO.

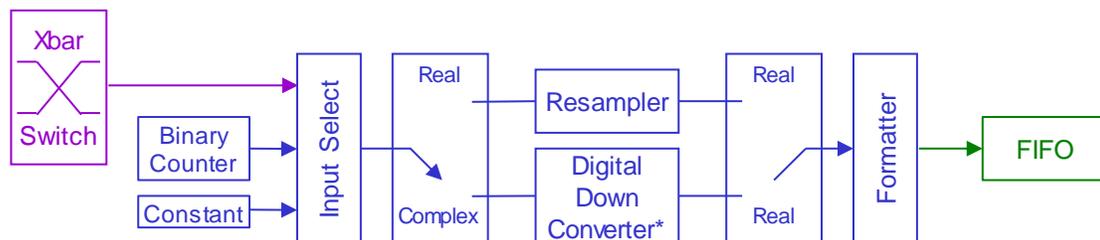


Figure 2-2 Receiver Datapath

*The digital down converter is unique to each SigStream product, refer to the Hardware Configuration Guide for details about a specific model number.

2.2.1 Input Select

In normal operation, SigStream input data is supplied by an ADC through the crossbar switch. There is an option to substitute ADC sample data with a binary count or a constant value to assist with debug of channel configuration settings.

The binary count can be easily parsed by diagnostic software to determine that the channel is behaving as expected. This data source can be used to verify resampler or snapshot settings in addition to proving that samples are not dropped or repeated anywhere in the processing chain. The digital down converter should not be used with this data source.

The constant input can be used to verify channel configurations involving the digital down converter. This is a handy diagnostic feature since the quadrature mixer can be used to generate a complex sinusoid from the constant. The remainder of the datapath will produce deterministic results that can be used to accurately analyze filter performance independent of ADC uncertainties.

2.2.2 Real or Complex Representation (Model Specific)

Application software can select either real or complex data representations on units equipped with a digital down converter. Complex samples will occupy the space of two real samples in the output payload, but will be spaced twice as far apart in time. The Nyquist criteria is not violated since the original signal samples can be reconstructed from the complex representation.

Raw ADC samples can bypass the digital down converter to retain their real representation. Although the data is routed through a resampler, there is no requirement to reduce the bandwidth. Setting the resampler to one will simply pass the raw samples to the output formatter.

The digital down converter generates real (in-phase) and imaginary (quadrature) components from the raw samples supplied by the ADC to produce a complex value. There is no limit imposed on the output bandwidth through this path. The full ADC signal bandwidth can pass through the digital down converter by setting the mixer to $F_s/4$ and selecting the first stage of filter output.

2.2.3 Resampler

The effective sample rate of the ADC can be throttled with the resampler. The resampler will discard data to produce an input to output ratio defined by a channel configuration setting. The ratio can be set to any integer from one to $2^{16}-1$. There is no filtering involved in this process, the ADC will simply appear to be operating at a lower sample rate if the ratio is greater than one.

2.2.1 DDC-1 Digital Down Converter (Model Specific)

The DDC-1 digital down converter shown in Figure 2-3 consists of a tuner followed by three stages of programmable FIR filters. Each FIR filter has a different structure, identified by a type number. Consult the DSP Function Library Reference Manual for details about the tuner and each filter type.

The down converter can produce a wide range of output bandwidths and sample rates by setting the desired filter characteristics and down sampling ratios through each stage. The full input signal bandwidth can be preserved by maintaining a complex sample rate of $F_s/2$ through the first two filter stages and bypassing the final stage. Very narrow output bandwidths can be achieved using all three stages.

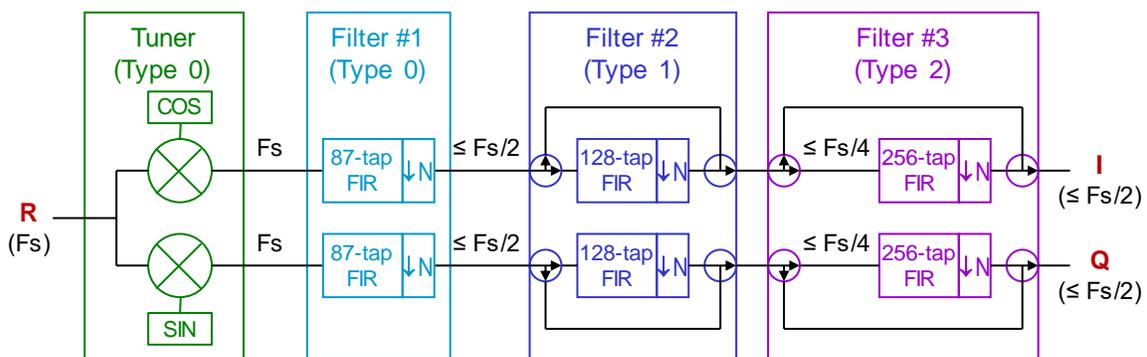


Figure 2-3 DDC-1 Digital Down Converter

The tuner uses a quadrature mixer to simultaneously frequency translate the input signal and generate in-phase (I) and quadrature (Q) components from the real (R) data stream. A digital numerically controlled oscillator generates the sine/cosine waveforms from a 23-bit phase increment value (Φ). This results in a tuning resolution of $F_s / 2^{23}$ and a tuning frequency of $\Phi * F_s / 2^{23}$, where F_s is the input sample rate.

The image created during the mixing process should be removed by the first filter stage. This is accomplished by setting the cut-off frequency less than or equal to half the normalized bandwidth. The downsampler (\downarrow) can be set to any integer between 2 and $2^{16}-1$. Setting the cut-off to half the bandwidth and down sampling by 2 will produce a complex representation of the full signal bandwidth supplied by the ADC.

The down sampler of the second filter stage can be set to any integer between 0 and $2^{16}-1$. Setting the value to 0 bypasses the filter entirely while a setting of 1 passes the filter output without down sampling. Either setting requires the third stage to be bypassed because the filter cannot accept an input rate greater than $F_s/4$.

The down sampler of the third filter stage can be set to any integer between 0 and $2^{16}-1$. Setting the value to 0 bypasses the filter entirely while a setting of 1 passes the filter output without down sampling.

2.2.2 DDC-2 Digital Down Converter (Model Specific)

The DDC-2 digital down converter shown in Figure 2-4 consists of a tuner followed by twelve stages of halfband FIR filters. The down converter can preserve the full input signal bandwidth through the entire channel by simply bypassing halfband filter number one through eleven. Consult the DSP Function Library Reference Manual for details about the tuner and filter type.

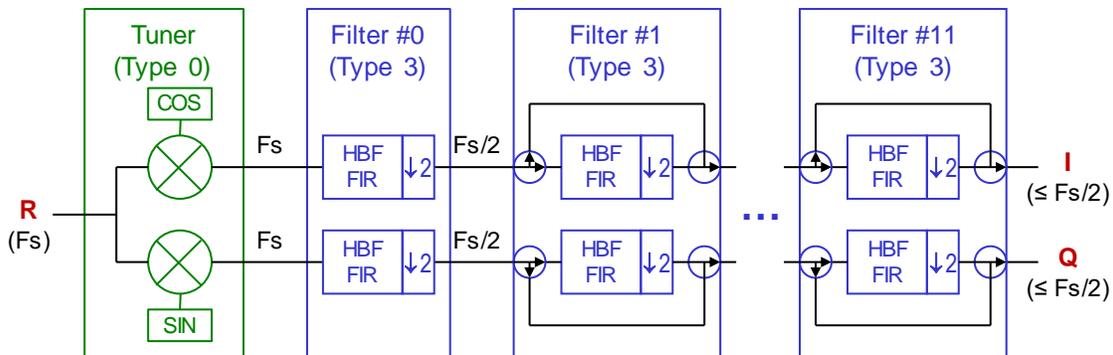


Figure 2-4 DDC-2 Digital Down Converter

The tuner uses a quadrature mixer to simultaneously frequency translate the input signal and generate in-phase (I) and quadrature (Q) components from the real (R) data stream. A digital numerically controlled oscillator generates the sine/cosine waveforms from a 23-bit phase increment value (Φ). This results in a tuning resolution of $F_s / 2^{23}$ and a tuning frequency of $\Phi * F_s / 2^{23}$, where F_s is the input sample rate.

