



November 2002

Gboard Feedback Processing System - Design Study and Proposal

L. Beckman, J. Fox, M. Tobiyama, D. Teytelman



Outline

- 1 Overview of Processing requirements for transverse and longitudinal processing
- 2 History of downsampled longitudinal processing development - why downsample? Past design decisions of 1992 - why not do digital transverse processing?
- 3 Technology options in 2002 - can we revisit non-downsampled architectures?
- 4 Gboard architecture and uneven stepping scheme
- 5 Detailed Gboard data flow, simulation results
- 6 Detailed design and implementation issues
- 7 Ideas for a collaborative future



Processing Requirements

For instability control, the processing channel must

- extract information at the appropriate synchrotron or betatron frequency,
- amplify it (a net loop gain must be generated, large enough to cause net damping for a given impedance)
- generate an output signal at an appropriate phase (nominally 90 degrees, but arbitrary if the system and cable delays, pickup and kicker locations are considered)

Some technical issues

- Bandwidth/sampling rate
- DC offset removal from the processing channel (e.g. from DC synchronous phase position, or static orbit offset)
- Saturation on large input errors
- Noise in the input channel (e.g. bandwidth reduction via processing filter)
- Maximum supportable gain
- Diagnostics (processing system and beam dynamics)



History of the LFB development collaboration, and important technical decisions

We examined a mix of all-analog, hybrid analog/digital, and all digital processing architectures in the original collaboration and study in the early 1990's. At that time, we decided:

Longitudinal processing was best addressed via a downsampled digital processing channel (where the downsampling better matched the synchrotron frequency to a sampling rate of $1/N$ revolutions). The general-purpose processing channel was implemented in an array of 80 commercial fixed-instruction DSPs (the "farm").

Transverse processing required a full-rate digital processing channel, and we could not see a technical means to implement a fully-programmable filter at the 500 MHz rate. The transverse systems were designed using a two-pickup front end, where the two pickups were separated in betatron phase - the correction signal was computed via a scaled sum of the two pickups, delayed by 1 turn. PEP-II used a digital delay mechanism - the ALS and smaller rings used analog delay cables.

KEK-B implemented a non-downsampled full rate digital filter, using a two-tap filter design suggested by Flemming Pedersen. This approach only required addition, not multiplication, though the filter characteristics (limited by DC offset constraints), and control of the output signal phase, are not as general purpose as a true FIR or IIR filter structure. The KEK implementation used full-custom GaAs circuitry to implement an 32 fold demultiplexer/multiplexer channel, with 32 filter channels. This architecture requires a harmonic number divisible by 32. In practice the sensitivity of the two-tap filter to changes in machine phase, and the total filter group delay require care in the operation of the KEK-B transverse system.

2002 Technology options

Do we really need 1 GSPS?

We think we really need **1.5 GSPS!**

Several machines in design right now are considering bunch repetition rates above 1 GHz.

- Photoinjected Energy Recovery Linac design at BNL is using TESLA superconducting 1.3 GHz RF cavities.
- IR ring at LBNL is considering 1-1.5 GHz RF.

Even for the 500 MHz and lower RF frequencies it is useful to be able to get two samples per bunch.

- I&Q detection using a single ADC
- Supporting a dual pickup transverse front-end

2002 technology supports very high-speed FPGA architectures, with special dedicated DSP functions

- High speed logic
- high speed multipliers as function block
- design tools for DSP functions



Specifications

Support bunch spacings down to 0.66 ns - sampling at 1.5 GHz.

Support arbitrary harmonic numbers (may be OK to support only even numbers?).

Independent processing for all bunches on all turns - required for transverse feedback.

Diagnostic memory capable of holding 20 ms of data at the full rate

Support downsampled processing - reuse the hardware to get longer filters

Support downsampling for diagnostics for studying slow events

Support long FIR or IIR filters

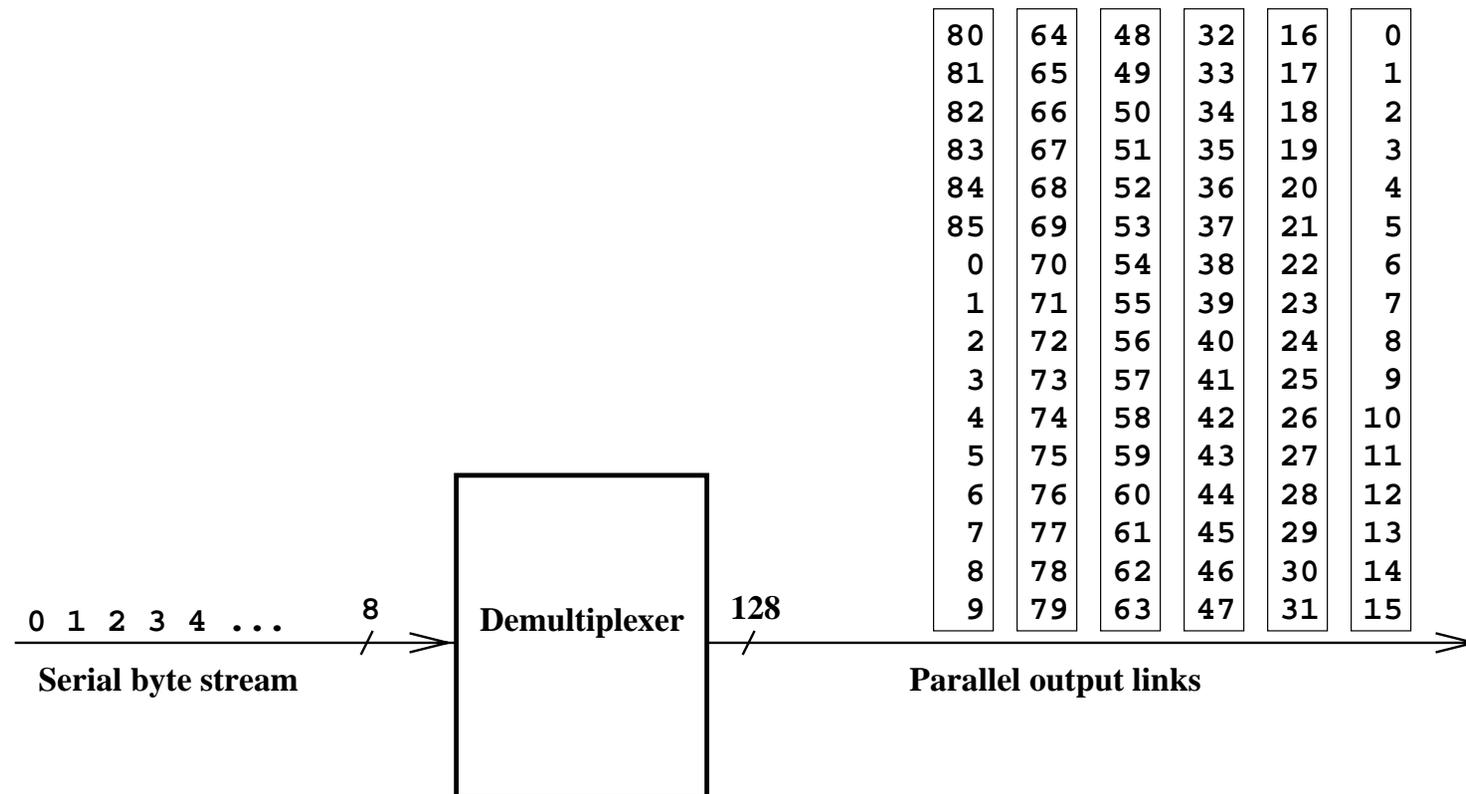
In longitudinal feedback non-downsampled processing allows one to better filter out the broadband noise, reduces loop delay somewhat. For example, the processing downsampling by 30 has 15 turn (0.5 sample) delay added to the filter group delay (90 turns for 6 tap filter) - a total of 105 turns. Running the same filter at full rate (180 taps!) the delay is only 90 turns. Of course, this extra complexity requires more computational resources.

