
Article:
Yearby, K.H., Walker, S.N. and Pickett, J.S. (2026) Signal processing for the cluster wideband data burst mode. Journal of Geophysical Research: Space Physics, 131 (2). e2025JA034623. ISSN: 2169-9380

https://doi.org/10.1029/2025ja034623

JGR Space Physics

METHOD
10.1029/2025JA034623

Special Collection:
Cluster mission: An in-depth view of instrument calibration and cross-calibration

Signal Processing for the Cluster Wideband Data Burst Mode

K. H. Yearby[1], **S. N. Walker**[1], **and J. S. Pickett**[2]

[1]School of Electrical and Electronic Engineering, The University of Sheffield, Sheffield, UK, [2]Department of Physics and Astronomy, University of Iowa, Iowa City, IA, USA

Abstract The Wideband Data instrument is part of the Cluster spacecraft Wave Experiment Consortium. Its primary data path is a direct connection to the spacecraft data handling system providing real time downlink to the ground stations of the Deep Space Network and Panska Ves Observatory. However, it was recognized during the mission design phase that this link may not always be available, especially given that simultaneous data acquisition from the four Cluster spacecraft required the use of four ground stations. Therefore, a secondary data path at reduced bit rate was included whereby the data was transferred to the Digital Wave Processor instrument and then to the spacecraft Solid State Recorder. Given that available resources were limited, both for onboard hardware and within the spacecraft assembly, integration and testing program, the design of this backup data path was less than optimal. Although it was verified during ground testing that data could be acquired via this route, the design did not make the best use of the available telemetry bandwidth, and the timing accuracy was too limited to support some multi-spacecraft observations. This paper describes work around solutions to optimize bandwidth utilization and timing accuracy. These involve patches to the onboard software of the Digital Wave Processor instrument and ingenious signal processing on the ground.

Key Points:

- The Wideband Data instrument on the Cluster spacecraft includes a secondary data path to enable data to be recorded onboard
- Limited resources meant that the design of this secondary data path was less than optimum for telemetry utilization and timing accuracy
- Work around solutions are described that enabled a substantial scientific data set to be acquired using this interface

Correspondence to:
S. N. Walker,
simon.walker@sheffield.ac.uk

Citation:
Yearby, K. H., Walker, S. N., & Pickett, J. S. (2026). Signal processing for the Cluster wideband data burst mode. *Journal of Geophysical Research: Space Physics*, *131*, e2025JA034623. https://doi.org/10.1029/2025JA034623

Received 27 AUG 2025
Accepted 9 JAN 2026

Author Contributions:
Conceptualization: K. H. Yearby, S. N. Walker
Formal analysis: K. H. Yearby
Funding acquisition: K. H. Yearby
Methodology: K. H. Yearby, J. S. Pickett
Software: K. H. Yearby
Validation: K. H. Yearby, J. S. Pickett
Visualization: K. H. Yearby
Writing – original draft: K. H. Yearby, S. N. Walker, J. S. Pickett
Writing – review & editing: K. H. Yearby, S. N. Walker, J. S. Pickett

1. Introduction

1.1. Wave Experiment Consortium

The four Cluster spacecraft were launched in July and August 2000 into elliptical polar orbits with an apogee up to 22 R_E. The initial planned mission duration was 2.5 years, but this was extended many times until science observations finally ended in September 2024 following the re-entry of spacecraft 2. During the mission, the separation between the spacecraft varied from a few km to more than 12,000 km. The primary scientific objective of the Cluster mission was to study in three dimensions the interaction of the solar wind with Earth's magnetosphere.

The suite of scientific instruments includes the Wave Experiment Consortium (WEC) (Pedersen et al., 1997), a collaboration to coordinate wave measurements and optimize usage of the limited spacecraft resources. WEC brings together the operations of the Electric Fields and Waves (EFW) (Gustafsson et al., 1997), Waves of High frequency and Sounder for Probing of Electron density by Relaxation (WHISPER) (Décréau et al., 1997), WideBand Data (WBD) (Gurnett et al., 1997), and Spatio-Temporal Analysis of Field Fluctuations (STAFF) (Cornilleau-Wehrlin et al., 1997) under the control of a central unit, the Digital Wave Processor (DWP) (Woolliscroft et al., 1997). Each spacecraft includes a pair of wire boom electric field antennae orientated in the spacecraft spin plane, measuring 88 m from tip to tip, as well as a three axis search coil antenna for measuring AC magnetic fields. The signal from either the electric field antenna or either of the two axes of the search coil may be connected to the WBD instrument (but only one at a time).

