

EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH

CERN - AB Division

CERN-AB-2003-065 BDI

**DIGITAL SIGNAL PROCESSING IN BEAM INSTRUMENTATION:
LATEST TRENDS AND TYPICAL APPLICATIONS**

**M.E. Angoletta
CERN, 1211 Geneva 23, Switzerland**

Abstract

The last decade has seen major improvements in digital hardware, algorithms and software, which have trickled down to the Beam Instrumentation (BI) area. An advantageous transition is taking place towards systems with an ever-stronger digital presence. Digital systems are assembled by means of a rather small number of basic building blocks, with improved speed, precision, signal-to-noise ratio, dynamic range, flexibility, and accompanied by a range of powerful and user-friendly development tools. The paper reviews current digital BI trends, including using Digital Signal Processors, Field Programmable Gate Arrays, Digital Receivers and General Purpose Processors as well as some useful processing algorithms. Selected digital applications are illustrated on control/feedback and beam diagnostics.

Presented at DIPAC 2003, Mainz, Germany, 5 - 7th May 2003

*Geneva, Switzerland
19 June 2003*

DIGITAL SIGNAL PROCESSING IN BEAM INSTRUMENTATION: LATEST TRENDS AND TYPICAL APPLICATIONS

M. E. Angoletta, CERN, Geneva, Switzerland

Abstract

The last decade has seen major improvements in digital hardware, algorithms and software, which have trickled down to the Beam Instrumentation (BI) area. An advantageous transition is taking place towards systems with an ever-stronger digital presence. Digital systems are assembled by means of a rather small number of basic building blocks, with improved speed, precision, signal-to-noise ratio, dynamic range, flexibility, and accompanied by a range of powerful and user-friendly development tools. The paper reviews current digital BI trends, including using Digital Signal Processors, Field Programmable Gate Arrays, Digital Receivers and General Purpose Processors as well as some useful processing algorithms. Selected digital applications are illustrated on control/feedback and beam diagnostics.

INTRODUCTION

This review addresses latest years' developments in the use of digital techniques in BI. Owing to the large body of data available on this topic, the reader is also referred to the proceedings of various conferences and workshops such as the BIW, DIPAC, PAC, EPAC, ICALEPCS and LINAC, widely available on the web, and to previous reviews on some aspects of this subject [1,2,3].

A "Digital Engineering" perspective is adopted here, through the concept of "Digital Building Block" (DBB). This indicates the smallest digital unit, a chip, which can accomplish a given function. DBBs are now profusely used as part of recent digital systems. There are several DBBs, important for the evolution of the digital signal processing field. The treatment here is limited to only four major programmable DBBs, having great flexibility through programmability, providing multifunctionality hence causing continuing system size shrinkage.

The digital signal processing area has been steadily evolving since the '60s [4]. It is witnessing a slow transition from an all-analogue approach to an analogue/digital balance, primarily in control/feedback applications, with an ever-stronger digital presence.

DIGITAL BUILDING BLOCKS IN BI

Nowadays, digital systems are assembled using a limited number of "miniaturised" DBBs, each accomplishing a range of functions and having a definite place in the signal processing chain (Figure 1). The four DBB considered in this paper, listed as one follows the signal from the detector, are 1) Digital Receiver (DRX), 2) Field Programmable Gate Array (FPGA), 3) Digital

Signal Processor (DSP) and 4) General Purpose Processor (GPP) and GPP-hybrid. DBB major functionalities include baseband translation and initial data reduction (for DRX), fast math-intensive processing (for DSP and FPGA), glue logic (for FPGA), slow processing, data management and high-level interfacing (for GPP). In general, system specification and final goal will influence and modify this generalised chain structure. For example, either the DSP or the FPGA may be absent depending on the type of processing and the time constraints the system imposes. Typically, FPGAs support faster sampling than DSPs. On-site availability of tools and expertise in DSP or FPGA development may tip the chain balance towards either a no-DSP or a no-FPGA implementation. A DBB not covered here is for example the Analogue-Digital Converter (ADC). This will also influence system




Figure 1: DAQ and processing chain DBBs.

architecture and its evolution will impact digital signal processing in the future. For example, the analogue/digital divide in the system signal chain may move closer and closer to the detector end as ADCs become faster and more powerful. Conceivably, future systems will directly feed the signal to the digitising ADC, without going through intermediate frequency (IF) translation.

Last but not least, developers should provide a digital system with enough diagnostic and troubleshoot access points, to avoid turning the system into a tightly sealed black box.