Standard demultiplexing

Standard 1 \rightarrow N demultiplexing only works for the machines with harmonic numbers divisible by N .

An example of 1:16 demultiplexing, harmonic number is 86.

The ring is not closed in this case - signal for a given bunch is sent to different processing channels on consecutive turns.

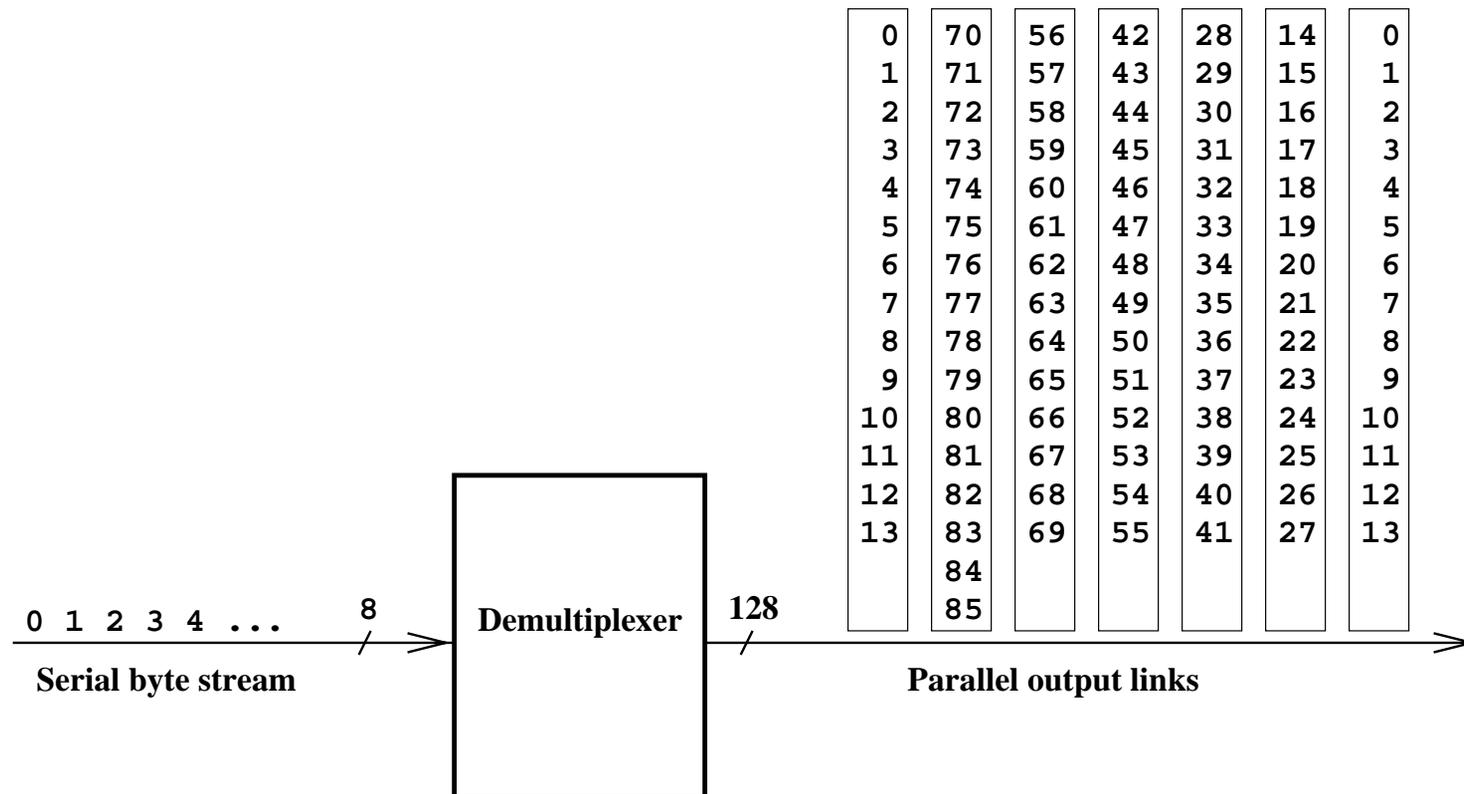


Uneven stepping

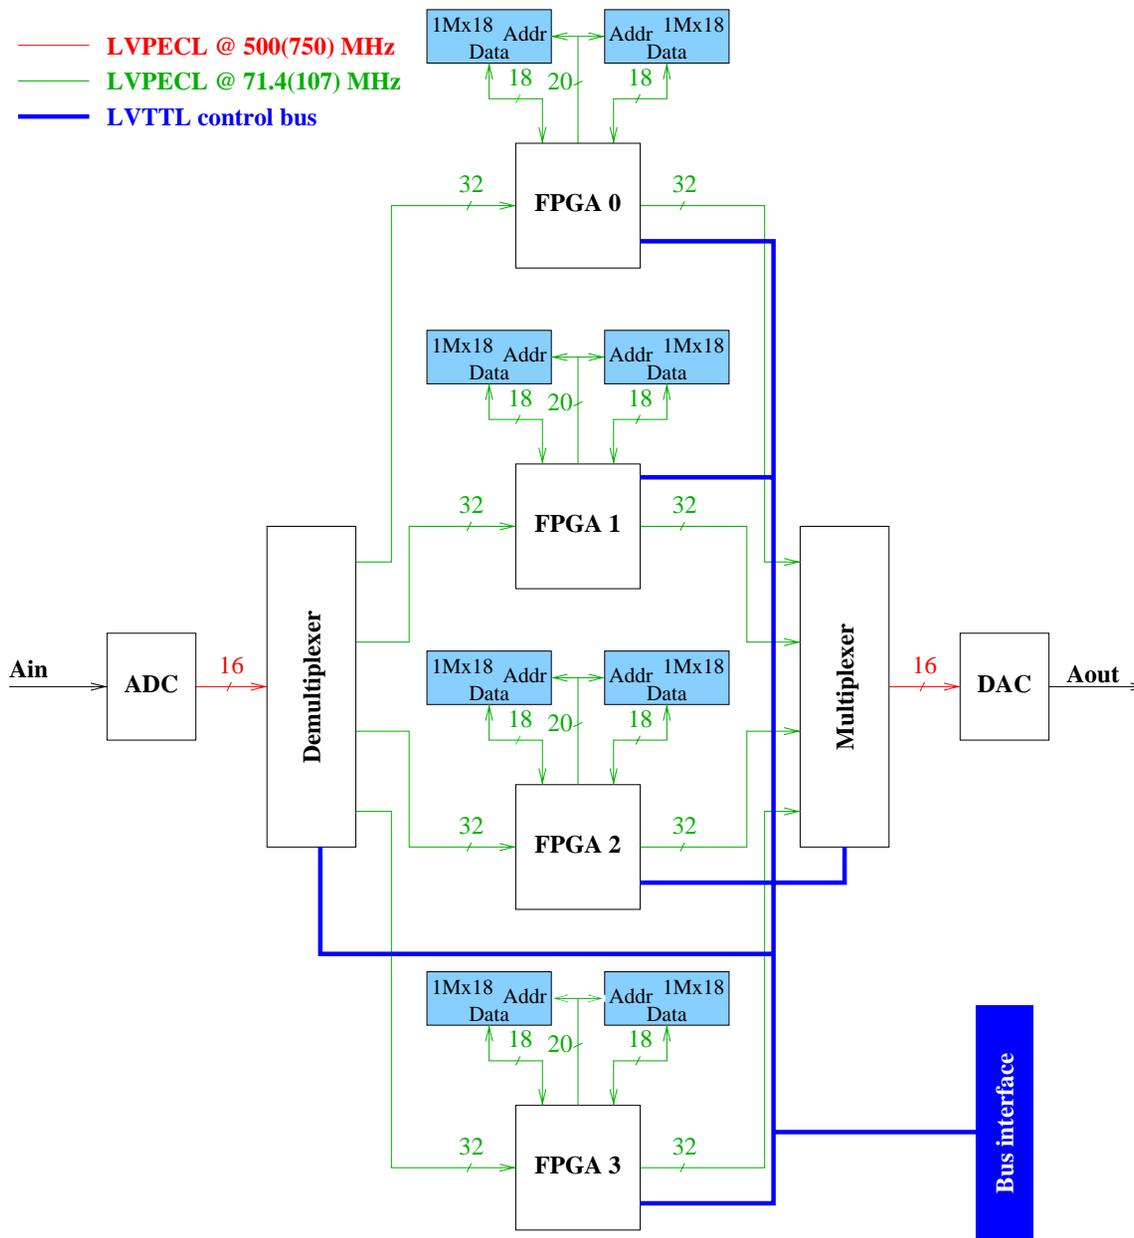
Uneven stepping - match the harmonic number with a combination of N and $N - 1$ wide transactions.

Any harmonic number larger than or equal to $(N - 1)(N - 2)$ can be matched.