As originally envisaged, DWP would receive commands from the spacecraft command and data handling system, using them to configure the WEC sensor instruments and synchronize their operation. The sensor data from EFW, STAFF, and WHISPER were to be collected by DWP and placed into the spacecraft telemetry stream. The primary route for data from the WBD was a direct link to the NASA Deep Space Network (DSN), and later in the mission to the Panska Ves Observatory, Czech Republic in real time. It was recognized during the design phase that sufficient DSN resources would not always be available, particularly for multi-spacecraft observations, as each downlink in general requires a separate DSN ground station. Therefore, a secondary data path was included whereby the data was transferred to the DWP and via the usual WEC data link to the spacecraft Solid State Recorder. This secondary data path is termed “burst mode two” (BM2) to reflect the operating mode of the

spacecraft data handling system, even though it provides a lower data rate than the primary real time mode. This paper examines the BM2 data path, beginning with an overview of the WBD and DWP instruments and their operations, the problems encountered with the original design and the work around steps that were implemented to yield a high quality electric and magnetic field data set.

1.2. WBD

The WBD instrument is a digital waveform receiver providing measurements of electric or magnetic fields in several bands covering in total a frequency range from 50 Hz to 577 kHz. The WBD analog to digital converter provides resolutions of 1, 4, or 8 bits, although the latter has been used almost exclusively for scientific observations. The sampling rate is selectable between 27.443, 54.886, or 219.544 kHz, providing a usable bandwidth of 9.5, 19, or 77 kHz. Frequencies above 77 kHz are handled by frequency conversion, utilizing conversion frequencies of approximately 125, 250, and 500 kHz. The receiver incorporates an automatic gain control (AGC) system covering 75 dB in 15 steps of 5 dB, which together with the analog to digital converter gives a dynamic range of over 100 dB. The instrument is designed to enable the highest possible timing accuracy, with all sampling and conversion clocks derived from a 14 MHz reference oscillator synchronized to the spacecraft Ultra Stable Oscillator (USO). Its primary data path is a direct connection to the spacecraft data handling system providing real time downlink to the ground stations of DSN and Panska Ves at a bit rate of 220 kbs^{-1}.

As well as the actual waveform data, the WBD instrument outputs status data including the selected instrument mode, and in particular the gain step of the AGC system, which is essential to be able to calibrate the data. In the real time mode, the status data are simply appended as 6 bytes in every data frame. In addition, the status data are transferred to the DWP and included within the WEC experiment housekeeping (HK) data to allow the instrument to be monitored through the usual spacecraft ground system. Given that the WBD status data are anyway contained within the WEC HK data, they are not included in the science data stream when the secondary BM2 data mode is in use.

1.3. DWP

The DWP is the central controller of the WEC and contains dual redundant interfaces to the spacecraft data handling system as well as dedicated interfaces to each WEC instrument. Here we concentrate on those aspects of the DWP design relevant to the processing of WBD data. For a detailed description of the DWP, the reader is invited to consult Woolliscroft et al. (1993), Woolliscroft et al. (1997), and Dunford et al. (1991).

The DWP includes three processor cores, each comprising an Inmos T222 transputer, 32K bytes of read only memory (ROM) and 32K bytes of volatile random access memory (RAM). The ROM, which is relatively slow, is only used to hold the code for the boot strap process—it is copied into the faster RAM for execution. There is no programmable non-volatile memory—the ROM contents cannot be changed and the RAM is reset to an undefined state every time power is removed. These facts represent a major limitation as to what can be achieved by patching the code. Any patch must be reloaded over the low rate command interface every time the instrument is switched on, which limits the total size of the patches to a few hundred bytes. A further limitation is that the RAM is subject to radiation induced single event upsets where random bit flips occur of the order of once every 250 hr (Yearby et al., 2014). Therefore, any code patch simply loaded into memory is not guaranteed to remain intact even while the instrument remains powered.