A cascade of halfband FIR filters allows the application software to select one of twelve possible output bandwidths. The first filter simply removes the image created by the mixing process, producing a complex representation of the full signal bandwidth supplied by the ADC. Subsequent stages are either bypassed or further reduce the signal bandwidth by factors of two. The resulting Nyquist bandwidths can be computed from the input sample rate as follows:

Nyquist Bandwidth = $F_s / 2^N$, where $N = 0, 1, 2, 3 \dots 11$

The complex output rate of each stage is half the input rate unless it is bypassed. The frequency of the complex data rate will always match the Nyquist bandwidth.

Keep in mind that signals within the filter transition band will fold around the half bandwidth frequency when the output is downsampled by two ($\downarrow 2$). This will corrupt frequencies near the upper edge of the output Nyquist bandwidth leaving a usable bandwidth equal to 42% of the filter input bandwidth.

2.2.1 Output Formatter

The SigStream output format is based on the VITA Radio Transport (VRT) Standard (ANSI/VITA 49.0). There is a great deal of flexibility within the standard to set payload parameters and packet variables. Section 3.0 of this document describes the relevant features of the VRT as they apply to a SigStream datapath. There is no requirement to packetize the output data, but the VRT payload format is still used even when not encapsulated within a packet.

2.3 SigStream Transmitter Datapath

The internal elements of a SigStream transmitter datapath are shown in Figure 2-5. Digital sample data supplied by the host enters the datapath from the FIFO and exits through the crossbar switch.

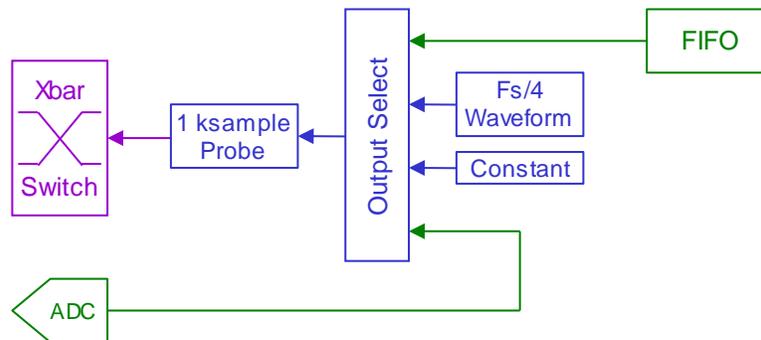


Figure 2-5 Transmitter Datapath

2.3.1 Output Select

Transmitter output data is supplied to a DAC from the crossbar switch. In normal operation, the crossbar switch receives sample data from the host computer through the FIFO. There are other options available to substitute sample data with a fixed frequency sinusoid, a constant value, or sample data routed from a receiver channel.

The $F_s/4$ waveform generator produces a full-scale digital sinusoid of frequency one quarter the sample clock. This well-defined pattern can be used to debug DAC register settings and verify the transmitter output amplitude.

The constant output can be used to verify the operation of the complex modulator internal to the DAC. A sinusoid can be generated by mixing the constant input value with the DAC NCO.

Transceiver products offer the additional option of routing ADC data from a receiver channel directly to the DAC of a transmitter channel to create a loopback. This can be used as a debug tool to verify the expected operation of the ADC and DAC with the SigStream datapath removed from the process.

2.3.2 Sample Probe

The sample probe is another debug feature that records the first 1024 samples sent to the DAC at the beginning of a transmission. This recording provides visibility at the end of the SigStream datapath that is otherwise not available in the digital domain. It can be a helpful resource if the DAC is producing an unexpected analog output.

2.4 Operating Modes

The SigStream datapath supports a variety of operating modes through channel configuration variables. The three most common type of configurations are continuous, snapshot, and periodic.

2.4.1 Continuous Mode

The continuous mode of operation is used to stream data samples from the ADC to host memory, or from host memory to the DAC. In this case, the PCIe bus must be able to sustain data transfers at the composite rate of all active channels.

The continuous operating mode will use the FIFO to rate buffer between the PCIe bus and ADC or DAC. The DMA engine will manage the transfer of data between the FIFO and the host computer.

Continuous mode operation can be initiated by any of the defined start events. Keep in mind that the FIFO needs to be primed with data prior to starting a DAC channel, otherwise there will be a short delay while the first block of data is transferred from the host.

2.4.2 Snapshot Mode

Snapshot mode is a special case of the continuous operating mode. The key difference is that the signal acquisition or generation will automatically halt after the specified number of samples.

There is a common misconception that the maximum snapshot size is equivalent to the FIFO capacity. It is perfectly valid to specify a snapshot that uses host memory for storage. The only restriction is that the composite data rate from all active channels cannot exceed the available PCIe bus bandwidth of the host.

It is important to consider the DMA block size when specifying snapshot length. There is a danger of leaving behind residue in the FIFO if the snapshot length includes a partial burst of data. These are samples collected in the receiver that do not get transferred to the host or extra samples in the transmitter that are not consumed by the DAC.

Snapshot operation can be initiated from any start event in the same manner used for continuous mode.



The snapshot length should be an integer multiple of the DMA burst size to avoid stranding data in the FIFO.

2.4.3 Periodic Mode

Periodic mode is a special case of the snapshot operating mode. There are two techniques available for receiving or transmitting data on a periodic interval.

In snapshot mode, each channel remains armed after the specified number of samples have been processed. This means that the channel will repeat the snapshot process in response to the next valid start event. The result is a periodic signal acquisition or generation sequence with a pulse repetition interval (PRI) that follows a trigger.

Alternatively, there is a PRI register assigned to each channel that can be used to set the collection interval based on the time of day clock. Rather than halting operation at the end of single snapshot, the same number of samples are collected at the start of every interval defined by this PRI value. The minimum PRI resolution is set by the clock period of the fractional seconds counter. This technique can be used to implement a PRI sequence that is guaranteed to maintain pulse to pulse coherence.

It is often convenient to enable headers when operating in periodic mode. These headers can be used as delimiters to separate data segments that are collected in each PRI.

3.0 VITA Radio Transport Standard (ANSI/VITA 49.0)

Red Rapids has adopted the IF Data Packet defined by the VITA Radio Transport (VRT) Standard (ANSI/VITA 49.0) as the data format for all products supported by the SigStream application software. Consequently, much of the nomenclature used throughout the code is derived from this standard. A detailed description of the IF Data Packet can be found in Section 6 of the VRT document available from VITA. A summary of the specification is provided here to assist in the explanation of the SigStream variables.

The template for an IF Data Packet is show in Figure 3-1, all fields are aligned to a 32-bit boundary. The packet consists of a mandatory Header and Data Payload along with five optional fields that convey additional information about the payload. Optional fields are completely absent from the packet if there is no information to convey. Remaining fields simply move up toward the header with no padding.

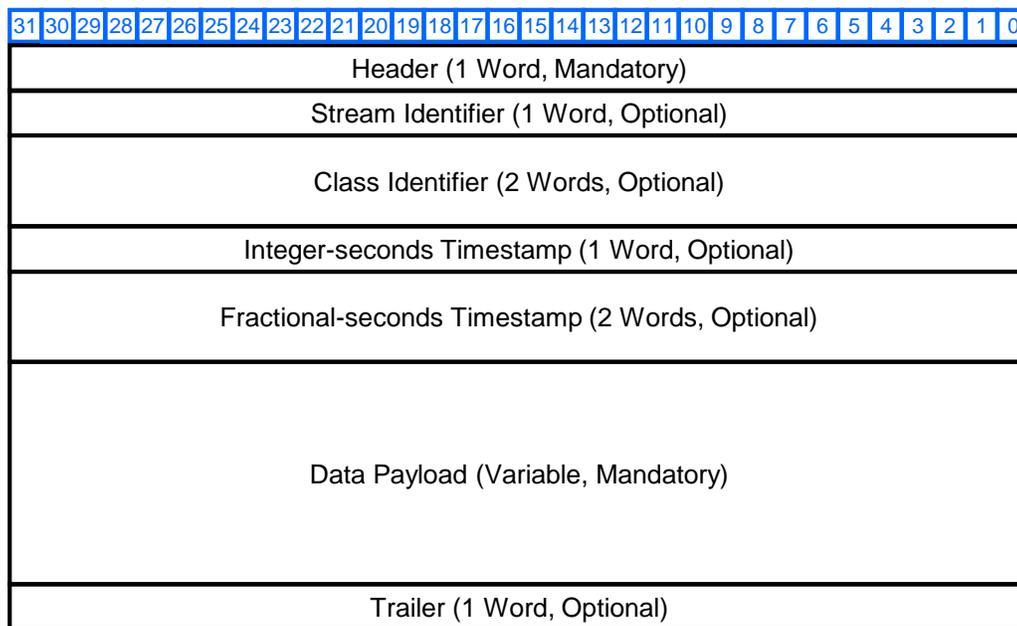


Figure 3-1 VRT (VITA 49.0) IF Data Packet Format

3.1 Header

The format of the 32-bit Data Packet Header is shown in Figure 3-2.