Another option is to use N and $N - 2$ combinations. Then any even harmonic number starting from $(N - 2)(N - 4)/2$ can be matched. For 16/14 uneven stepping all even harmonic numbers $h \geq 84$ can be matched.



Baseband processing architecture



Hardware description

Baseband architecture with 1.5 GHz maximum processing rate implemented as a single VME64X module.

Data Flow Processing is implemented in 4 Xilinx Virtex-II FPGA devices.

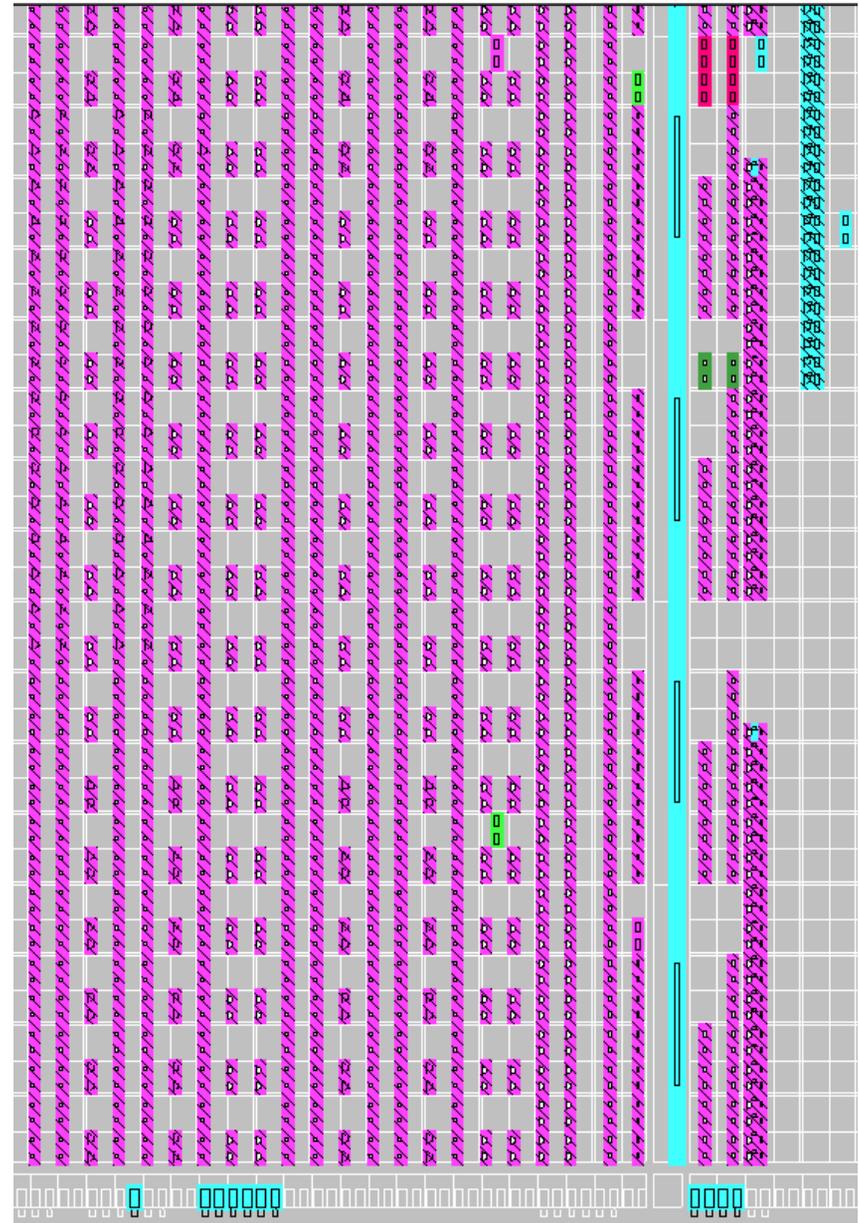
Each chip handles 4 data samples in parallel. With uneven stepping parallel stream alternates between 16 and 14 samples (94-107 MHz clock rates at 1.5 GSPS).