A key function of the DWP is to synchronize the data sampling of the WEC instruments. To this end, it generates a 900 Hz master clock which is divided down to produce the EFW and STAFF waveform data sampling clocks of 25 Hz (normal mode) or 450 Hz (burst mode), a 75 Hz clock for WHISPER and a variable rate clock between 1 and 10 Hz controlling the WBD AGC system. The 900 Hz frequency was chosen to allow the required WEC sampling frequencies to be easily derived by simple division. It is in turn derived from a crystal oscillator in DWP. Regrettably, resources did not allow it to be synchronized to the spacecraft USO or the WBD data sampling clock.

In the DWP, there were insufficient high rate data interfaces to allow one to be dedicated to the WBD secondary data path considering that it was expected this would rarely be used. Instead, it was decided that WBD would share the interface with the WHISPER instrument (Décréau et al., 1997), an instrument that covered a similar frequency range (1–80 kHz), so that only one of the two instruments could be used at once.

1.4. WEC Operations

DWP coordinates WEC observations using a facility termed "macros." These are sequences of commands, including time delays and looping, stored in the DWP memory (Woolliscroft et al., 1993). An initial set of macros is loaded from the ROM, but these may be modified by commands sent from the ground via the spacecraft data handling system. To ensure the integrity of the macros in case of single event upsets, they are loaded in an area of memory termed "checksummed memory" protected by software checksums and with a redundant copy as backup. The checksum and redundant copy are regularly checked and updated whenever this memory is modified. An unexpected change in the checksum triggers a reboot and if necessary, the working copy of the checksummed memory will be restored from the backup.

1.5. Difficulties

The Cluster spacecraft went through an extensive assembly, integration and test program. This ensured that all the instrument modes worked as designed, and in particular for the subject of this paper, that WBD data could be acquired via the secondary interface in the BM2 mode. There was also an extensive effort in planning the Cluster scientific observations, but for WBD these concentrated on the use of the primary real time interface given that this provided a much higher bandwidth. It was only in 2004, after the mission had been extended beyond its original 2 years that detailed plans were made to use the secondary interface for science observations and then a number of difficulties became apparent. These were:

- Being unable to take WBD and WHISPER data simultaneously was a significant drawback.
- Use of the WEC telemetry was not optimized in that there was insufficient telemetry to allow the other WEC instruments to operate at high bit rate (450 Hz sampling) when WBD was in use, while on the other hand there was substantial unused telemetry if they operated at low bit rate (25 Hz sampling).
- The timing accuracy of the WBD science data in BM2 was not sufficient to support multi-spacecraft observations.
- The timing accuracy of the gain information in the WBD status data was not sufficient to allow the science data to be properly calibrated.

Work around solutions were designed to address all these problems, involving patches to the onboard software of the DWP instrument and ingenious signal processing on the ground. These are described in the following sections.

The WBD data, and all data from the Cluster mission, are archived at the Cluster Science Archive (CSA, https://csa.esac.esa.int) (Laakso et al., 2010; Pickett et al., 2010). Over the full duration of the Cluster mission, a total of around 4,000 hr of data have been collected in BM2 by the four spacecraft, compared to around 37,300 hr in the real time mode. The BM2 data were primarily taken on all four spacecraft simultaneously, unlike the real time mode when this was often not possible, as it depended on having three or four antennas on the ground available to receive the data. The WBD data sets are numerical waveform data and graphical spectra derived from them.

2. Secure Code Patching

A pre-requisite to any solution involving patches to the onboard code is a way to implement patches securely so that the modified code will survive radiation induced upsets in the memory. This is achieved by loading the patch instructions into the "checksummed memory' primarily used to store the WEC macros. It may also be used to hold code patch instructions configured to load into memory automatically during a reboot. The limited size of this memory (256 words), which is also needed to implement the WEC modes, is a further constraint on the complexity of patching.

Once loaded into the main memory, the patches and all other code are protected by a software checksum, recomputed in the background around every 5 s. In the event that it changes due to a single event upset, a reboot is triggered. This reloads the main code from the ROM and then overlays the patches from the checksummed memory.