Digital Receivers

The digital receiver chip, also called Digital Down Converter (DDC), is the evolution of the classical analogue superheterodyne receiver [1]. It extracts a narrow frequency band of interest from a wide-band input by down mixing, filtering and decimation. Figure 2 shows the block diagram for a generic DRX. Its main blocks are the local oscillator (LO), the complex (or I/Q) mixer and the low-pass filter. LO is a direct digital synthesiser generating a quadrature sinusoid, with programmable frequency and phase, from a look-up table. The mixer consists of two digital multipliers and outputs only the

sum and the difference frequency signals, unlike analogue mixers that generate many unwanted mixer products. The third block in Figure 2 performs filtering and decimation.

This kind of digital processing has many advantages. For example, the chip's behaviour can be changed by software, adding flexibility through programmability. Second, the chip is highly linear. Spurious signals are strongly rejected, based on the accuracy of the sinusoidal waves generated by the local oscillator and by the mathematical precision of the mixer. Such rejection is a difficult task if performed by analogue electronics.

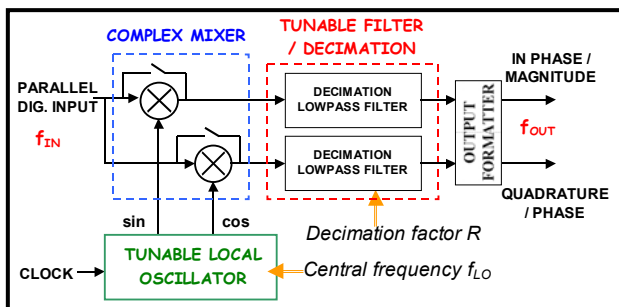


Figure 2: Typical DRX schematic layout, showing its 3 building blocks and two of the control parameters.

Finally the decimation action sensibly cuts down the number of data samples, allowing a subsequent data processing step. DRX are mostly used together with DSPs and/or FPGAs in BPMs [5] and beam parameters measurement systems [6, 7].

FPGAs

FPGAs are high-performance programmable logic devices, i.e. integrated circuits made of an array of logic cells that can be programmed to be interconnected to achieve different designs. FPGAs feature very high capacity (now in the 10^6 gate range) and flexibility. Altera and Xilinx are the two main FPGA producers, their more powerful families being Altera's Stratix and Xilinx' Virtex-II and Virtex-II Pro. FPGAs lead the way to System-On-a-Programmable-Chip (SOC) technology, which combines in a chip a large amount of programmable logic with memory, a processing engine and possibly additional Intellectual Property (IP) core. FPGAs are limited in that usually the floating-point logic is not easily implemented and its execution is slow. However, some core libraries are on the market now to address this need.

FPGAs are widely used in BI for two main functions: glue logic and very fast data processing. In fact, they support and sustain much faster sampling than afforded by DSPs. Examples of FPGA BI applications are the LHC orbit trajectory system [8], the SNS Front-End low level RF system (FELLRF) [9] and SNS BPMs [10]. In the SNS FELLRF, FPGAs are in charge of fast intra-pulse processing (pulse length is about 1 ms), while inter-pulse slower processing is left to the host processor. In SNS BPMs, FPGAs are again in charge of fast processing, with slower processing left to rack-mounted PCs.

Digital Signal Processors

DSPs first BI applications date back to the '80s [11]. DSPs are microprocessors tailored for fast, math-intensive tasks and may "service" sampling frequencies up to a few MHz. Their architecture, shaped by digital signal processing algorithms, include specialised units for fast execution of mathematical tasks and to keep numerical fidelity. Specialised I/O interfaces and handling schemes, such as Direct Memory Access, ensure fast data transfer without DSP main processor's intervening. The code's predictable execution time and the DSP's low-latency response to interrupts are enhanced by the frequent lack of Operating System (OS), as occurs in interrupt-driven systems. Conventional pre-mid-90s DSP architectures were based on separate data and address buses - the Harvard model [12]. Two main enhanced architectures have emerged since then: Very Long Instruction Word (VLIW) and Single Instruction Multiple Data (SIMD). They increased the execution speed by executing in parallel either different instructions or the same instruction on different data. VLIW is used in TI's 'C6x DSPs and SIMD in the Analog Devices ADSP-2116x family. VLIW and SIMD can coexist in the same DSP, like in Analog Device's TigerSHARC. With these architectures, DSPs are more efficient *but* optimising the associated Assembly code is more complex. This forces developers to write almost entirely C-code and to rely on compiler/optimiser performance. To a higher abstraction level, developers use automatic code-generation tools, such as Matlab Real-time Workshop Embedded Coder and Hypersignal RIDE, allowing code construction by drag-and-drop of pre-defined standard-function boxes.