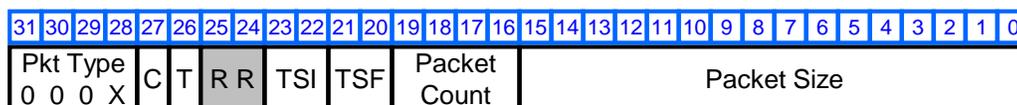


Figure 3-2 Data Packet Header Format

There are only two valid packet type codes assigned to an IF Data Packet. The value “0000” designates a packet without the optional Stream Identifier and “0001” indicates that a Stream Identifier is present in the packet.

The “C” bit field is set to one if a Class Identifier is present in the packet or set to zero if it is not included.

The “T” bit field is set to one if a Trailer is included at the end of the packet or set to zero if it is not.

The TSI and TSF codes indicate whether the timestamps are included in the packet. If timestamps are present, these codes also convey the origin of each value as shown in Table 3-1.

Table 3-1 TSI/TSF Code Meaning

Code	TSI	TSF
0	Integer Timestamp Not Present	Fractional Timestamp Not Present
1	Coordinated Universal Time (UTC)	Sample Count
2	GPS Time	Real Time (Not Supported)
3	Other	Free Running Count

The Packet Count field increments with each new IF Data Packet produced for a specific stream. The least significant bit of the modulo-16 count is the right-most bit in the field.

The Packet Size field indicates the total number of 32-bit words in the current IF Data Packet; including header, payload, and optional fields. The 16-bit field size limits the maximum packet size to 65,535 words.

3.2 Stream and Class Identifier

The Stream Identifier is an optional 32-bit number assigned to a packet stream. This designation can help back-end processes identify the source of IF data in a multi-stream environment.

The optional 64-bit Class Identifier field makes it possible to identify the Information Class used for the application and the Packet Class used to create each packet. Figure 6.1.3-1 shows the the three subfields included in the Class ID field. The first subfield contains the Organizationally Unique Identifier (OUI), an IEEE-assigned 24-bit number which indicates the identity of the company that created the Information Class and the Packet Class generating the IF Data Packet. The second field contains the Information Class code indicating which of that company’s Information Classes defines the Information Stream containing the Packet Stream. The third field contains the Packet Class code that identifies which of the company’s Packet Classes was used to make the packet.

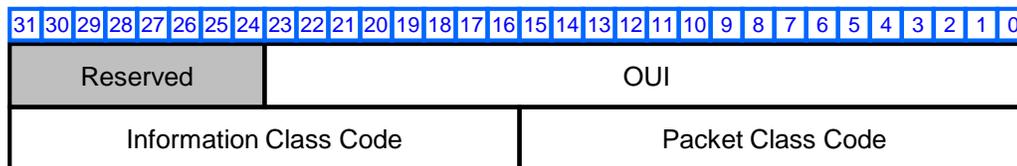


Figure 3-3 Class Identifier Contents

3.3 Timestamps

The optional Timestamp in an IF Data Packet is divided into an integer-seconds component and a fractional-seconds component. Together these two components precisely specify a reference time point for the first data sample contained in the packet.

The integer-seconds part consists of a single 32-bit word that specifies the reference time point only to one-second resolution. It may be used to convey UTC time, GPS time, or some user-specified time-code.

The fractional-seconds part consists of an unsigned 64-bit integer which occupies two consecutive 32-bit words. The first 32-bit word in the packet contains the most significant bits of the overall 64-bit value. It may be expressed as a number of sample count periods, a real-time value in picoseconds, or a free running modulo-N counter where N can be any positive integer up to 64.

3.4 Data Payload

The Data Payload contains a contiguous sequence of data samples encapsulated in a virtual container called the Item Packing Field. The Item Packing Field, shown in Figure 3-4, can occupy any size from one to sixty-four bits.

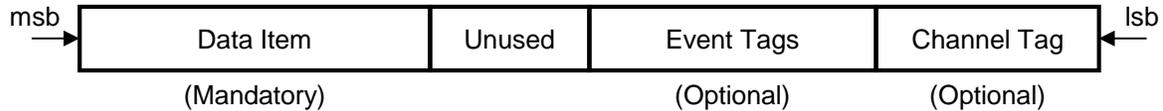


Figure 3-4 Item Packing Field

An Item Packing Field must contain exactly one Data Item that can occupy the entire field if no Event Tags or Channel Tags are present. The Data Item can be a real-valued sample, the real or imaginary component of a complex Cartesian sample, or the amplitude or phase component of a complex polar sample. The Data Item is always left justified in the Item Packing Field.

A Channel Tag is a label associating a Data Item with a particular signal conveyed by an IF Data Packet. Only one optional channel tag may be included in the Item Packing Field. It is always right justified in the Item Packing Field and must be no larger than fifteen bits.

An Event Tag is a bit used to indicate that a signal-related or processing-related event has occurred coincident with a Data Item in the payload. The Event Tag field can be no larger than seven bits and is located immediately to the left of the Channel Tag if it is present.

The size of the Item Packing Field may be larger than required to contain the included fields. In this case, the unused bits shall reside immediately to the right of the Data Item.

There are two methods of organizing the Item Packing Fields within the 32-bit payload words. Figure 3-5 illustrates the processing efficient method using a 15-bit Item Packing Field as an example. This approach simply left justifies the maximum number of Item Packing Fields that will fit in a 32-bit word without truncation, any remaining bits are set to zero.

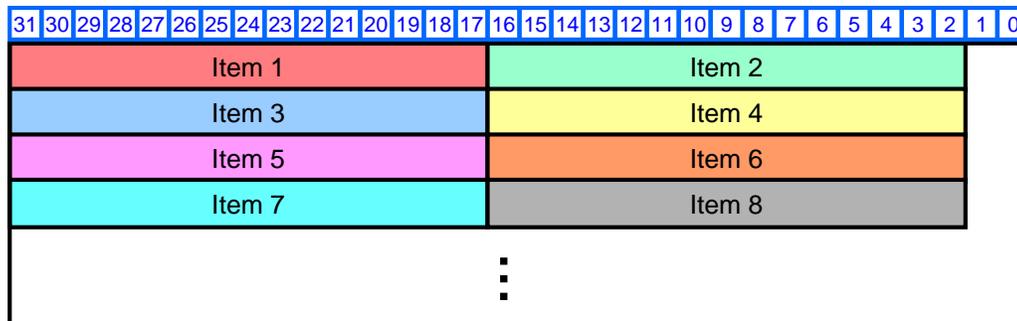


Figure 3-5 Processing Efficient Payload Organization

Figure 3-6 illustrates the link efficient method using the same 15-bit Item Packing Field as an example. This approach utilizes all of the available payload space by wrapping Item

Packing Fields between 32-bit word boundaries. This eliminates the zero padding of the processing efficient organization. The payload always ends with a complete Item Packing Field, so there may be some zero padding at the end of the final 32-bit word.

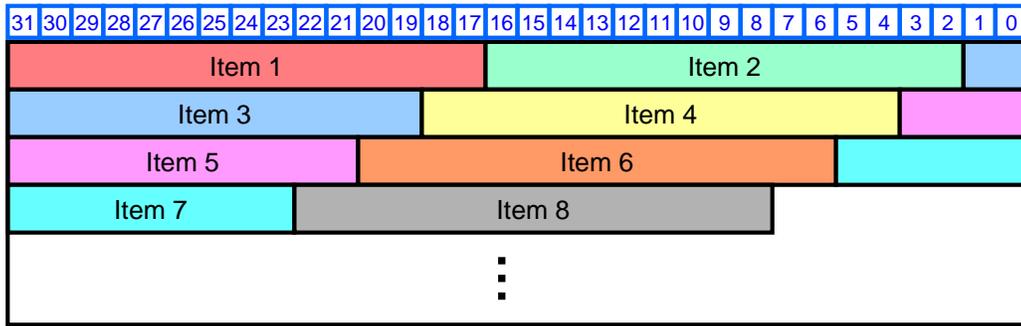


Figure 3-6 Link Efficient Payload Organization

When the Item Packing Fields contain complex samples, the in-phase value always precedes the quadrature value for Cartesian coordinates and the magnitude always precedes phase for polar coordinates.