3. Telemetry Optimization

Whilst the DWP hardware precludes simultaneous acquisition of WBD and WHISPER data, time sharing is possible. With the help of a simple code patch, the dead time between WBD and WHISPER acquisition can be

Table 1
WBD Digital Filter Characteristics

Implementation	Sampling rate *fz* (kHz)	Pass band *fp* (kHz)	Stop band *fs* (kHz)	Usable bandwidth (kHz)
Original in/3 mode	9.148	3.860	4.500	3.860
Original in/4 mode	6.861	3.860	4.500	2.361
Optimized/3	9.148	4.254	4.912	4.236
Optimized/4	6.861	3.101	3.787	3.074

reduced to less than 1 s. WEC observation modes are usually based on a 52 s cycle or a multiple thereof, and it was decided to use a 5 s WHISPER data acquisition comprising two active sweeps and 16 passive sweeps, followed by 47 s of WBD acquisition.

The WEC data link supports a maximum of 91 kbs^{-1}. To accommodate the WBD data and still leave room for data from the other WEC instruments, in the original design the WBD data are reduced by a factor of three to 73 kbs^{-1}, either by digital filtering (reducing the sampling rate) or duty cycling (selecting only one out of three frames). The digital filtering option is used when it is required to maintain continuous waveform acquisition, but this is achieved at the cost of a reduction in the usable bandwidth. Digital filtering may only be used when the WBD is operated at its lowest sampling rate of 27.443 kHz, which provides continuous waveform acquisition. Alternatively, the duty cycle option allows the full bandwidth to be maintained at the cost of lower time resolution. This option may be used with all WBD sampling modes. With a WBD data rate of 73 kbs^{-1} and an overall WEC data link rate of 91 kbs^{-1}, there is 18 kbs^{-1} for the other WEC instruments (EFW and STAFF). This is far more than required when they operate at normal bit rate (25 samples per second—4,136 bs^{-1}) but not sufficient for high bit rate (450 samples per second—minimum of 33,864 bs^{-1}). Hardware limitations preclude operating at intermediate bit rates. To allow WBD observations with EFW and STAFF at high bit rates, a "divide by 4" WBD down sampling mode was implemented in DWP using a simple code patch. Like the standard "divide by 3" mode, this can be achieved either by digital filtering or duty cycling. The resulting bit rate is 55 kbs^{-1} and leaves 36 kbs^{-1} for EFW and STAFF.

4. Digital Filter Optimization

A "finite impulse response"(FIR) digital filter is employed to reduce the bandwidth of the WBD data as a prelude to reducing the bit rate by decimation. A FIR filter works by calculating the convolution of the input time series with a series of coefficients. The frequency response of the filter is given by the Fourier transform of the coefficients. It is easy to modify the response of such a filter by uploading a new set of coefficients. The provenance of the original coefficients is unclear, but analysis showed they were not optimum.

The response of a low pass filter is specified by four parameters: the pass band ripple, the frequency of the pass band edge (*fp*), the stop band edge (*fs*), and the stop band attenuation. An optimum filter must have an acceptably low pass band ripple, a transition region between *fp* and *fs* as narrow as possible, and a stop band attenuation sufficient so that any frequency aliasing effects folded into the pass band are reduced to an acceptable level (ideally below the noise level).

The filter response impacts the maximum usable bandwidth in two ways. Firstly, the usable bandwidth is directly limited by the passband (*fp*) of the filter. Secondly, we must consider the impact of aliases (frequencies above half the sampling rate, *fz*, being reflected into the passband).

A frequency *fa* may be reflected into the pass band at a frequency of $fz - fa$. This effect is constrained by the filter stop band (*fs*). In such a case, the maximum usable bandwidth is given by the minimum of *fp*, and $fz - fs$.

The filter characteristics are listed numerically in Table 1 and shown graphically in Figure 1.

If we take the case of the WBD divide by 3 mode with the original filter (Figure 1a), the pass band edge is 3.86 kHz, whereas the stop band edge is 4.50 kHz. The sampling rate is 9.15 kHz, so protection from aliases is achieved up to 4.65 kHz. The maximum usable bandwidth is the lower of these two figures that is 3.86 kHz. If we use the same filter with the divide by 4 mode, sampling rate 6.86 kHz, the alias limit is just 2.36 kHz and this sets the maximum usable bandwidth.