DSPs are widely used in BI, alone or with other DBBs, typically FPGAs and/or DRX. Usage includes tune, intensity and emittance measurement [6,11]. Beam Position Monitors (BPM) also use DSPs: here the processing carried out is typically signal difference-over-sum plus offset subtraction [5,13]. DSPs occur in Beam Current and Beam Loss Monitors (BCM, BLM), either alone or preceded by a DRX, to perform a differential current loss measurement and often to interface to a fast machine-protection system [14,15]. DSPs roles in feedback/feedforward control systems are beam control, a complex task where beam dynamics is a crucial factor for the control loop design [16] and, less demanding, power converter control and regulation [17].

GPP-based hybrids

GPPs are designed for a variety of computational tasks. From the end of the '80s on, GPPs are mostly Reduced Instruction Set Computers (RISC), characterised by a small set of instructions executable in a single cycle. Examples of RISC-based GPPs are Motorola's PowerPC and Sun's SPARC. GPPs are quite popular in the BI area, since they rely on standard OS and development systems. Typical BI applications use GPPs in master-VME boards to carry out data management and to interface Front-Ends to Control/GUI levels. In such capacity, GPPs depend on

well known Real-Time Operating Systems (RTOS), such as VxWorks, Lynx and, lately, Linux RTAI [18]. Nowadays, the increase in clock speed allows GPPs to carry out data processing and slow control actions, allowing taking over roles once reserved to DSPs [18,19]. This crossover is now even more favoured by the tendency to integrate DSP hardware features and specialised instructions into GPPs, yielding *GPP-hybrids*. Examples are the Pentium with MMX extension and the PowerPC with Motorola's AltiVec extension, the latter's core architecture being shown in Figure 3. The AltiVec 128-bit SIMD unit adds up to 16 operations per clock




Figure 3: GPP-hybrids. The core of Motorola's PowerPC with AltiVec extension, schematic view.

cycle, in parallel to the Integer and Floating Point units, and 162 instructions to the existing RISC architecture.

DIGITAL SIGNAL PROCESSING IN BI

Of the many techniques useful to BI this paper only covers four: 1) Digital Filters, 2) Wavelets, 3) Spectral Analysis and 4) Algebraic Reconstruction Technique (ART). Of these, 1) and Fast Fourier Transform (FFT) as part of 3), are well-known BI tools. Others, such as 2) and 4), are less frequent and used as off-line post-processing schemes. Still, they show promise for integration in real-time online processing.

Finite Impulse Response (FIR) and Infinite Impulse Response (IIR) filters are widely used to filter unwanted spectral components out of a signal. In adaptive filtering techniques the FIR and the IIR filter parameters are varied in real-time, depending on "system unknowns", to achieve a specific criterion. For example, CERN's DSP-based current monitoring system [14], aims at maximising S/N. Digital FIR filters with 2 to 5 taps typically integrate feedback loops. The filter aims to determine a suitable damping action by performing DC rejection and providing a $\pi/2$ phase shift at the oscillation frequency.

Wavelet algorithms expand signals in terms of small wavelike components, i.e. wavelets, with finite energy concentrated in time or space [20]. Suited to approximating data with sharp discontinuities, they are applied to digital filtering [21] and to non-parametric system identification for feedforward control [22].

In spectral analysis a signal is analysed in the frequency domain, splitting it into possibly infinite spectral components, each with a different amplitude and phase. Fourier Series and Transforms and the FFT are the tools of the trade. FFT alone may not guarantee a sufficient

frequency resolution, unless a very large number of data points is used, to the detriment of memory and computation time. Interpolating around the FFT-magnitude peak of interest overcomes this problem. Other methods, such as the Numerical Analysis of Fundamental Frequency [23], allow similar refinements and BI applications abound [24, 25]. A less-known analysis method is Lomb's Periodogram. This is particularly suited to non-equispaced data, as it weighs the data on a "per point" rather than on a "per time interval" basis. Hence, it can resolve frequencies with a number of points insufficient for frequency resolution by standard FFT. Another advantage over the FFT is the lack of constraints on the number of data points. A drawback is the lack of phase information, as only amplitudes are obtained. An application to tune measurements by multiturn beam position analysis is given in [26].