3.5 Trailer

The trailer, shown in Figure 3-7, contains fields that indicate the validity of the data and the status of the processes producing that data. It also contains a field that indicates whether related context is being sent in one or more separate Context Packets.

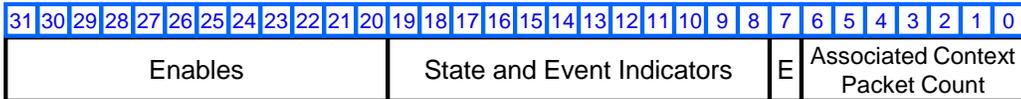


Figure 3-7 Trailer

The bits assigned to the Enables field act as a mask to the corresponding bit location in the State and Event Indicators field. An enable bit is set to zero when there is nothing to report in the associated State and Event Indicators bit. The enable bit is set to one when there is information available in one of the indicator bits listed in Table 3-2.

Table 3-2 State and Event Indicators

Enables Bit	Indicators Bit	Description
31	19	Calibrated Time Indicator
30	18	Valid Data Indicator
29	17	Reference Lock Indicator
28	16	AGC/MGC Indicator
27	15	Detected Signal Indicator
26	14	Spectral Inversion Indicator
25	13	Over-range Indicator
24	12	Sample Loss Indicator
23...20	11...8	User Defined Indicators

When the “E” bit is set to one the Associated Context Packet Count field contains a count of all of transmitted Context packets that are directly or indirectly associated with the IF Data Packet. When the “E” bit is set to zero, the Associated Context Packet Count is undefined.

4.0 Hardware Configuration

Default variable settings are assigned to the hardware configuration structures contained in the [default_hardware.config.c](#) source file. The exact contents of this file will vary between product families due to the characteristics of specific chips used in the design. There may also be model number specific structures within a product family.

4.1 Time of Day (TOD) Clock Settings

```
typedef struct
{
    unsigned TODSeconds;
    unsigned ClockFrequency;
    unsigned TSFCode;
    unsigned TSICode;
} s_TODSettings;
```

An embedded time of day (TOD) clock is used to produce IF data packet timestamps and start/stop triggers for the synchronizer. Time is measured in two components; integer seconds and fractional seconds. The `TODSeconds` variable is used to set the integer seconds value. This value is often synchronized to an external time source such as GPS or UTC. The absolute time of day is read from the external source and then written to the TOD clock. The entire operation must be completed before the external source increments to the next second to achieve synchronization.

An external source usually supplies a 1 PPS strobe that can be connected to the Red Rapids device to increment the integer seconds value. The `ClockFrequency` variable must be set to zero if an external 1 PPS signal is active. If an external 1 PPS is not available, the `ClockFrequency` variable should be set to a value equivalent to the frequency of the ADC/DAC sample clock. An internal counter will then be used to increment the seconds counter autonomously when the specified number of periods have passed. The fractional seconds value will be reset to zero coincident with every increment of the integer seconds value.

It is impossible to achieve exact alignment between an external source and the TOD clock. If a 1 PPS strobe is used, there is no way to predict the phase relationship between the leading edge of the strobe and the leading edge of the clock that will be used to increment the integer seconds. In general, the TOD clock will lag the external source by two sample clock periods with an additional ambiguity of plus or minus one sample clock period. If a 1 PPS strobe is not used, the fractional seconds counter will start from zero when the integer seconds value is loaded. This will cause the TOD clock to lag the external source by the time it took to read and write the value, which is likely some fraction of a second. This latency will remain constant if the sample clock is phase locked to the external time source and there will be no ambiguity if the sample clock is set to an integer value. If the sample clock is non-integer, then the TOD clock will accumulate a small error every second. The error will be equivalent to the difference between the non-integer sample clock frequency and the integer value assigned to the `ClockFrequency` variable.

The integer seconds component of the TOD clock can be read from software but the fractional seconds component cannot. It is impractical to supply a fractional seconds value since it transitions within a clock domain that is not coherent to the bus clock. The

nondeterministic nature of the bus timing also contributes an ambiguity that is much greater than the fractional seconds resolution.

Timestamps are only available when data is received from the hardware in IF Data Packets as discussed in Section 2.0. This is accomplished by setting the `PacketSize` variable to a non-zero value in the IF Data Packet structure (see section TBD). The integer-seconds and fractional-seconds timestamps will only be included in the IF Data Packet if their corresponding codes (`TSICode/TSFCode`) are also non-zero.

Setting the `TSICode` variable to one indicates that the TOD clock integer seconds value is synchronized to Coordinated Universal Time (UTC). Setting the variable to two indicates that GPS Time is the source. The `TSICode` should be set to three if the TOD clock is synchronized to any other source.

Setting the `TSFCode` variable to one will report fractional-second timestamps measured in periods of the sample clock. This value will never exceed the sample clock frequency since that is the one second roll-over value. Setting the `TSFCode` to three will produce a free running 64-bit fractional seconds count that increments with each period of the sample clock. The counter will reset back to zero when the maximum 64-bit value is reached. The integer seconds timestamp is automatically suppressed when the `TSFCode` is set to three since it has no meaning in the context of a free running count.

4.2 LED Source Selection

```
typedef struct
{
    unsigned Yellow[ ];
    unsigned Green[ ];
} s_LEDSource;
```

There are two status LED indicators available to visually monitor trigger and channel activity. Trigger activity is displayed on the yellow LED and channel activity is displayed on the green LED.

There are multiple triggers available to synchronize processing through a channel. The `Yellow[]` array allows any selected trigger shown in Table 4-1 to flash the light. Setting a variable in the array to zero disables monitoring of the trigger assigned to that index. Setting a variable to one causes the corresponding trigger to flash the LED when an active edge is detected. Multiple trigger monitors can be enabled simultaneously, but there will be no way to determine which specific trigger produced a flash.

Table 4-1 LED A Trigger Monitor Assignment

Index	Trigger Source
0	Software
1	Time of Day
2	Coax Trigger
3	GPIO Trigger

The `Green[]` array is used to enable activity monitoring in each channel. A channel is considered active when the synchronizer state machine is operating between start and stop events. Setting a variable in the array to zero disables monitoring of the channel assigned to that index, as shown in Table 4-2. Setting a variable to one causes the

assigned channel to flash the LED when the channel is actively processing data. Multiple channel monitors can be enabled simultaneously, but there will be no way to determine which specific channel produced a flash.

Table 4-2 LED B Channel Monitor Assignment

Index	Trigger Source
0	Channel #1
1	Channel #2
2	Channel #3
3	Channel #4

4.3 GPIO Settings

```
typedef struct
{
    unsigned Disconnect;
    unsigned Direction[ ];
} s_GPIOSettings;
```

The GPIO connector is a bidirectional interface that can be used to accept external input triggers or drive an external output signal. Setting the **Disconnect** variable to one will electrically isolate (tristate) all of the pins on the GPIO connector. The variable must be set to zero if any of the GPIO signals are going to be used as inputs or outputs.

The **Direction[]** array sets each signal assigned to the connector as an input when the variable is set to zero, an output when the variable is set to one, or a debounced input when the variable is set to three. The debounce option applies a hysteresis equivalent to four sample clock periods. This prevents ringing on the input from producing multiple unintended . Table 4-3 shows the mapping of array index to GPIO signal.

Table 4-3 GPIO Setting Assignment

Index	Connector Pin
0	GPIO #1
1	GPIO #2
2	GPIO #3
3	GPIO #4
4	GPIO #5
5	GPIO #6

4.4 Device Configuration Settings

```
typedef struct
{
    chip_type chip_function;
    .
    .
    .
    chip_type chip_function;
} s_MxxxDevices;
```

Every Red Rapids product is assigned a three-digit model number that uniquely identifies the chips used on that product. A device configuration structure that lists all of the chips (`chip_type`) is assigned to each model number (`s_MxxxDevices`). The name of each `chip_type` corresponds to the vendor part number for that device. The `chip_function` variable assigns values to the configuration registers internal to each `chip_type`. The following is an example of the structure that would be used for a design that includes the AnalogDevices AD9512 clock distribution chip and two Texas Instruments ADS42LB69 dual channel ADC chips.

```
typedef struct
{
    s_AD9512 ClockDist;
    s_ADS42LB69 RX1TO2_ADC;
    s_ADS42LB69 RX3TO4_ADC;
} s_MxxxDevices;
```

Please refer to the device vendor datasheet for a list of configuration registers specific to the part number. These registers are typically accessed through a serial port (SPI, I2C, etc.) provided on the device. Figure 4-1 contains a partial listing of the configuration registers internal to the Analog Devices AD9512 clock distribution chip that is used on several Red Rapids products. Each register is assigned an address and configuration variable, which are both eight bits in this example.