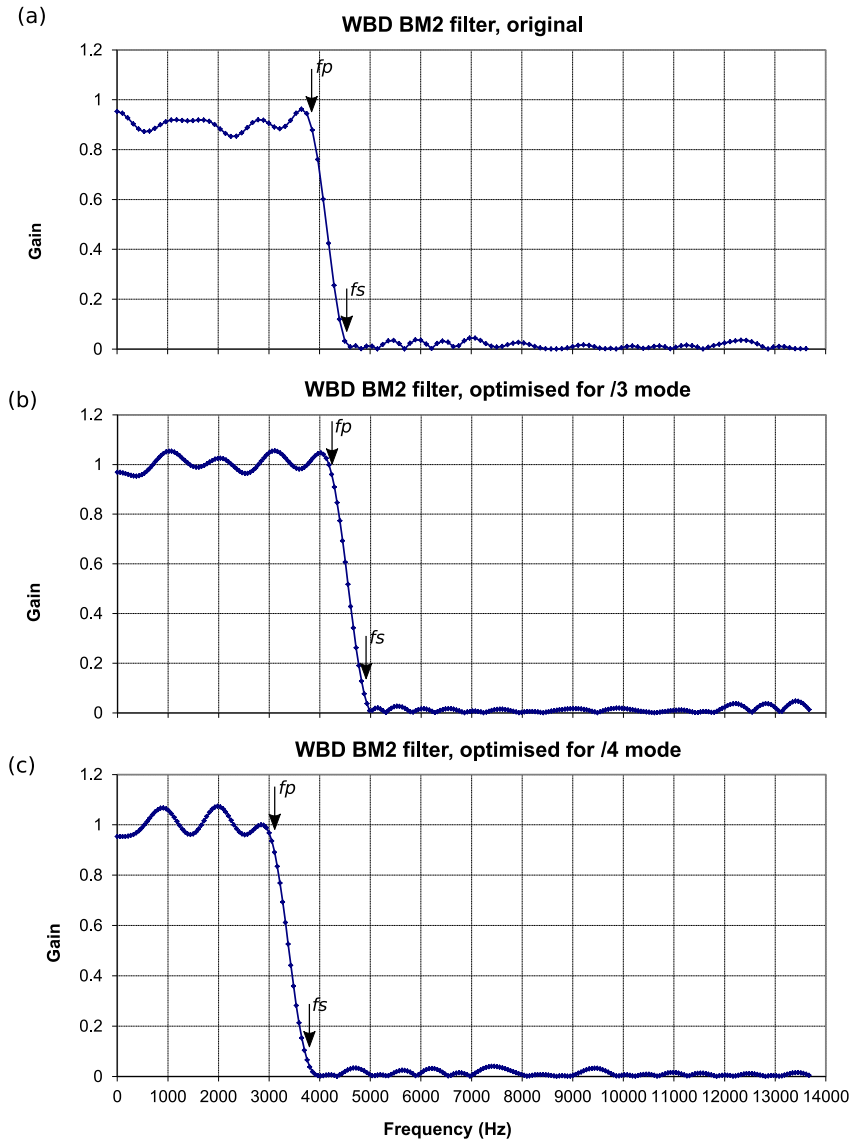

Figure 1. WBD filter gain as a function of frequency. Panel (a) original filter, (b) optimized/3, (c) optimized/4. The pass band edge *fp* and stop band edge *fs* are marked in each case.

For an optimum filter, we should aim for the pass band edge to be the same as the alias limit. For the divide by 3 mode, this is implemented with the filter in Figure 1b. The pass band edge is 4.254 kHz and the stop band edge 4.912 kHz. The alias limit is thus 4.236 kHz. For the divide by 4 mode, we use the filter in Figure 1c. This has a pass band edge at 3.101 kHz and a stop band edge at 3.787 kHz. The alias limit is thus 3.074 kHz. The optimized filters also show some improvement in the pass band gain (closer to unity) and for the divide by 3 mode better stop band attenuation.

The optimized filters were designed using the Constrained Least Squares method (Adams, 1991). This allows each parameter to be separately specified and uses an iterative method to attempt to achieve the desired response. This is an improvement on the classical method using a Fourier transform, which produces filters with symmetrical pass band and stop band ripple (the two cannot be separately specified) and tends to require more coefficients (and hence processing power) to achieve the same performance. The Adams method provides real valued coefficients—these were truncated to integer values required by the DWP software and listed in Table 2. The response was checked using a Fourier transform to produce the graphs in Figure 1.