Use Xilinx Virtex-II FPGA XC2V8000

- 112 × 104 CLB array
- 3024 Kbits of RAM
- 168 18 × 18 multipliers - up to 210 MHz clock

Each FPGA controls two synchronous SRAMS of 512K x 36

System can acquire 7 x 2 x 1M=14M samples of transient data (worst case) - this corresponds to 14 ms data record at 1 GHz.





Example of PEP-II/KEK-B/DAΦNE processing

KEK-B and PEP-II are the most processing intensive machines.

Table illustrates processing loads and limitations of the FIR algorithm.

Table 1: PEP-II, KEK-B, and DAΦNE processing

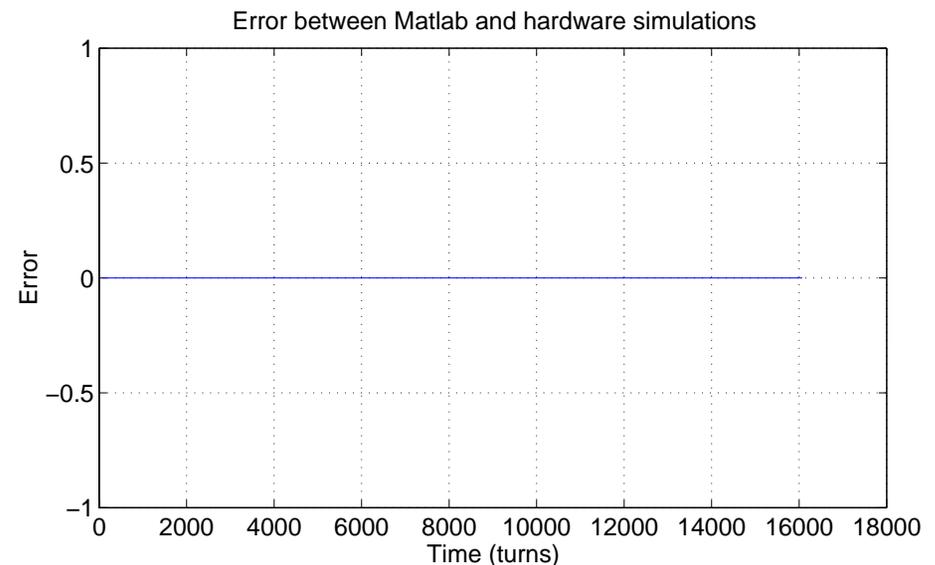
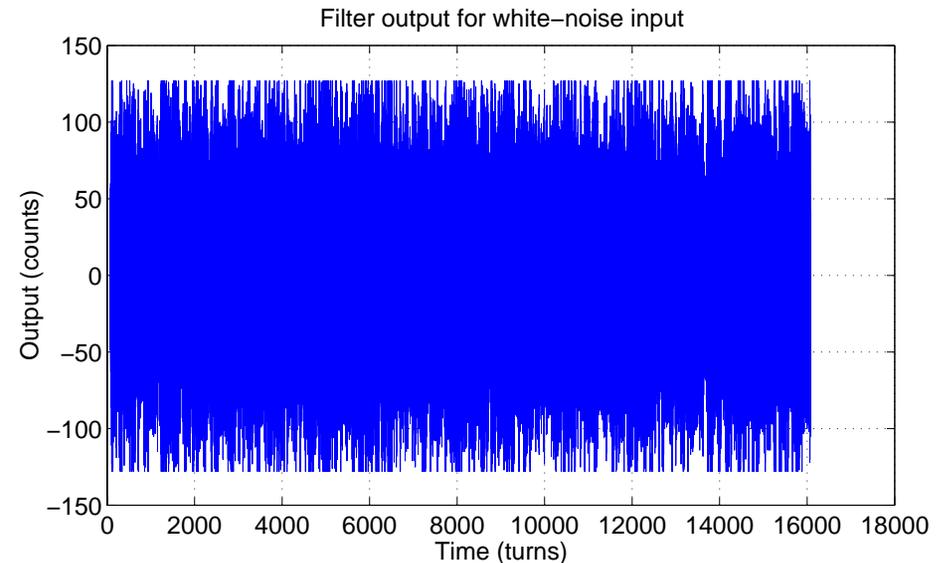
Parameter	PEP-II	KEK-B	DAΦNE
RF frequency, MHz	476	508.9	368
Harmonic number	3492	5120	120
Stepping selection	213@16, 6@14	320@16, 0@14	4@16, 4@14
Groups per turn	219	320	8
Multiplier limit on FIR filter taps	42	42	42
FIR processing rate for I&Q sampling, MHz	34	36.35	26.3
Full I&Q channel rate	68	72.7	52.6

In case of DAΦNE multipliers can be used several times per sample to make longer FIR filters

Functional simulation of a 20-tap FIR channel

A 20-tap FIR channel was designed for the FPGA implementation.

- 8 bit ADC data
- 8 bit DAC output
- 16 bit coefficients
- Full width accumulators (24-29 bits)
- Shift gain of 0-7 bits
- Output saturation to 8 bits
- FPGA resource usage - 13%
- Compiled implementation has 6.3 ns cycle time - for 1.5 GHz we need 9.3 ns.
- Functional simulation (Innoveda Fusion) for 8 groups (DAΦNE case) using white noise input signal
- Compared to bit-true MATLAB simulation

