Multichannel Downconverter Cores for FPGA

Product Information Sheet

RF Engines Ltd (RFEL) has developed high specification FPGA-based downconverter channelisers for a range of different applications. The cores are designed to extract a number of signals from a wideband input source, filter and downconvert to baseband. The range of cores available includes fixed configuration cores that are optimal for a specific fixed requirement and configurable cores that provide real-time control of channel parameters such as centre frequency, bandwidth and output rate.

Being FPGA-based, RFEL's existing reference designs can be tailored to meet particular processing and interface requirements, such as input/output format, number of channels, dynamic range and filter shape.

The designs are based on novel frequency transform techniques that allow downconversion of a large number of channels in a highly efficient manner, reducing system cost and complexity. The designs are sold as Intellectual Property (IP) cores to equipment manufacturers and allow a reduced risk route and a faster design cycle.

The adjacent figure shows the general block diagram for downconversion of N channels from a wideband input.

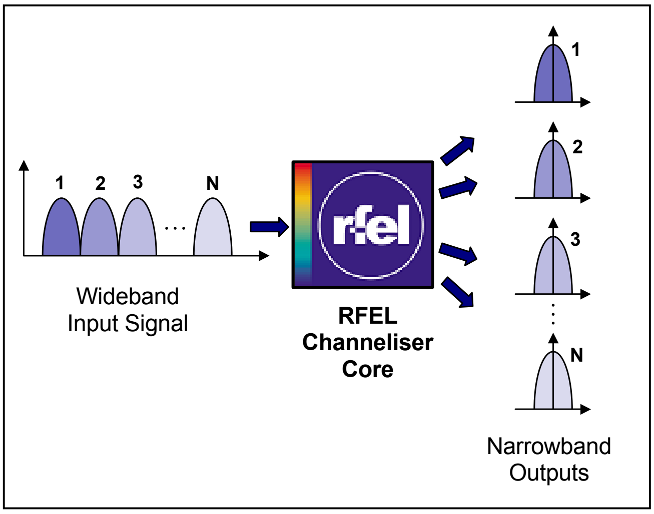

Figure 1 : Multichannel Downconversion

Key Features

- Downconversion of 1000s of channels in a single FPGA.
- Wide input bandwidths supported, up to GHz
- Use of silicon efficient frequency transform techniques for channelisation.
- Bit true Matlab models allow simulation of the channeliser, reducing the development risks.
- Xilinx and Altera FPGA devices supported
- Designs are customised for each application to minimise silicon usage

Key Applications

- Telecommunications base stations
- Satellite communication systems
- Software defined radio
- Military communications systems
- Radio monitoring systems

Overview of the Range

The channeliser range includes the following classes of core:

1) **Fixed multichannel downconverter cores** that channelise a very large number (up to thousands) of channels from fixed channel locations, where the channels have a fixed spacing and all share the same filter shape and output sample rate. (See Page 3)
2) **Reconfigurable multichannel downconverter cores** providing channelisation of a large number (up to hundreds) of relatively narrowband channels from a wideband input. These channeliser cores can be used to efficiently extract channels from any dynamically selectable frequency with a very wide variety of channel sample rates and filter characteristics/bandwidths. (See Page 6)
3) **Wideband DDC cores** providing downconversion of a few relatively wideband channels from a wideband input. (See Page 9)

For each application, the appropriate class may be selected to provide the most silicon-efficient solution; Figure 2 summarises the applicability for the different classes. Wideband DDC cores provide significant flexibility for a limited number of channels, whereas fixed downconverter cores provide limited flexibility for a large channel capacity. The flexible downconverter cores occupy the parameter space between; a core can be tailored to meet the required flexibility versus efficiency trade-off for a particular application. In general, silicon usage increases both with the number of channels and the required flexibility.