Table 2
WBD Digital Filter Coefficients

Index	Original	Optimized/3	Optimized/4
1	1	0	0
2	1	0	1
3	0	1	1
4	−1	1	1
5	−1	1	0
6	0	−1	−1
7	1	−1	−1
8	1	−1	−1
9	1	1	−1
10	0	1	0
11	−1	1	1
12	−1	−1	1
13	0	−2	0
14	2	−1	−1
15	2	1	−2
16	1	2	−2
17	−2	1	−1
18	−3	−2	1
19	−2	−4	3
20	1	−2	3
21	4	2	1
22	4	5	−2
23	0	3	−6
24	−6	−4	−7
25	−7	−9	−3
26	−2	−6	5
27	10	8	15
28	24	27	25
29	34	41	31

Note. The table shows the integer coefficients used in the DWP onboard software. The filter is symmetrical of length 2N, where N is 29, such that the coefficient at index 30 is the same as that at index 29, and so on.

5. Data Timing Accuracy

Data timing on the Cluster spacecraft is implemented as follows. The onboard data handling system uses the ESA Packet Telemetry Standard (PSS-04-106, available at http://microelectronics.esa.int/vhdl/pss/PSS-04-106.pdf). In the WBD data modes, this uses a raw bit rate of 262,144 (2^{18}) bs^{-1} synchronized to the USO. Data are packed into transfer frames with a data field of 1,096 bytes and an overall size including sync marker, headers and Reed Solomon code block of 1,279 bytes. The maximum usable data rate is thus $2^{18} \times 1{,}096/1{,}279 \approx 224.6$ kbs^{-1}. For the WBD instrument, a constant bit rate just below this, synchronized to the USO, of $2^{22}/19 \approx 220.8$ kbs^{-1} is used with idle frames inserted to pad the remainder. A total of 132 transfer frames were packed into a "format." The format duration is thus $132 \times 1{,}279 \times 8/2^{18} \approx 5.152$ s. The onboard data handling system issues a timing reference pulse synchronized with the transmission of the first bit of the first transfer frame in each format. The time of each pulse is measured onboard by counting the cycles of a high frequency clock and this count is included in the header of the data packets as the on-board time (OBT). For more details, the reader is referred to the Cluster Experiment Interface Document Part A archived at the Cluster Science Archive (https://caa.esac.esa.int/documents/CM_CD_AUX_EID_A.pdf).

The ground system calibrates the OBT with respect to Coordinated Universal Time (UTC) and records the calibration in time calibration (TCAL) files included with the raw data. As originally designed, the error between the calibrated OBT and UTC was allowed to increase to 2 ms before performing a new calibration. Later in the mission, calibration was performed every orbit, which was usually sufficient to keep the error at the level of a few tens of microseconds.

The DWP was designed in accordance with the original stated timing accuracy for the Cluster mission of ±2 ms, whereby the time stamps for science data packets are specified as the number of cycles of the WEC 900 Hz master sampling clock that have occurred since the most recent timing reference pulse issued by the onboard data handling system. It then became apparent that a much better mission timing accuracy (~20 microseconds) could be achieved in practice and that this accuracy was indeed required for inter-spacecraft comparison of data acquired in burst modes. Within DWP, this required an additional parameter giving the time delay between the timing reference pulse and the first cycle of the WEC 900 Hz master sampling clock that follows. For further details, see Yearby et al. (2010, 2013).

However, WBD burst data is not synchronized to the WEC sample clock, and thus is not accurately time tagged even using the timing enhancements above. The time of each WBD packet is well determined, but there is an unknown offset between 0 and 1,111 μs following this event and the 900 Hz master clock cycle used to generate the time tag. Two methods have been devised to address this. Both depend on the fact that the WBD data are acquired at a very stable rate, one packet every 0.039719 s (WBD packets are comprised of 1,096 bytes at a bit rate of $2^{22}/19$ s^{-1}, hence the packet duration is $1{,}096 \times 8 \times 19/2^{22}$ s) This means that if we can determine accurate times even just occasionally, the times of all the other packets can be interpolated.