REGISTER MAP AND DESCRIPTION

SUMMARY TABLE

Table 17. AD9512 Register Map

Addr (Hex)	Parameter	Bit 7 (MSB)	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0 (LSB)	Def. Value (Hex)	Notes	
00	Serial Control Port Configuration	SDO Inactive (Bidirectional Mode)	LSB First	Soft Reset	Long Instruction	Not Used				10		
01 to 33		Not Used										
	FINE DELAY ADJUST											Fine Delays Bypassed
34	Delay Bypass 4	Not Used							Bypass		01	Bypass Delay
35	Delay Full-Scale 4	Not Used		Ramp Capacitor <5:3>			Ramp Current <2:0>			00	Max. Delay Full-Scale	
36	Delay Fine Adjust 4	Not Used		5-Bit Fine Delay <5:1>				Not Used		00	Min. Delay Value	
37, 38, 39, 3A, 3B, 3C		Not Used										

Figure 4-1 Analog Devices AD9512 Configuration Registers

The `chip_function` variable is a two dimensional array that stores the address and configuration value for every internal register of the device. There is a total of twenty-two registers in the AD9512 example, but the addresses are not contiguous. Referring to Figure 4-1, notice that addresses 0x01 through 0x33 are not used. This address range is skipped in the `chip_function` array, so index #1 in the first dimension of the array is assigned to address 0x34 in the AD9512 register map.

The SigStream application code will include a default device configuration structure for every model number available in a product family, but only one will be used to configure the hardware based on the model number returned by a query to the feature identification register.

5.0 Operating Configuration

Default variable settings are assigned to the operating configuration structures contained in the [default_operating_config.c](#) source file. There are separate default values assigned to receiver (RX) and transmitter (TX) channels, but both use the same `s_Channel` structure type. Some products are only equipped with one type of channel, so the default values assigned to the other type may be ignored. The SigStream application will query a feature identification register in the Red Rapids device to determine which channel types are available.

Although there may be multiple RX or TX channels in a device, there is only one set of default settings defined for each channel type. This means that all channels of the same type (RX or TX) will be configured identically for demonstration purposes. This is not a limitation of the hardware or software; each channel can be configured with unique settings.

5.1 Channel Settings

```
typedef struct
{
    s_SwitchSelect SwitchSelect;
    s_PayloadFormat PayloadFormat;
    s_DatapathControl DatapathControl;
    s_Synchronizer Synchronizer;
    s_TODTriggers TODTriggers;
    s_EventDuration EventDuration;
    s_IFpktFormat IFpktFormat;
} s_Channel;
```

Each member of the channel structure is another structure that defines a specific channel characteristic. The channel structure is identical for RX or TX channels, although some variables may have slightly different meaning.

The Channel[] array declared in main() is allocated space to store all of the available RX and TX channel configuration settings using the `s_Channel` structure.

The following sections explain the function of each channel structure member and the differences in how they are applied to an RX channel versus a TX channel.

5.2 Switch Settings

```
typedef struct
{
    unsigned ChannelSelect;
    unsigned ChannelBypass;
    unsigned Amplitude;
} s_SwitchSelect;
```

Any RX channel can be connected to any RX input (ADC) and any TX output (DAC) can be connected to any TX channel. There are also special settings available to assist debug and hardware initialization. For demonstration purposes, there is one RX channel connected to each RX input and one TX channel connected to each TX output in the SigStream application.

The **ChannelSelect** and **ChannelBypass** variables determines the mapping between ADC/DAC hardware resources and the available RX/TX channels. The settings have different meaning when applied within the context of an RX channel versus a TX channel.

All of the valid RX switch select settings are shown in Figure 5-1. The function of the switch is to select which source should supply data to a particular RX channel. Not all products will be equipped with the eight RX inputs shown in the drawing. The number of available RX channels will always match the number of RX inputs on the product.

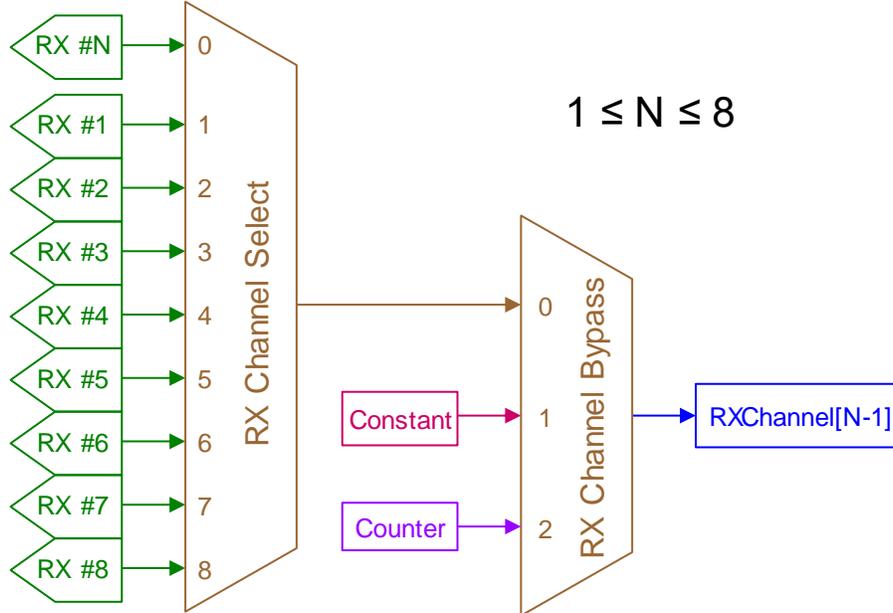


Figure 5-1 RX Switch Select Settings

The zero setting is the default used to demonstrate the SigStream application. This value connects the RX #N input to RXChannel[N-1], where N is any integer between one and the total number of RX channels available on the product. For example, RX #1 connects to RXChannel[0] and RX #2 connects to RXChannel[1].

Setting the **ChannelSelect** variable between one and eight will connect any RX channel to a specific RX input. For example, setting **ChannelSelect** to one in both RXChannel[0] and RXChannel[1] will route the RX #1 input to both RX channels. This will leave one input stranded since the number of available RX channels matches the number of inputs.

There are two additional **ChannelBypass** settings that will produce a defined input pattern for test and debug purposes. Setting the **ChannelBypass** variable to one will connect the RX channel to the constant value set by the **Amplitude** variable. The two setting will supply a binary count to the channel.

All of the valid TX switch select settings are shown in Figure 5-2. The function of the switch is to select which source should supply data to a particular TX output. Not all products will be equipped with the eight TX channels shown in the drawing. The number of available TX channels will always match the number of TX outputs on the product. If the product also contains RX channels, TX channel numbering will begin after the last RX channel. For example, a board with two RX channels will begin with TX #3.