ART is an algorithm used in beam tomography, to reconstruct 2-dimensional images from 1-dimensional beam projections [27], similarly to Computer Assisted Tomography. The difference from the medical application is in the added complexity caused by the beam motion and deformation. A modified ART accounts for this by a particle-tracking algorithm. Figure 4 shows the measured CERN PS Booster injected proton particle density as a function of particle energy and time of arrival at the




Figure 4: Longitudinal phase space density of a bunched beam during adiabatic capture in the PS Booster [28].

detector [28]. Beam tomography is applied both to the longitudinal [29] and to the transverse case [30, 31].

APPLICATIONS

A small-scale literature search places control & feedback BI applications at the forefront of digital technology, whereas diagnostics *per se* are somewhat lagging behind. This section highlights a few of the main digital applications in diagnostics and control/feedback.

Beam Position Monitors

In BPMs the beam position is determined from the ratio between the difference and the sum of two voltages induced in an electrode [3]. Figure 5 outlines a BPM variation involving DSPs and DRX, developed in jointly by SLS and ELETTRA [5]. Four voltage signals proportional to the beam's horizontal and vertical

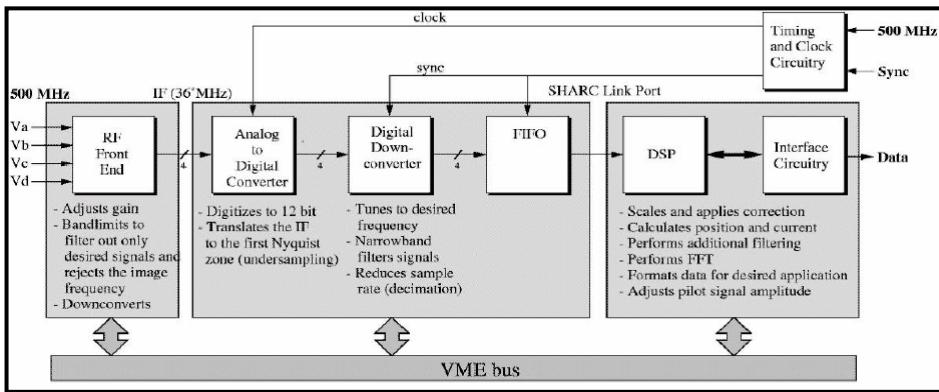


Figure 5: The digital BPM system used on Elettra and SLS – Block diagram [5].

positions are downconverted to IF = 36 MHz and low-pass filtered. They are then sampled, translated to baseband and decimated in the DRX (Digital Down Converter in Figure 5). The DRX bandwidth (BW) is user selectable from a few hundred Hz to over 1 MHz; this allows several measurements modes such as turn-by-turn, feedback, closed orbit and tune modes. When set to low-bandwidth (typically $BW < 10$ kHz) it allows following “slow” aspects of the beam evolution with high precision. An example is the closed-orbit mode, which allows resolution of less than $1 \mu\text{m}$. When BW is wider it is possible to look at “fast” aspects of the beam position with a resolution lower than $20 \mu\text{m}$ (turn-by-turn mode). The DSP can also carry out FFTs to determine the tune from turn-by-turn data.

Other examples of BPM systems include that under development for FNAL’s Recycler Ring [32], the beam position and phase detection system developed for FNAL’s Main Injector/Tevatron/Recycler Low-Level RF [33] and the system developed for KEKB [13].

Control and Feedback

Transverse and/or longitudinal beam oscillations may arise due to various factors and may have to be damped by external means, to prevent beam loss. Damping these

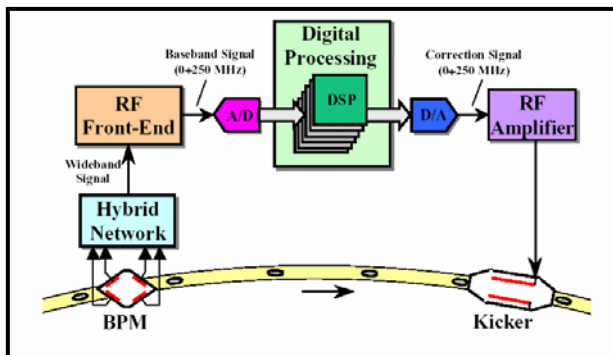


Figure 6: ELETTRA’s transverse multi-bunch feedback system – schematic view [16].

oscillations is the task of the beam feedback system. In addition, the data processing part of a digital feedback system may be a very useful tool to investigate beam instabilities. Feedback systems may be studied from a

frequency-domain or from a time-domain viewpoint (“bunch-by-bunch” feedback). The former approach, the “mode-based feedback”, requires the detailed knowledge of the spectral form of the external forces (forcing and damping), hence is somewhat more complex. Mixed digital/analogue control/feedback systems date back to the early ‘90s [34].