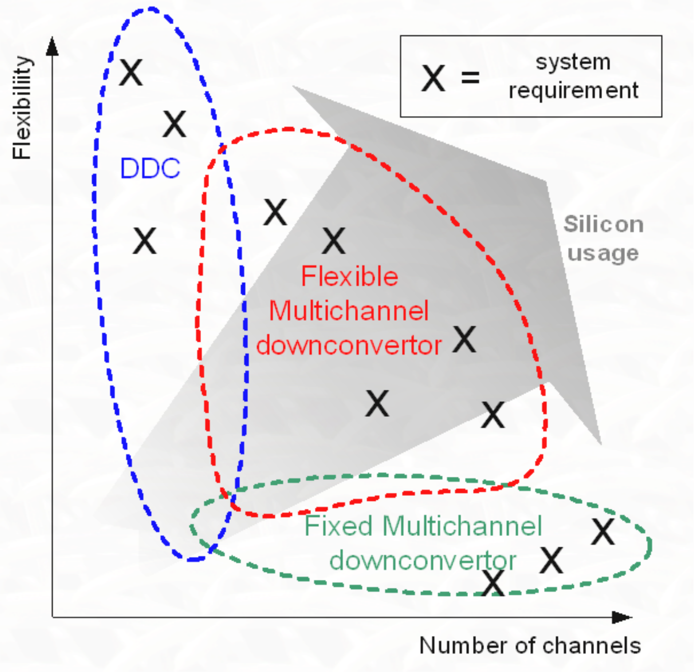

Figure 2: Suitability of channeliser architectures

For more information about the architectures and techniques for multi-channel downconversion, please refer to the paper “Comparison of Wideband Channelisation Architectures”, John Lillington, RF Engines Limited. This paper is available from RF Engines website (www.rfel.com).

Fixed multichannel downconverters

The structure of a typical fixed multichannel downconverter is shown in Figure 3. The input data stream may be either real or complex. The figure shows the structure for real input data (typically an ADC output), where the processes applied are as follows:

1) DHBF (Distributed Halfband Filter): efficiently converts the real input to complex, and centres the spectrum on DC. This includes filtering and decimate-by-2 processes.
2) Frequency offset correction: A complex frequency shift is applied that removes a systematic frequency offset across all channels (for instance, due to Doppler shift). The frequency shift to be applied is dynamically controllable via the core interface.
3) Parallel filter bank: extracts a large number of equally spaced channels from the input; all channels are identically filtered and have the same output sample rate.

Note that both the DHBF and the frequency offset correction are optional components that may be removed from the design if they are not required.

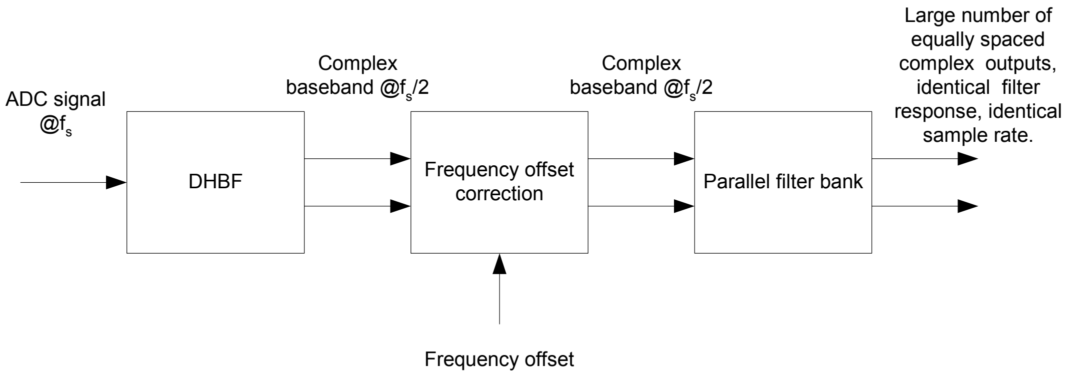

Figure 3: Typical fixed multichannel downconverter structure

Main features

- Downconversion of a very large number of channels (typically 100s to 1000s).
- Single digitised (real or complex) data input - input bandwidths up to 160MHz (wider bandwidths can be provided - up to GHz).
- Optional DHBF to accommodate real data input.
- Optional "systematic" frequency offset correction applied across all channels.
- Equal channel spacing, and fixed frequency spacing between channels.
- Identical filtering on each channel (filter shape customisable at design time). Designs may also be provided with run-time loadable filter coefficients.
- Identical output sample rate for all channels.
- SFDR of 90dB (higher SFDRs can be provided).

Main benefits

- Designs are more efficient than a stack of DDCs, typically when the number of channels is greater than 8.
- Parameterised HDL design may be modified by RFEL to meet specific requirements.
- Fully tested Netlist/Bitstream delivery.
- Fixed point Matlab models provided to de-risk core integration process.
- Reduced technical and timescale risk to project.