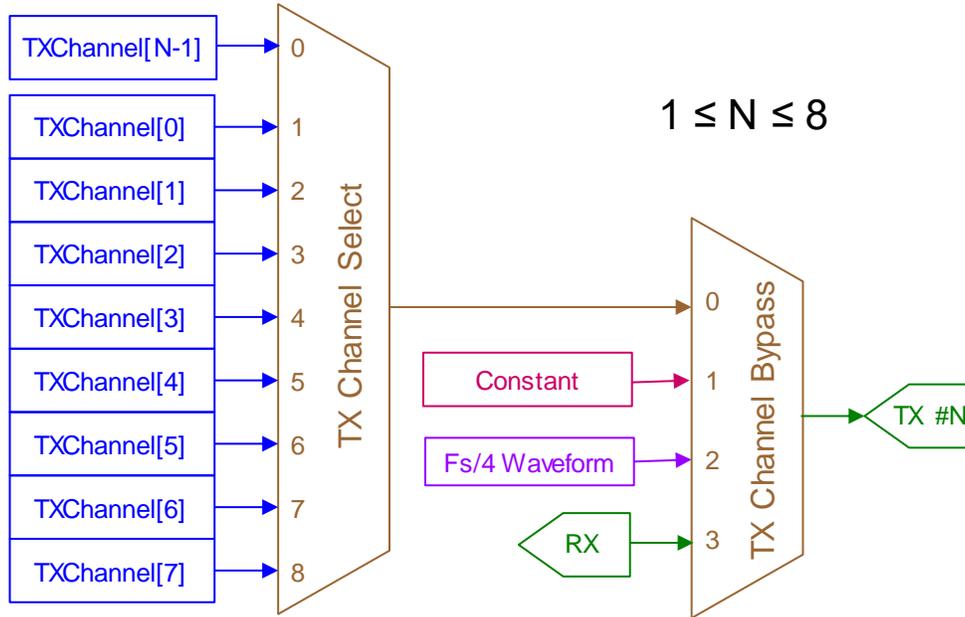


Figure 5-2 TX Switch Select Settings

The zero setting is the default used to demonstrate the SigStream application. This value connects TXChannel[N-1] to the TX #N output. The range of N depends on whether RX channels are also available on the product. If there are no RX channels, N is any integer between one and the total number of TX channels available. For example, TXChannel[0] connects to TX #1 and TXChannel[1] connects to TX #2. If RX channels are present, N is any integer between the total number of RX channels plus one and the total number of RX channels plus the total number of TX channels. So a product with two RX channels would connect TXChannel[2] to TX #3 and TXChannel[3] to TX #4. Any configuration settings assigned to TXChannel[0] and TXChannel[1] would be ignored since they are not physically present in the product.

Setting the **ChannelSelect** variable between one and eight will connect any TX output to a specific TX channel. For example, setting **ChannelSelect** to zero in both TXChannel[0] and TXChannel[1] will route TXChannel[0] to both the TX #1 and TX #2 outputs. The remaining TXChannel[1] settings will be ignored since the TX #1 output will be supplied from the TXChannel[0] datapath.

There are three additional **ChannelBypass** settings that will produce a defined output pattern for test and debug purposes. Setting the **ChannelBypass** variable to one will connect the TX output to the constant value set by the **Amplitude** variable. The two setting will supply a sinusoid waveform of frequency $F_s/4$ to the TX output.

Setting the **ChannelBypass** variable to three is only valid on products that have both RX and TX channels. This setting will route signal data directly from each RX input to each TX output, pairing the channels in sequential order. For example, a product with two RX channels and two TX channels will route RX #1 to TX #3 and RX #2 to TX #4. The pairing continues until either channel type is exhausted.

5.3 Payload Format Settings

```
typedef struct
{
    unsigned PayloadPacking;
    unsigned Endianness;
    unsigned DataItemSize;
    unsigned EventTagOrder;
    unsigned EventTagMask;
    unsigned EventTagIndex;
    unsigned ItemPackingFieldSize;
    unsigned ChannelTag;
} s_PayloadFormat;
```

The VRT Data Payload consists of a virtual container called the Item Packing Field as described in Section 3.4. The contents of the Item Packing Field are determined by the settings in the `s_PayloadFormat` structure. The VRT standard limits the size of an Item Packing Field to sixty-four bits, but the SigStream hardware imposes a further restriction of thirty-two bits since that is twice the size of the largest Data Item produced by a Red Rapids product.

The `s_PayloadFormat` structure only applies to RX channels. The VRT processing efficient payload organization shown in Figure 3-5 is applied to TX channels, but there is no option to resize Sata Items or add Event Tags.

The `PayloadPacking` variable determines how the Item Packing Fields are arranged within the 32-bit payload words. Setting the variable to zero results in processing efficient organization (Figure 3-5) and setting the variable to one results in link efficient organization (Figure 3-6).



The `PacketSize` variable must be set to a value greater than zero if the link efficient organization is selected. See section 5.8 for further details.

The `Endianness` variable selects big endian or little endian data order, but this feature is not yet supported.

The `DataItemSize` variable sets the number of bits allocated to the Data Item that appears in each Item Packing Field. The default value matches the number of bits produced by the ADC in an RX channel. The supported range of values includes even numbers from two to thirty-two. If the Data Item is set to a number less than the default, then the values will be rounded to the requested number of bits. If the Data Item is set to a number greater than the default, then values will be sign extended to the requested number of bits. When a Data Item is rounded, fractional values less than 0.5 will be rounded toward zero to the nearest integer and fraction values greater than or equal to 0.5 will be rounded to the nearest integer away from zero.

Optional Event Tags can be included in the Item Packing Field. There are two Event Tags produced by Red Rapids hardware; ADC over-range and channel synchronization. The ADC over-range tag is set to one when a Data Item is produced while the ADC is driven

beyond the specified analog input range. The channel synchronization tag is set to one when a sync strobe is issued to initialize elements of the datapath.

The **EventTagOrder** variable is used to select the position of the tags in the Event Tags field. When set to zero, the channel synchronization tag is located in the LSB of the 2-bit field and the ADC over-range tag is located in the MSB. Setting the value to one will swap the positions.

The **EventTagMask** will suppress specific tags based on the settings listed in Table 5-1.

Table 5-1 Event Tag Mask Settings

Value	Result
0	Suppress both tags.
1	Suppress the ADC over-range tag.
2	Suppress the channel synchronization tag.
3	Both tags are active.

Suppressing a tag does not remove it from the Item Packing Field, it will simply report zero at all times.

The **EventTagIndex** defines the location of the Event Tags field within the Item Packing Field as an offset from the LSB. The Event Tags field is located to the immediate left of the optional Channel Tag field in the Item Packing Field. Since the Channel Tag field is right justified within the Item Packing Field, the **EventTagIndex** effectively defines the size of the Channel Tag. The SigStream hardware limits the **EventTagIndex** to any even integer between zero and fourteen, which limits the size of the optional Channel Tag to fourteen bits even though the VRT standard allows fifteen bits. Setting the **EventTagIndex** to the Data Item LSB position in the Item Packing Field will allow the Channel Tag to pass without the Event Tags.

The **ItemPackingFieldSize** variable sets the number of bits allocated to the Item Packing Field; which is the sum of the Data Item field, Event Tag field, Channel Tag field, and any unused bits. It is very common to set the Item Packing Field size equal to the Data Item size when the optional Event Tags or the Channel Tag are not needed, this is the minimum value. The **ItemPackingFieldSize** effectively defines the number of unused bits since the Event Tags and Channel Tag field sizes are set by other variables.

 The **ItemPackingFieldSize** options are limited to 8 bits, 10 bits, 12 bits, 14 bits, 16 bits, 20 bits, 24 bits, and 32 bits. Setting the value to zero will size the data item to match the ADC output or DAC input. The Model 276 is further restricted to 8 bits, 10 bits, 12 bits, or 16 bits.

The **ChannelTag** variable can be any integer that fits within the Channel Tag field size as defined by the **EventTagIndex**.

5.4 Datapath Control Settings

```
typedef struct
{
    unsigned Transform[];
```

```

unsigned Residue[];
unsigned SyncInitMask[];
unsigned SyncCycleMask[];
} s_Datapath Control;

```

The Transform[] array activates optional signal conditioning functions that can be applied to the sample data. The first index of the array, Bandwidth Select, selects an optional digital down converter and resampler at the output of the ADC in RX channels. Setting this variable to zero bypasses all filters and streams real samples through the datapath. The data flows at a rate equal to the ADC sample clock (Fs) and the signal bandwidth is defined by the Nyquist theorem (Fs/2).

Setting Bandwidth Select to an integer greater than zero will stream data in complex format instead of real samples. Setting this variable to one produces an analytic signal equivalent to the real signal on the in-phase (I) component and the Hilbert Transform of the real signal on the quadrature (Q) component of the complex value.

The first index in the array, Bandwidth Select, selects an optional digital down converter and resampler at the output of the ADC. Setting this variable to zero bypasses all filters and streams real samples through the datapath. The data flows at a rate equal to the ADC sample clock (Fs) and the signal bandwidth is defined by the Nyquist theorem (Fs/2).

Setting Bandwidth Select to an integer greater than zero will stream data in complex format instead of real samples. Setting this variable to one produces an analytic signal equivalent to the real signal on the in-phase (I) component and the Hilbert Transform of the real signal on the quadrature (Q) component of the complex value.