Figure 6 shows ELETTRA’s multibunch feedback system, which deals with 432 bunches 2 ns apart [16]. A baseband signal corresponding to the position error of the bunches passing through the BPM is sampled at 500 MHz by one ADC and de-multiplexed to six COTS DSP boards working in parallel. Each board includes four DSPs, three for feedback tasks and the last working on diagnostic tasks. The three feedback DSPs process data




Figure 7: Transverse cross-section of the ELETTRA beam. Left – no feedback, right - vertical feedback [36].

corresponding to 72 bunches per turn. Processing is by a 5-tap FIR filter and corrective actions are applied with a 4-turn delay owing to ADC/DAC delays. The diagnostic DSP can acquire the bunch position for over 200,000 turns, allowing analysing the evolution of transverse coupled multibunch instabilities [35]. Figure 7 is a transverse cross-section of ELETTRA’s beam via a synchrotron radiation profile monitor, showing the effect of the multi-bunch transverse feedback [36].

The longitudinal feedback system developed for PEP II, DAΦNE and ALS [37], also used as diagnostics, is another multi-bunch feedback example. Other feedback applications include RHIC’s beam control and feedback system [38], implementing a state-space technique and BNL’s ABS Booster upgrade, now being developed [39]. Orbit feedback systems are presented in [40,41,42].

Betatron Tune Measurement by PLLs

The betatron tune is among the key machine operation parameters. Several measurement methods are available, differing in the beam excitation used and the processing technique. An important aspect is the excitation strength required to achieve a given tune resolution, keeping in mind that a weak excitation is preferred to avoid emittance growth. The method based on using a Phase-

Locked Loop (PLL) needs usually a continuous excitation with amplitude smaller than needed for standard FFT-based methods. In addition, it provides a continuous measurement of the tune, hence it can be used for feedback. Figure 8 shows PLL-based tune measurement principles. The beam experiences a transverse sinusoidal excitation, with radial frequency ω , generated by a

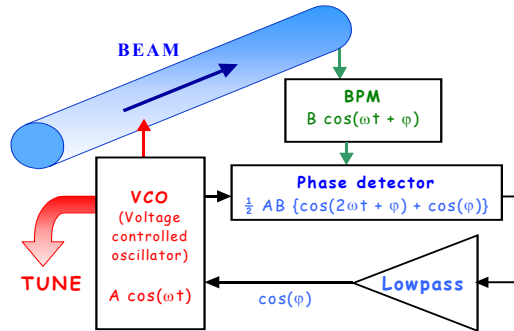


Figure 8: PLL-based tune measurement principles.

Voltage Controlled Oscillator (VCO). A BPM detects the ensuing betatron oscillation; its signal $B \cdot \cos(\omega t + \varphi)$, φ being the phase difference between beam response and excitation, is multiplied by the excitation signal $A \cdot \cos(\omega t)$ in the Phase Detector module. The result has a component with radial frequency $2 \cdot \omega$, which is eliminated by means of a low-pass filter, leaving a $\cos(\varphi)$ term. The system changes the excitation frequency until the φ equals 90 degrees. At that time the PLL is locked and the betatron tune can be read directly from the VCO as the excitation signal frequency.

In software PLLs (SPLL), features such as the auto-regulation of the excitation amplitude, based on the beam response amplitude, can be added thanks to a DSP. An “intelligent behaviour” may be implemented, with several states and working modes, such as swept mode or PLL-locked/unlocked status. The latter feature is useful when a tune-correction feedback loop relies on the PLL tune. These features make the SPLL state diagram much more complex than that used for BPMs or feedback loop systems. RHIC have been fine-tuning a similar SPLL-based tune measurement system [43].

CONCLUSION

A major transition is slowly taking place in the BI area, from all analogue to mostly-digital philosophies. This sea change imposes a concurrent change of culture. The many advantages of this transition include an ever-decreasing system size, increased flexibility through reconfigurability, the possibility to benefit from other labs’ development and easier inter- and multi-labs collaboration. Digital system development still requires a certain degree of expertise, but is becoming easier thanks to powerful development systems available on the market.