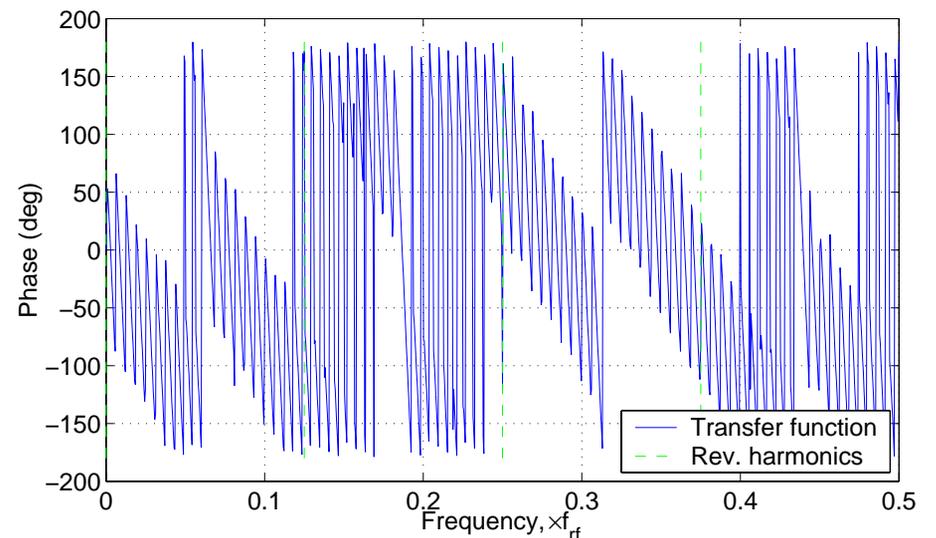
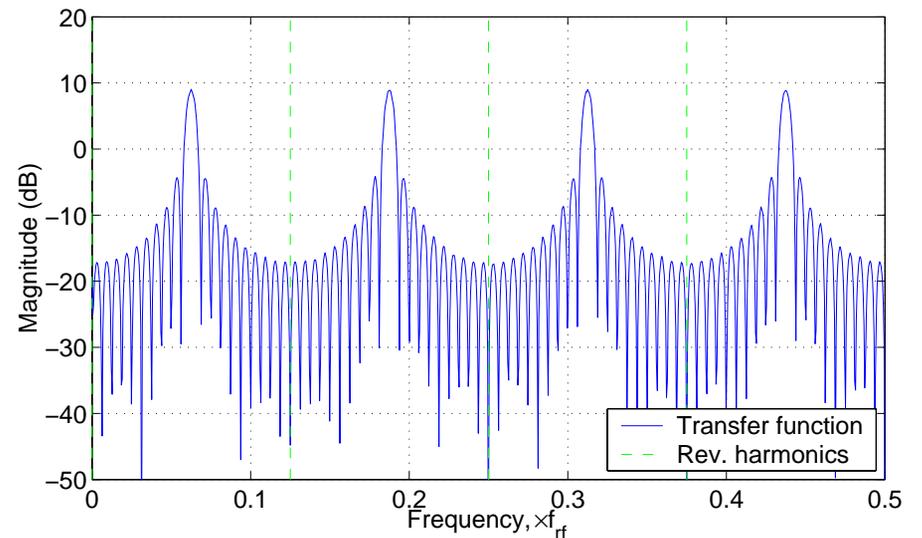


Transfer function estimate for the simulation output

Using input and output vectors we estimate the channel transfer function

Filter coefficients are chosen to alternate between 1 and -1 - peak response midway between revolution harmonics.

Periodicity of 4 corresponds to 8 bunches.



ADC to DAC functional simulation

Simulation includes

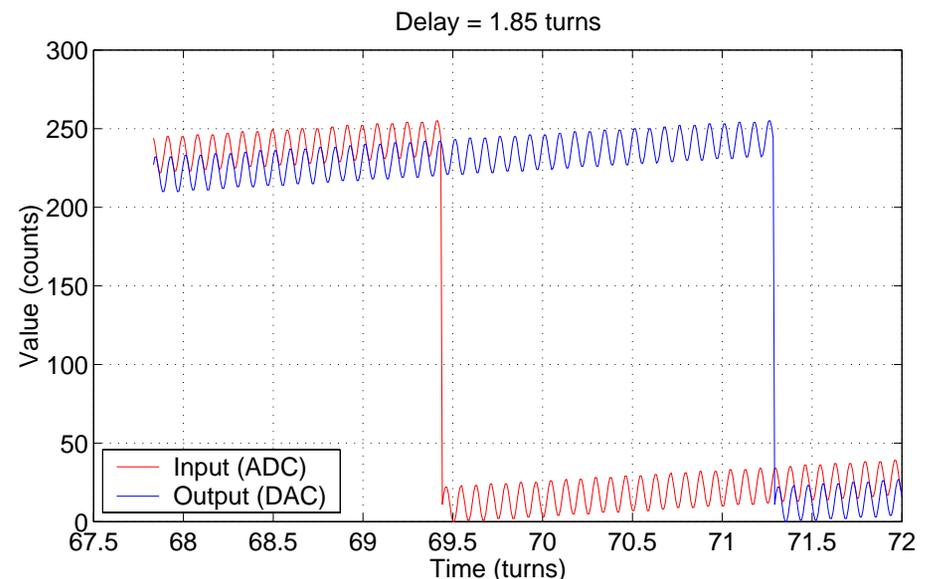
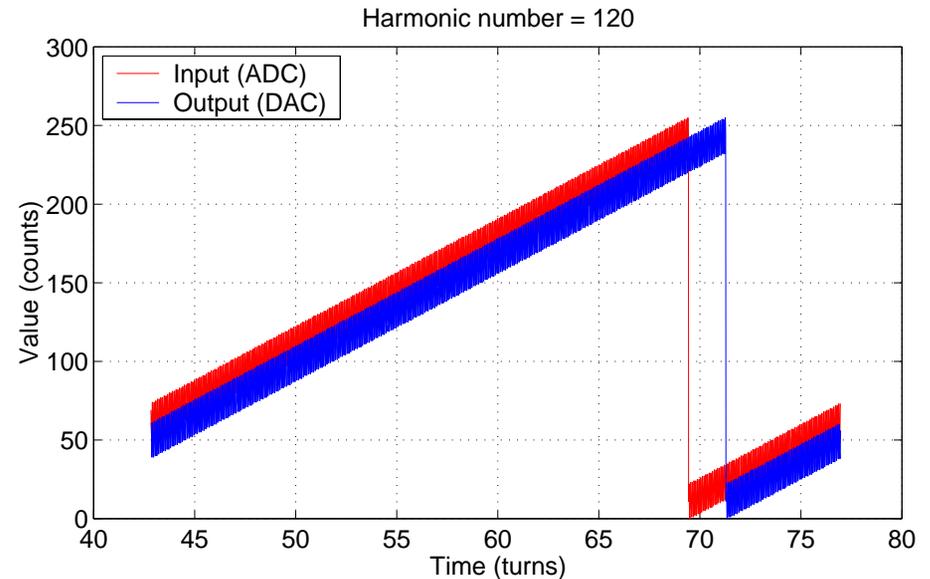
- 1 ADC model (MAX108)
- 2 Uneven stepping demultiplexer implemented in ECLinPS Plustm logic
- 3 4 Xilinx Virtex-II FPGAs
- 4 Uneven stepping multiplexer
- 5 DAC data stream generator