Table 5-2 Transform Array Definition

Index	Function	Integer Range	Transform
0	DDC-1 Bandwidth Select	$0 \leq n \leq 11$	n = 0: BW = Fs/2 (real) n > 0: BW = Fs/2 ⁿ (complex)
	DDC-2 Bandwidth Select	$0 \leq n \leq 2$	n = 0: BW = Fs/2 (real) n = 0: BW = Fs/2 (complex) n = 1: BW determined by filters
1	Real Resampler	$0 \leq n \leq 2^{16} - 1$	n = 0: Product default setting n > 0: Channel rate = Fs / n
2	Frequency Translation	$0 \leq n \leq 2^{22}$	Fcenter = (n * Fs) / 2 ²³
3	Phase Offset	$0 \leq n \leq 2^{23} - 1$	Phase = (n * π) / 2 ²³
4	Reserved		
5	Reserved		

Units equipped with DDC-1 will load filter configuration files when Bandwidth Select is set to the value two. The output bandwidth and sample rate are determined by the cut-off frequency and down sample ratio programmed into each filter stage.

Units equipped with DDC-2 will produce a complex output sample rate of Fs/2ⁿ, where n is a Bandwidth Select value greater than zero. The Nyquist bandwidth is also Fs/2ⁿ, but the usable signal bandwidth is 0.42*Fs/2ⁿ due to the roll-off of the digital lowpass filter. For example, setting the variable to one will still pass the maximum available Nyquist bandwidth (Fs/2) but the complex samples will flow at half the rate of real samples carrying the same bandwidth.

The second index in the array, Real Resampler, will discard samples out of the ADC. The variable simply defines the ratio of samples to discard or replicate. There is no filtering applied to this operation, so the signal bandwidth of the channel is effectively reduced. For example, setting the variable to two will produce a result equivalent to the ADC operating at half the actual sample rate. The Real Resampler can only be used when Bandwidth Select is set to zero.

The third index in the array, Frequency Translation, programs a complex mixer that will down convert an RX spectrum. This variable will only be recognized if Bandwidth Select is set to a non-zero value. The translation operation will effectively map a center frequency of $(n * F_s) / 2^{23}$ to zero in the complex baseband spectrum.

The fourth index in the array, Phase Offset, controls the starting phase of the digital local oscillator connected to the complex mixer. Each count adds a phase increment offset of $(n * \pi) / 2^{23}$ from zero.

Figure 5-3 illustrates the primary components involved in the flow of data from source to destination through a TX or RX channel. It is important to realize that the ADC or DAC is always producing or consuming sample data, there is no way to turn them off. Instead, the flow control function effectively isolates the ADC or DAC from the datapath when the channel is inactive. It is possible to strand samples already in the pipeline when the flow of data is suspended.

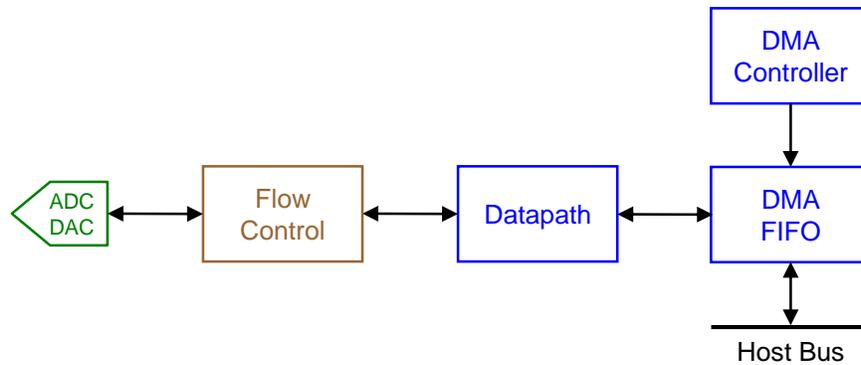


Figure 5-3 TX/RX Data Flow

The `Residue[]` array determines the disposition of data remaining in the pipeline when a channel is stopped or a cycle completes (`CycleCount` \neq 0). It also determines the state of the DMA controller when the channel is started or each new cycle begins. The array stores two variables as listed in Table 5-3.

Table 5-3 Residue Array Definition

Index	Function	Integer Range
0	Channel Stop	$0 \leq n \leq 7$
1	Channel Cycle	$0 \leq n \leq 7$

The settings listed in Table 5-4 can be applied to both the Channel Stop and Channel Cycle variables. Each value represents a unique combination of actions that are applied to the datapath registers, DMA FIFO, and DMA controller.

Table 5-4 Residue Variable Settings

Value	Channel Type	Datapath	DMA FIFO	DMA Controller
0	RX & TX	Clear	Clear	Reset
1	RX & TX	Clear	Clear	Hold
2	RX & TX	Hold	Hold	Reset
3	RX & TX	Hold	Hold	Hold
4	RX & TX	Flush	Flush	Reset
5	RX & TX	Flush	Flush	Hold
6	RX	Flush	Hold	Reset
7	RX	Flush	Hold	Hold

There are three actions that can be applied to residue. First, the pipeline can be cleared to remove any remaining sample data. This action ensures that the pipeline will be primed with new samples the next time the channel is started. Notice that the datapath and FIFO are always cleared together. The second option is to simply hold the data samples in place. As the name implies, the residue will remain in the pipeline and will be processed ahead of fresh data when the channel is restarted. Finally, an RX pipeline can be flushed. Residue will be forced to a destination instead of pushed along by new samples entering the pipeline. The details of how each setting will affect operation depends on whether the channel is RX or TX.

Samples cannot be stranded in an RX datapath, data will flow to the FIFO even after the channel is shut off. However if the data is packetized, there is always a possibility that the channel will stop on an incomplete packet boundary. The datapath residue setting determines how an incomplete packet is handled. Clearing the datapath will reset the packetizer and eliminate any partial packet already stored in the FIFO. This may result in a partial packet stored in the host DMA buffer. Holding the datapath will allow the packet to be completed with data produced when the channel restarts. This will obviously result in a waveform discontinuity within the packet. Flushing the datapath will force the packet to complete by simply replicating the final payload data. The completed packet will enter the FIFO where it can either be held or flushed along with other residue. Flushing the FIFO will append a repeating pattern of data to the end of the last DMA burst after the final sample or packet has exited the FIFO.

Stopping a TX channel halts the flow of data from the FIFO to the DAC, potentially stranding samples anywhere in between. Clearing the datapath and FIFO will remove all residue. Holding the datapath and FIFO will allow the residue to eventually reach the DAC when the channel is restarted. It will be followed by new data once the residue has been depleted. Flushing both the datapath and FIFO will force the entire pipeline to empty. Flushing the datapath while holding the FIFO is not an option for TX channels.

The DMA controller continuously cycles through DMA buffer pages as data is transferred. Each page consists of a programmed number of bursts. The state of the controller can be set to hold the current page and burst pointer when a channel is stopped or reset both to zero. Reference the DMA on Demand Operating Guide for further details.

The `SyncInitMask[]` array determines which elements of the datapath will respond to the synchronization strobe immediately preceding the first start event. The `SyncCycleMask[]` array performs the same function with respect to the synchronization strobe immediately preceding subsequent start events (`CycleCount` \neq 0). The arrays store two variables as listed in Table 5-5.

Table 5-5 SyncInitMask Array Definition

Index	RX Function	TX Function	Integer Range
0	Binary Data Counter	Reserved	0 (no sync) or 1 (sync)
1	Resampler or DDC	Resampler or DUC	0 (no sync) or 1 (sync)
2	Reserved	Reserved	
3	Reserved	Reserved	

A binary counter can be selected as the input to an RX channel in place of ADC sample data as described in section 0. This feature is provided as a debug function to analyze channel behavior with a known input pattern. The first index of the two synchronization arrays is used to select when this counter is reset to zero.

The second index of each array determines when the resampler or down converter of any channel will be initialized. The resampler discards ADC samples associated produced by an RX channel. Initializing the resampler establishes a starting point for this process. For example, the resampler will pass the sample aligned with the synchronization strobe and start the discard process with the following sample. The digital down converter behaves much like the resampler in each channel except that it processes complex pairs instead of real samples. It also has a mixer that is reset to the programmed starting phase when the synchronization strobe is applied.

5.5 Synchronizer

```
typedef struct
{
    unsigned StartEvent;
    unsigned StartOffset;
    unsigned StopEvent;
    unsigned StopOffset;
    unsigned CycleSelect;
    unsigned CycleStop;
} s_Synchronizer;
```

The synchronizer supervises all activity associated with data movement through a channel. It manages flow control and ensures that DMA transfers occur only when expected by the host. The `s_Synchronizer` structure defines the specific actions that will initiate and terminate channel processing.