ACKNOWLEDGMENTS

I am grateful to many colleagues at CERN, in particular

to H. Schmickler, F. Pedersen, S. Hancock, M. Giovannozzi, J. Serrano and P. Baudrenghien. Others, from all over the globe, also deserve sincere thanks: P. Cameron and J. DeLong (BNL), J. Crisp (FNAL), L. Doolittle, A. Ratti and M. Placidi (LBNL), C. Hovater (TJNAF), K.T. Hsu (SRRC), M. Lonza and M. Ferianis (Sincrotrone Trieste), M. Plum, A. Regan, J. Power (LANL), V. Schlott (PSI) and R. Ursic (Instr. Technol.).

REFERENCES

- [1] R. C. Kimball, BIW’02, Upton, 2002, p. 79.
- [2] I. Linscott, BIW’00, Cambridge, 2000, p. 61.
- [3] G. Vismara, BIW’00, Cambridge, 2000, p. 36.
- [4] P. Zanella, CERN Computer Newsletter, Vol. 36, Issue 2, April-June 2001.
- [5] M. Dehler et al., PAC’99, New York, 1999, p. 2087.
- [6] M.E. Angoletta et al., PAC’03, Portland, WPPB023.
- [7] C.H. Kuo et al, EPAC’98, Stockholm, 1998, p. 1529.
- [8] E. Calvo-Girardo, CERN, private communication.
- [9] L. Doolittle et al., LINAC’02, Gyeongju, TU428.
- [10] T. J. Shea et al., PAC’01, Chicago, 2001, p. 512.
- [11] K. D. Lohmann et al., CERN/LEP-BI/88-45 (1988).
- [12] S. W. Smith, “The Scientist And Engineer’s Guide To Digital Signal Processing”, California Technical Publishing, ISBN 0-9660176-3-3 (1997), p. 509.
- [13] M. Tejima et al., EPAC’00, Vienna, 2000, p. 1815.
- [14] J. Musson et al., PAC’01, Chicago, 2001, p. 2329.
- [15] D. Barr et al., BIW’00, Cambridge, 2000, p. 242.
- [16] M. Lonza et al., ICALEPCS’99, Trieste, 1999, p. 255.
- [17] J. Carwardine et al., ICALEPCS’99, Trieste, p. 171.
- [18] M. Ferianis et al., DIPAC’01, Grenoble, p. 174.
- [19] K. T. Hsu, SRRC/TLS, private communication.
- [20] M. Vetterli, C. Herley, IEEE Trans. Signal Proc., 40 (1992) 2207.
- [21] Y. Yin et al., PAC’95, Dallas, 1995, p. 2622.
- [22] Y.-M. Wang et al., PAC’01, Chicago, p. 1613.
- [23] J. Laskar et al., Physica D, 56 (1992) 253.
- [24] M. Belgroune et al., EPAC’02, Paris, 2002, p. 1232.
- [25] R. Bartolini et al., PAC’99, New York, p. 1557.
- [26] M.E. Angoletta et al., EPAC’02, Paris, p. 1948.
- [27] D. Raparia et al., PAC’97, Vancouver, 1997, p. 2023.
- [28] <http://tomograp.home.cern.ch>.
- [29] C. Montag et al., EPAC’02, Paris, 2002, p. 1942.
- [30] M. Geitz et al., PAC’99, New York, 1999, p. 2175.
- [31] V. Yakimenko et al., EPAC’98, Stockholm, p. 1641.
- [32] J. Crisp, FNAL, private communication.
- [33] B. E. Chase et al., BIW’02, Upton, 2002, p. 368.
- [34] E. Kikutani et al., EPAC’00, Vienna, 2000, p. 192.
- [35] L. Tosi et al., PAC’01, Chicago, 2001, p. 2066.
- [36] S. N. Thanos et al., BIW’02, Upton, 2002, p. 248.
- [37] J. Fox et al., PAC’99, New York, 1999, p. 636.
- [38] M. Brennan et al., EPAC’98, Stockholm, p. 1705.
- [39] J. DeLong, BNL, private communication.
- [40] J. Carwardine et al., BIW’98, Stanford, 1998, p. 125.
- [41] M. Böge et al., EPAC’02, Paris, 2002, p. 2067.
- [42] B. Podobedov et al., PAC’01, Chicago, 2001, p. 396.
- [43] P. Cameron et al., EPAC’02, Paris, 2002, p. 1855.