Simulation clocked at 1 GHz, 120 bunches per turn

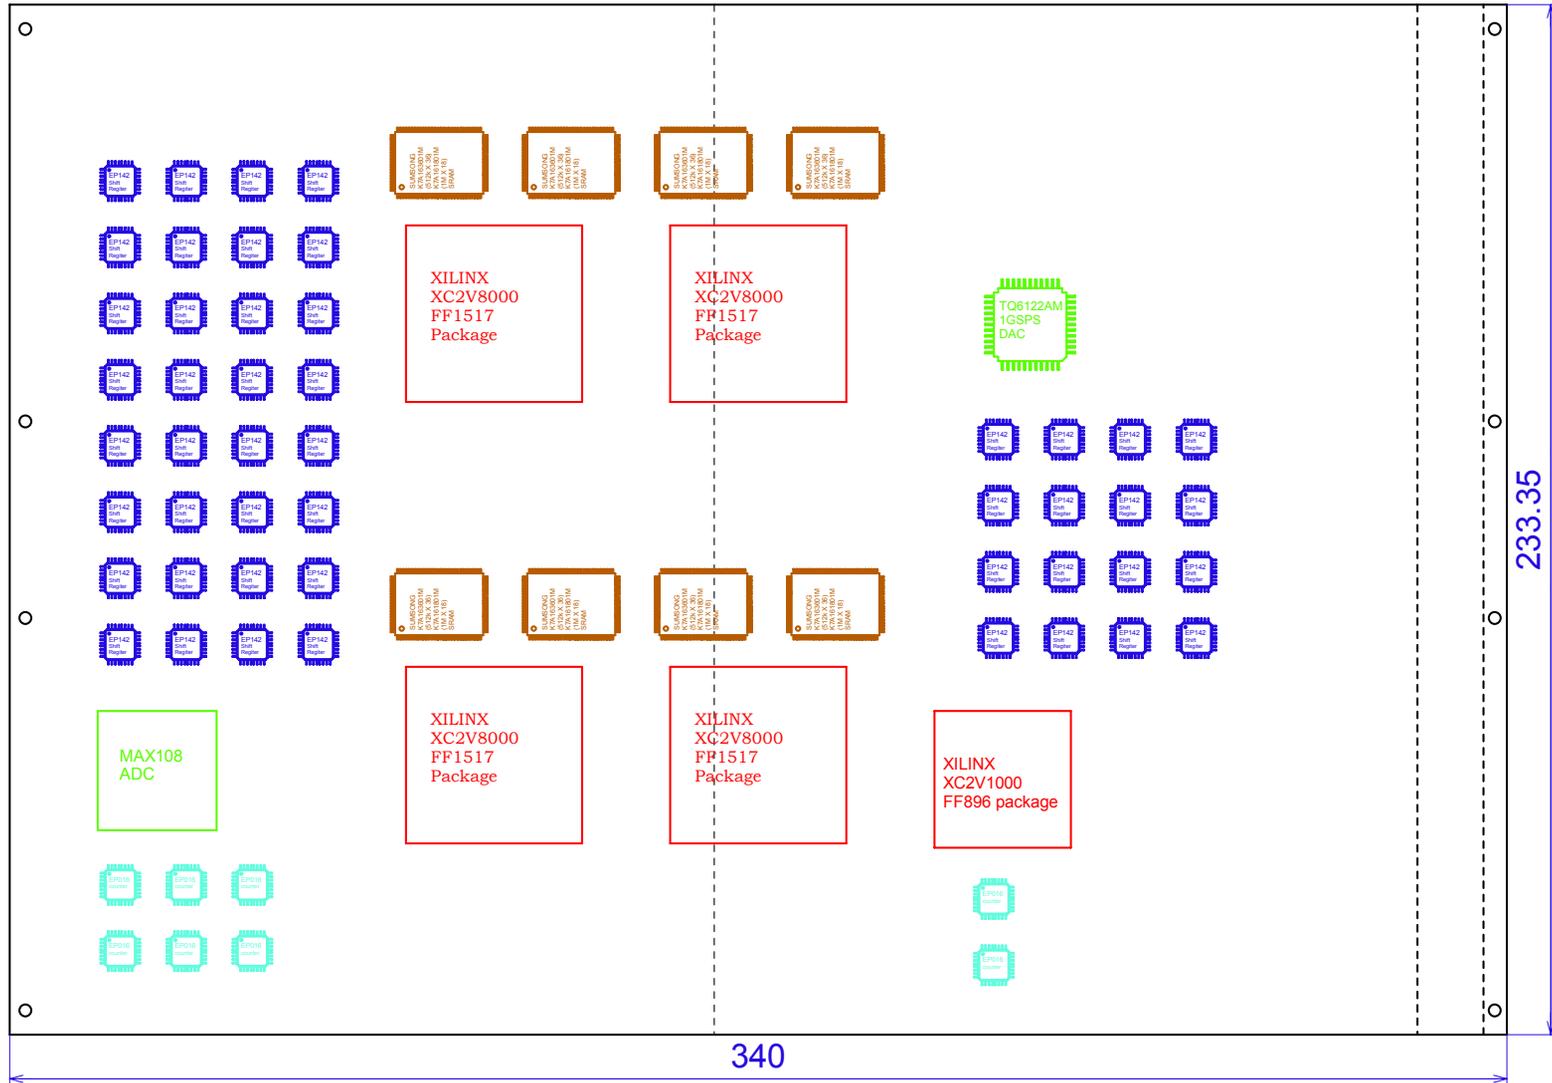
FIR coefficients set to pass-through (1,0,0,...)

Delay through the processing channel is 1.85 turns (222 clock cycles) is defined by the adjustable fiducial delay counter.

Minimum delay is around 176 clock cycles - if necessary can be reduced for DAΦNE by removing pipelining stages from FPGA. The pipelining is needed to get 6.3 ns processing time (1.5 GSPS) vs. 38 ns processing time required in DAΦNE (368 MSPS).