Required parameters

The following parameters are required to enable RFEL to provide a silicon estimate and costing for a given design:

Channel spacing:	The frequency spacing between adjacent channels.
Output rate:	The required output sample rate, including any tolerance.
Filter shape:	The filter shape to be applied to each of the channels. This may be specified either as a set of taps (designed at the input sample rate) or as a set of filter parameters (e.g. RRC with cutoff frequency and excess bandwidth specification). In practice, the filter length is all that is required to provide a silicon estimate.
Number of channels:	The number of channels to be extracted from the input signal.
Input sample rate/IF:	If either the input sample rate or the IF is fixed then this must be specified at design time.
Input format:	The precision and format of the input signal. If the signal is coming from a known ADC, then the part number is helpful.
Output format:	The required output precision and format.
SFDR:	The required spurious free dynamic range (SFDR). Typical SFDRs are in the range of 70-90dB.
Frequency resolution:	The required frequency resolution. This is used to determine the precision required for the frequency offset compensator.
Device requirements:	Intended target device/family and any other constraints.

This class of design is highly silicon efficient, because it uses information about the signal configuration to remove redundancy in the processing architecture. Therefore, where parameters are flexible it is helpful to know this. For instance, it is possible that a change in the input sample rate would provide a much more silicon efficient implementation for a given channel spacing/output sample rate combination.

Example applications

In one example application, a fixed channeliser was designed to extract 200 GSM channel outputs. For GSM, the channel spacing is 200kHz and for this application, the required sample rate was twice the GSM symbol rate of 1625/6 = 270.833kHz. The filter shape was chosen to meet the GSM receiver specification. The solution to this problem is as shown in Figure 4. RF Engines have generated a design for this channeliser and have shown that for a 14-bit ADC, it will fit comfortably on about 50% of a Xilinx Virtex-II 6000 FPGA.

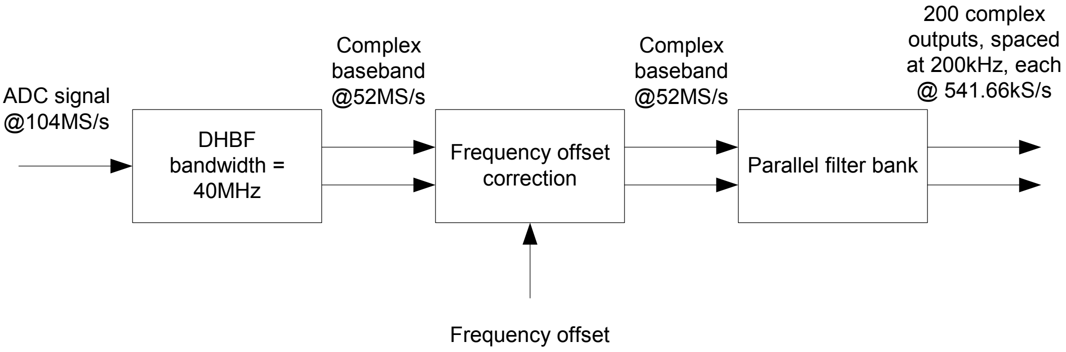

Figure 4: Solution for 200 channel GSM channeliser

In another recent application the requirement was to extract more than 1500 channels from an input bandwidth in excess of 40MHz with a given output sample rate and channel spacing. This design fitted comfortably within a Xilinx Virtex Pro50 FPGA.

Reconfigurable multichannel downconverters

The second class of multichannel downconverter channelisers provides an alternative trade-off between flexibility and channel capacity. In this case some of the flexibility of a DDC is removed in order to provide a more silicon efficient solution to meet a particular requirement. The downconverter is shown in Figure 5. The example shown is for a real input signal, with the complex output data provided in a time-multiplexed form, though alternative variants can be provided.

Parameters for each output channel can be programmed individually via the control interface. Generally, the channel centre frequency, output rate, filter selection and gain are programmable separately for each channel, whereas the Nyquist selection and Spectral flip option apply to the ADC data so a single control is provided that applies to all channels.

The channel filter shape is typically defined relative to the output rate and this filter is applied to all output channels, scaled to the required output rate. As shown in the figure, it is possible to provide multiple filter shapes to accommodate different requirements of the core.