The `StartEvent` variable selects which action will begin the flow of data through a channel. First the channel is enabled using the `FlowControl()` API function. This places the synchronizer in an armed state waiting for the selected start event. Once the event occurs, data will begin to flow after the time interval specified by the `StartOffset` variable.

The `StopEvent` variable selects which action will halt the flow of data through a channel. The synchronizer begins monitoring for the `StopEvent` immediately following a `StartEvent`. Once the `StopEvent` event occurs, data flow will be shut off after the time interval specified by the `StopOffset` variable.

The `StartEvent` and `StopEvent` variables share almost identical definitions as listed in Table 5-6. Most of the events are associated with external signals supplied through the coaxial or GPIO hardware connectors. Each hardware input can be set to detect the rising

or falling edge of the signal. The `SoftwareTrigger()` API function can also be chosen as an event. This option allows multiple channels to respond synchronously to a single software command instead of commanding each channel individually.

Table 5-6 StartEvent and StopEvent Variable Settings

Value	Start Event	Stop Event
0	Immediate	Software Command
1	Software Trigger	Software Trigger
2	Coax Trigger Rising	Coax Trigger Rising
3	Coax Trigger Falling	Coax Trigger Falling
4	Reserved	Sample Count
5	Time of Day	Time of Day
6	GPIO Trigger Rising	GPIO Trigger Rising
7	GPIO Trigger Falling	GPIO Trigger Falling

The zero setting shown in Table 5-6 produces a different result when it is applied to a `StartEvent` versus a `StopEvent`. Data will begin flowing immediately after the channel is enabled if the `StartEvent` is set to zero. The same value will cause the channel to wait for software to issue a command through the `FlowControl()` API function when it is assigned to a `StopEvent`. Setting the `StopEvent` variable to four will activate a sample counter that terminates data flow when the requested number of samples have transferred. Prior to 2017 firmware release dates the samples are counted at the output of the ADC in an RX channel or at the input to the DAC in a TX channel and not at the lower rates that may be selected through the `Transform[]` array. Firmware releases dated 2017 or later count samples at the rates selected through the `Transform[]` array, which may be lower than the ADC output or DAC input rates. The time of day may also be used as a `StartEvent` or `StopEvent`. See Section 4.1 for more information about the time of day clock operation.

The `StartOffset` and `StopOffset` variables perform identically. They both measure time in increments of sample clock periods to establish a delay interval. The range can be set from 0 to $2^{16}-1$. The computed time interval will be referenced to the ADC sample clock in an RX channel and the DAC sample clock in a TX channel.

Each channel can be programmed to cycle on and off automatically. The `CycleSelect` variable selects what criteria, if any, should mark the beginning of each new cycle. The three available options are listed in Table 5-7. The zero setting disables the channel immediately after the first `StopEvent` is detected, resulting in a single cycle. Setting the variable to one causes each new cycle to begin with whatever action is selected by the `StartEvent` variable. The final option, setting two, commands the synchronizer to start each new cycle on a periodic frame interval.

Table 5-7 CycleSelect Variable Settings

Value	Cycle Start Event
0	None / Single Cycle
1	<code>StartEvent</code>
2	Frame

The `CycleStop` variable selects the criteria that will terminate cycling when `CycleSelect` is set to a non-zero value. The three available options are listed in Table 5-8. The zero setting instructs the synchronizer to continue cycling until software issues a command

through the `FlowControl()` API function. Setting the variable to one activates a counter that terminates cycling when the requested number of cycles have completed. The final option, setting two, terminates cycling based on a time of day. See Section 4.1 for more information about the time of day clock operation.

Table 5-8 CycleStop Variable Settings

Value	Cycle Start Event
0	Software Command
1	Cycle Count
2	Time of Day

5.6 Time of Day Triggers

```
typedef struct
{
    unsigned IntegerSecondsStart;
    unsigned FractionalSecondsStart;
    unsigned IntegerSecondsStop;
    unsigned FractionalSecondsStop;
} s_TODTriggers;
```

The `s_TODTriggers` structure stores the time of day start value and stop value that can trigger events in the synchronizer as described in Section 5.5. Each value consists of an integer seconds component and fractional seconds component. All of the variables will accept an input range from 0 to $2^{32}-1$. See Section 4.1 for more information about the time of day clock operation.

5.7 Event Duration

```
typedef struct
{
    unsigned SampleCount;
    unsigned CycleCount;
    unsigned FrameCount;
} s_EventDuration;
```

The `s_EventDuration` structure stores count values that are used to time various events in the synchronizer as described in Section 5.5. The `SampleCount` variable is associated with two of the `StopEvent` options listed in Table 5-6. The `CycleCount` variable is associated with one of the `CycleStop` options listed in Table 5-8. The `FrameCount` variable is used to define the length of a periodic frame that can be used as one of the `CycleSelect` options listed in Table 5-7. All of the variables will accept an input range from 0 to $2^{32}-1$.

5.8 IF Data Packet Format

```
typedef struct
{
    unsigned PacketSize;
    unsigned TSFSelect;
    unsigned TSISelect;
    unsigned TrailerSelect;
```

```

unsigned ClassIDSelect;
unsigned ClassID;
unsigned StreamIDSelect;
unsigned StreamID;
unsigned ContextPktCntEnable;
unsigned IndicatorEnables;
unsigned SpectralInversion;
unsigned UserIndicators[];
} s_IFdpktFormat;

```

The `s_IFdpktFormat` structure selects optional information to be included in the VRT IF Data Packet. Only RX channels support packetized data.

The `PacketSize` variable conveys the total number of 32-bit words necessary to form the complete IF Data Packet, not just the number of payload words. The value must include the header, payload, and all optional fields. A value of zero will result in a continuous stream of payload fields with no actual packet information. This option is not allowed if payload packing is enabled (`PayloadPacking = 1`). Setting the variable to one is impractical since the packet would only contain a header and no payload data. The maximum `PacketSize` value is $2^{16}-1$, which is a limit imposed by the number of bits allocated to the header field as discussed in Section 3.1.



The following restrictions are imposed on the Model 276:

- (a) `PacketSize` must include an even number of payloads when `ItemPackingFieldSize > 8`.
- (b) When payload packing is enabled (`PayloadPacking = 1`), the number of payloads included in `PacketSize` must be a multiple of ten when `ItemPackingFieldSize = 10` and a multiple of six when `ItemPackingFieldSize = 12`.

The `TSFSelect` and `TSISelect` variables determine whether the fractional and integer seconds components of the timestamp are included in the packet, respectively. Setting the variable to zero suppresses the designated field while setting to one inserts the field. These values are not equivalent to the code displayed in the TSF and TSI header fields described in Section 3.1.

The `TrailerSelect` variable determines whether a trailer is included in the packet. Setting the variable to zero suppresses the trailer while setting to one inserts the field. All of the bits contained in the trailer are assigned to optional fields as discussed in Section 3.5. The hardware does not generate context packets, so the `ContextPktCntEnable` variable is ignored and the lower eight bits of the trailer are always zero. The remaining bits are assigned to the State and Event Indicators listed in Table 3-2. There are two indicators produced by Red Rapids hardware; ADC over-range (bit 13) and channel synchronization (bit 11). The remaining User Defined Indicators (bits 10:8) can be set to zero or one through the `UserIndicators[]` array. The Spectral Inversion Indicator (bit 14) can be set to zero or one through the `SpectralInversion` variable. The `IndicatorEnables` variable determines which of the indicator bits should be masked as explained in Section 3.5.

The `ClassIDSelect` variable determines whether the value stored in the `ClassID` variable is included in the packet. Setting the variable to zero suppresses the Class Identifier while

setting to one inserts the field. Refer to Section 3.2 for additional information about the Class Identifier.

The `StreamIDSelect` variable determines whether the value stored in the `StreamID` variable is included in the packet. Setting the variable to zero suppresses the Stream Identifier while setting to one inserts the field. Refer to Section 3.2 for additional information about the Stream ID.

5.9 DMA Buffer

```
typedef struct
{
    unsigned PageCount;
    unsigned PagesPerMark;
    unsigned BurstSizeB;
    unsigned BurstCount;
} s_DMABuffer;
```

The `s_DMABuffer` structure sets the operating parameters for DMA transfers. Refer to the DMA on Demand Operating Guide (REF-004-000-Rxx) for a detailed description of each variable.

The `PageCount` variable sets the number of pages available in the buffer.

The `PagesPerMark` variable sets the number of pages between marker flags.

The `BurstSizeB` variable sets the size of each DMA burst in bytes.

The `BurstCount` variable sets the number of bursts in a page.