Board size estimate

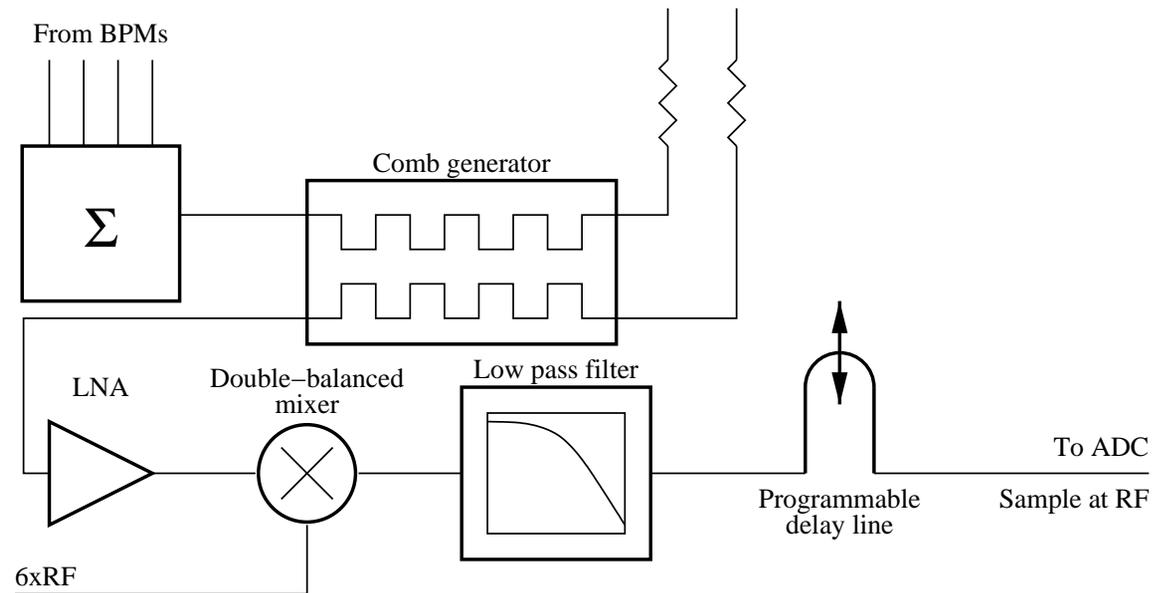


Analog front-end: current design

Here is the conventional analog front-end

Design Issues:

- Sensitivity to DC beam phase - need phase servo feedback loop
- Sensitivity to gap transients - loss of gain at the edges, feedback sign flip
- Bunch currents can only be determined indirectly



Is it better to have I&Q front-end?

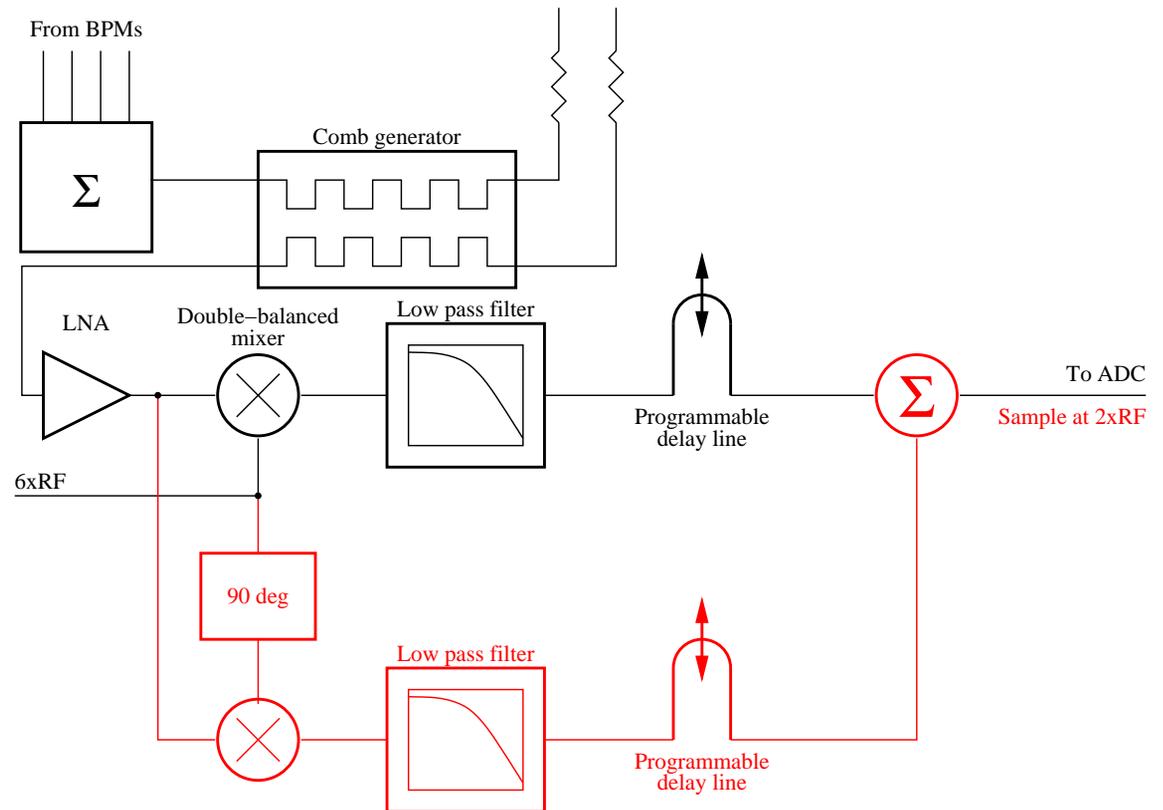
Analog front-end: brute force I&Q

Sampling twice per RF period we can get I&Q information for each bunch

Here is a brute force approach to I&Q front-end design

Problems of the brute-force approach:

- Gain errors in the two paths
- Phase error in the 90 degree splitter
- Requires precise timing adjustment of the two pulses