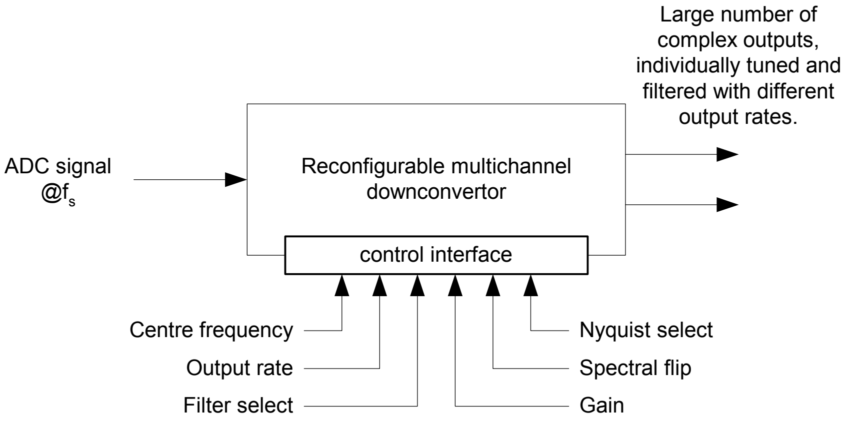

Figure 5: Reconfigurable multichannel downconverter

Main features

- Downconversion of a large number (up to 100s) of different channels.
- Independent control of:
  - Centre frequency;
  - Output rate;
  - Filter selection;
  - Gain (adjustable in 6dB increments).
- Real or complex data inputs.
- All channels independently reconfigurable during runtime (no interruption of existing data outputs).
- Multiple selectable filter shapes may be provided.
- Multiple ADC inputs can be accommodated in a single design.
- SFDR of 90dB (higher SFDRs can be provided).

Main benefits

- Highly optimised design to provide maximum channel capacity/minimum silicon for required flexibility.
- Parameterised HDL design may be modified by RFEL to meet specific requirements.
- Fully tested Netlist/Bitstream delivery.
- Fixed point Matlab models provided to de-risk core integration process.
- Reduced technical and timescale risk to project.

Required parameters

The following parameters are required to enable RFEL to provide a silicon estimate and costing for a given design:

Number of channels:	The number of channels to be extracted from the input signal.
Bandwidths/rates:	Required output channel bandwidths and corresponding sample rates, including any tolerances. The maximum output channel width is particularly significant in dictating the silicon usage.
Number of inputs:	Number of ADC inputs to the core.
Input bandwidth/ sample rate:	Bandwidth of input signal and/or input sample rate.
Filter shape(s):	The shapes of the filters to be applied to the output channels. One or more filter shapes can be specified either as a spectral mask or as a set of filter parameters (e.g. RRC with cutoff frequency and excess bandwidth specification). The filter shape should be specified with respect to the output sample rate (see for example Figure 6).
Frequency resolution:	The required centre frequency resolution.
Input format:	The precision and format of the input signal. If the signal is coming from a known ADC, then the part number is helpful.
Output format:	The required output precision and format.
SFDR:	The required spurious free dynamic range (SFDR). Typical SFDRs are in the range of 70-90dB.
Device requirements:	Intended target device/family and any other constraints.

Again, information about the signal configuration can be exploited to increase the efficiency of the design. Therefore, any additional information (e.g. valid channel/rate/filter combinations, any flexibility in parameter values) is useful in optimising the design.

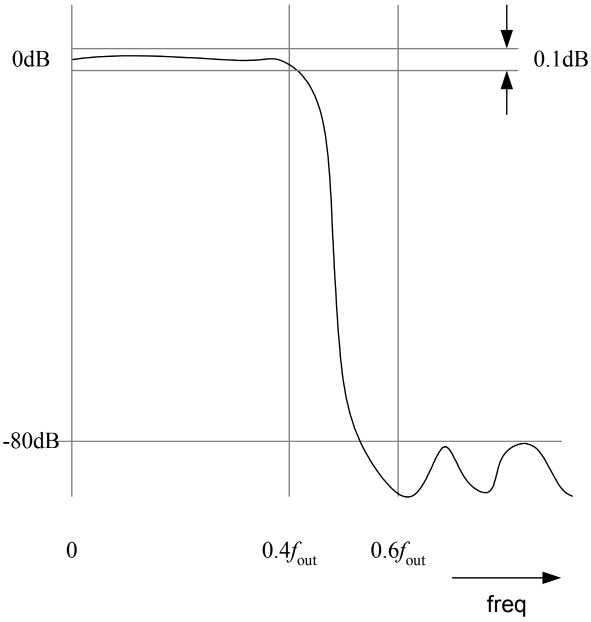

Figure 6: Example filter shape defined relative to output sample rate

Example applications

In one application, shown in Figure 7, 128 channels were required to be downconverted from a real input signal sampled at around 100MHz, with an input bandwidth of 40MHz. Two channel options were defined, where an option comprises a channel bandwidth and output sample rate, such that each of the 128 output channels could be selected to either option. The resulting core requires about 50% of a Xilinx XC2V6000 FPGA.

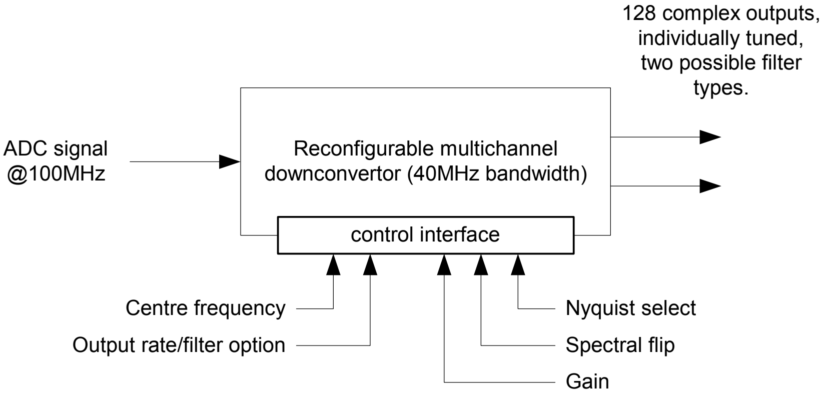

Figure 7: Reconfigurable multichannel downconverter example 1

In another example, a core was generated that would extract 64 channels from an input bandwidth of around 70MHz. In this case, the core was designed to have two separate IF inputs. The 64 output channels could be chosen from either of the two IF inputs with each channel having any of 10 different bandwidth/rate combinations, ranging from a few kHz to a few tens of kHz. A single set of filter parameters were used for all channels, so that the signal bandwidth was proportional to the output sample rate. The resultant FPGA design fits comfortably onto a single Xilinx Virtex-II Pro 30 FPGA.

Wideband DDC cores

General-purpose DDC cores are available from FPGA and board manufacturers, so there is no requirement for RFEL to replicate this functionality. For instance, the Xilinx LogiCORE DDC accepts real signals sampled at rates up to 100MHz and provide integer decimation rates from 4 to 1048512.

However, for specialist applications, RFEL have the capability to provide customised DDC cores, running at input sample rates into the GHz, with selectable filters and non-integer sample rate conversion factors.

Main features

- Input/output sample rates from Hz to GHz.
- Non-integer decimation rates (fractional resampling) - rates programmable at run-time.
- Multiple-channel DDC cores exploit commonalities to provide more silicon-efficient structure without compromising flexibility[1].
- All channels independently reconfigurable at runtime (no interruption of existing data outputs).
- Multiple selectable filter shapes may be provided.
- Multiple ADC inputs can be accommodated in a single design.

Main benefits

- Highly optimised designs provide minimum silicon usage to meet required specification.
- Re-use of RFEL's existing tried-and-tested IP blocks.
- Fully tested Netlist/Bitstream delivery.
- Fixed point Matlab models provided to de-risk core integration process.

Required parameters

Number of channels:	The number of channels to be extracted from the input signal.
Bandwidths/rates:	Required output channel bandwidths and corresponding sample rates, including any tolerances.
Number of inputs:	Number of ADC inputs to the core.
Input bandwidth/ sample rate:	Bandwidth of input signal and/or input sample rate.
Filter shape(s):	The shapes of the filters to be applied to the output channels.
Frequency resolution:	The required centre frequency resolution.
Input format:	The precision and format of the input signal. If the signal is coming from a known ADC, then the part number is helpful.
Output format:	The required output precision and format.
SFDR:	The required spurious free dynamic range (SFDR). Typical SFDRs are in the range of 70-90dB.
Device requirements:	Intended target device/family and any other constraints.

---

[1] Commonalities can be used to provide more efficient DDC cores for a relatively small number of channels. A larger number of channels can be accommodated using one of RFEL's multichannel downconverter cores described above.

rfengines limited

Innovation Centre
St Cross Business Park
Newport, Isle of Wight
PO30 5WB
United Kingdom

Tel: +44 1983 550330
Fax: +44 1983 550340
Email: info@rfel.com
Website: www.rfel.com