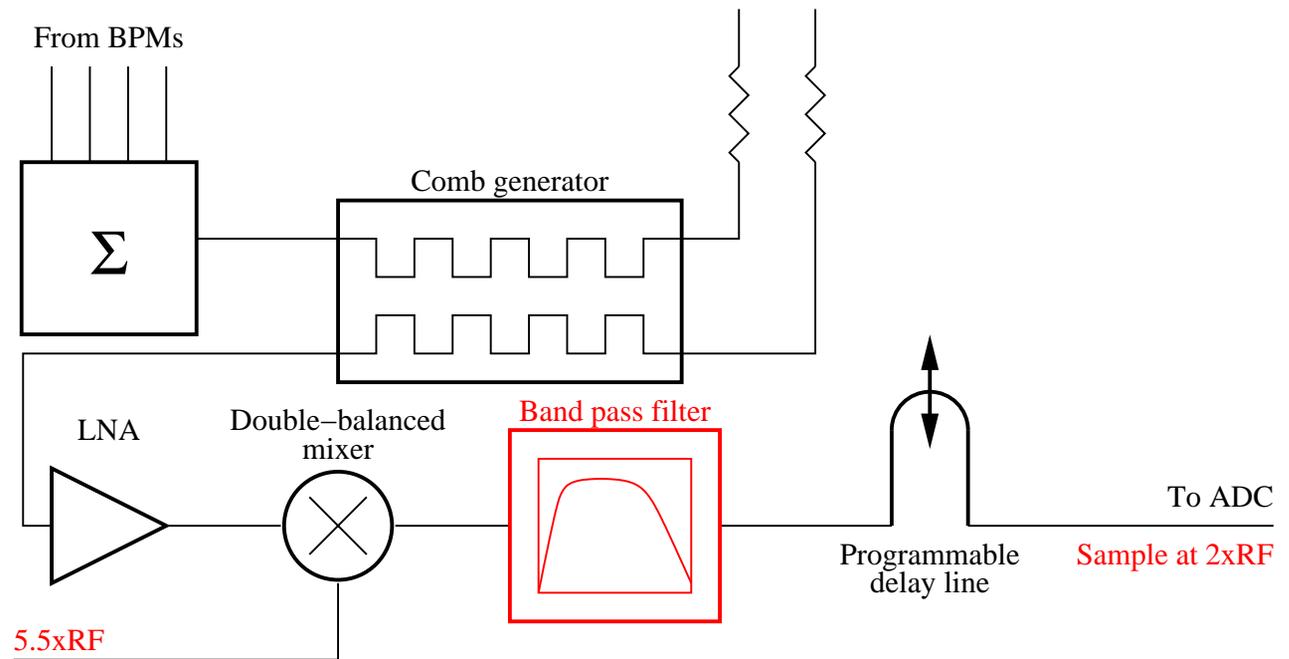


Analog front-end: IF I&Q

Here is an alternate way to design an I&Q front-end using a digital IF

This design places oscillation data in the $f_{rf}/2$ wide band around $f_{rf}/2$ (band from $f_{rf}/4$ to $3f_{rf}/4$).

Sampling that waveform at $2f_{rf}$ we get two quadrature samples per bunch.



Pluses and minuses of IF I&Q processing

Advantages

- Fast bunch-by-bunch current monitoring is trivial
- No I or Q path errors due to phase or gain asymmetries
- No DC at the ADC input
- Lower $1/f$ noise than baseband sampler
- Works for any gap transient
- No need for phase servo - less phase noise in the LO carrier
- DC beam phase shifts can be tracked
- Input FPGA samples (I&Q) are rotated at full processing rate (4 multiplies, 2 adds) to get $Gi_b \cos \phi_{ac}$ and $Gi_b \sin \phi_{ac}$
- The rotation angle is computed by a slow software process at 1-10 Hz.



Pluses and minuses of IF I&Q processing (cont'd)

Disadvantages

- Needs 2x sampling - won't work for $f_{rf} > 750$ MHz
- Requires tracking of synchronous (DC) angles and computation of the per-bunch rotation matrices.
- Need to reject image frequency at $5f_{rf}$
- Two consecutive samples must land in the same processing channel

The baseband processing architecture (from the A/D through the processing channel to the D/A) is the same hardware (with some firmware changes)- the issue is the RF and analog processing structure.

Lessons learned with the existing LFB, continued

The most obvious problems we've had with the existing LFB architecture, as well as suggestions to correct these problems in the new design.

Problem: Data conversion (the `gd_post` problem)

- Convert the data to Matlab-accessible format within the crate controller avoiding all synchronization problems.

Problem: Dataset management and labeling

- Each transient will be automatically entered into a database. User interface should allow one to enter comments for each dataset. After a timeout period (several weeks?) un-commented datasets are deleted.

Problem: Pre-trigger acquisition

- Include continuous acquisition with stop trigger from software or hardware in the new design

Lessons learned with the existing LFB

Problem: Loss of control in a grow/damp

- Automated feedback loop closure on the bunch-by-bunch basis. During the open-loop portion of the transient we still compute the correction signal. When this signal for a given bunch approaches some threshold (below DAC saturation) we close the loop on that bunch. The resulting dataset will have 3 sections: open-loop on all bunches, transition, closed-loop on all bunches.

Problem: Getting grow/damp statistics

- Often we need multiple grow-damps with fast rise/fall times. In the new system we should provide for the multi-transient diagnostics where the feedback loop is opened and closed several times in a row per data record. Then a single dataset can contain multiple fast grow/damp transients.

Problem: Gap transients

- Addressed by the I&Q processing

Problem: Timing by programmable delay lines

- ADC and DAC clocks in the new system will be adjustable over one RF period to position sampling clock and baseband output relative to the beam. Back-end carrier phase must be adjusted to place the bunch on the maximum. Then back-end timing sweep will produce a flat pulse rather than rectified carrier.



Detailed Design and Development Issues

Critical high-speed signal processing channel is verified via functional and timing verification - what's left to do?

Signal processing - add downsampling features to allow best use of the processing capability for longitudinal processing in large machines

Remaining tasks - the real detailed engineering

Control interface, user interface

High-speed timing and clock distribution design

physical layout, circuit PC board design, controlled impedance design

(the physical layout, delays, skews, etc. are NOT simulated in the board-level simulation)

Packaging format - VME64X? 400 mm depth? What supplies to use?

Component choices - largely made. Issue of use of Triquint D/A?

Board size issues - density, thermal management, use of BGA technology (limits useful board size)

Front panel - connectors? Monitor points or functions? Temperature monitoring? Shutdown?

Initial module prototype with baseband processing - use existing detection, front end and back-end functions as they are implemented. Later implement new VME64X functions?



Ideas for a Collaboration

We would very much like to do this processing system development and technology effort as a multi-lab collaboration. Our successes with the original LFB collaboration have taught us the numerous advantages.

Likely labs - SLAC, LNF-INFN, KEK, BESSY-II (?), APS(?), ALS(?)

We must plan a schedule and work plan - we would like to complete the detailed design this winter, and have functioning prototypes in fall 2003 to evaluate at one or more labs.

Best use of each labs' skills, resources, and people?

Economic issues -

This one baseband processing module effectively replaces the existing downsampler/holdbuffer and the VME crates of DSP boards (plus controllers, interface cards, etc.). We hope it won't cost as much as all these other functions did - but it's not going to be a cheap module. The Virtex-II FPGA devices alone on the card will cost \$25K.



Acknowledgements

We are grateful for the numerous contributions from A. Young (SLAC), J. Olsen (SLAC), R. Larsen (SLAC), S. Prabhakar (Stanford) and L. Sapozhnikov (SLAC).

We are appreciative of the longstanding interest and collaboration with the LNF-INFN team. We especially thank M. Serio, A. Drago, A. Ghigo, O. Coiro and P. Possanza for all the hospitality and friendly help over so many years. Our thanks also to F. Marcellini for his continuous and essential help with the overdamped cavity kicker effort at LNF and SLAC in the past year.