The first method works by analysis of the time stamps of a long series of WBD data packets. Taking the time modulo 0.079437 s (twice the interval between packets due to double buffering), gives a number which varies in the short term over a 1,111 μs range, as the unknown offset varies. Packets with a time modulo at the lower end of this range must have a near zero offset, so are the most accurate points. So we just have to select the packets with the lowest time modulo. This process is illustrated in Figure 2.

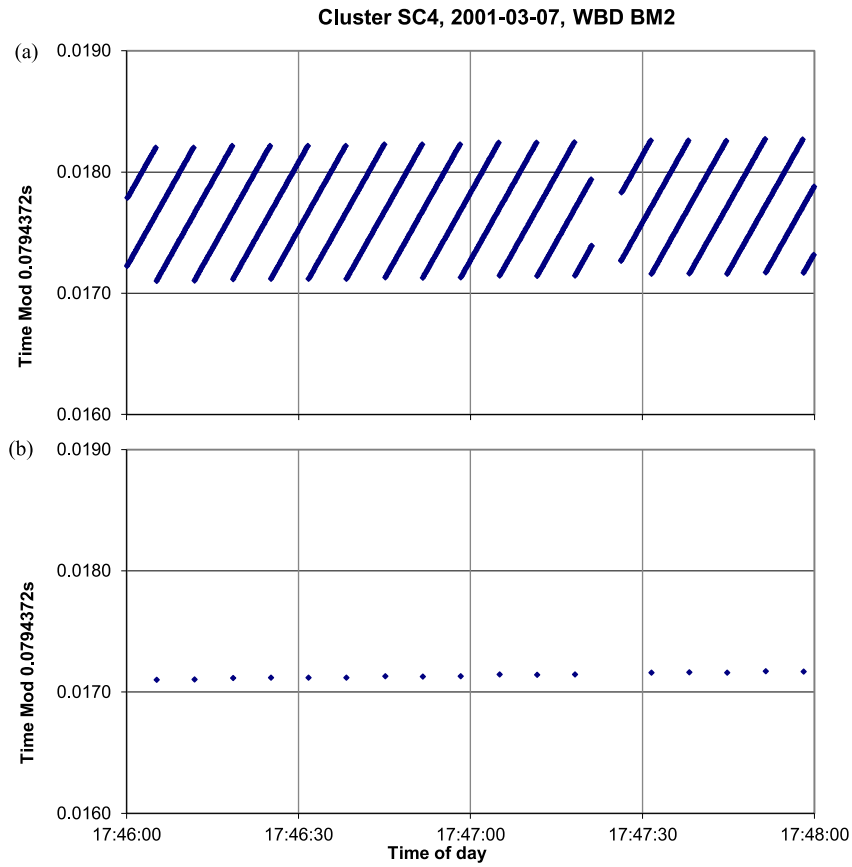

Figure 2. Cluster SC4 2001-03-07 17:46:00 to 17:48:00 time modulo (a) all packets, (b) selected packets with accurate time tags.

The second method exploits the fact that when measuring time on the OBT scale, the frame interval is **exactly** 2,665,472 ($1{,}096 \times 19 \times 2^7$) counts at 2^{24} Hz, because both the WBD clock and OBT clock are derived from the same USO. Using this fact, it is possible to extrapolate the frame times over periods of several hours or even days from the previous or following real time DSN or Panska Ves acquisition.

It is easy to convert the UTC time tags to OBT by applying the current TCAL backwards. The information needed to do this is in the TCAL files included with the raw data. This process is not exact, which will lead to a small jitter in the OBT frame interval, but over a long period the regular frame cycle is preserved. In principle, this method gives a more accurate result, but it only works if the WBD instrument remains powered the whole time between the current burst mode operation and the reference real time (DSN or Panska Ves) operation. In practice, both methods are applied when possible and compared to check the consistency of the result.

To demonstrate the timing accuracy, Figure 3 shows spectrograms of whistler mode chorus observed on Clusters C3 (a) and C4 (b). These appear almost identical except for glitches at the lowest frequencies due to gain changes. To further demonstrate the relative timing accuracy, we performed a cross-correlation of the amplitude envelope in two bands (c). The spacecraft were separated by 6.5 km, roughly along the field line, which in free space would give a propagation time of 22 μs. The actual measured time delays average 49 μs at 1 kHz and 143 μs at 2 kHz. There are no electron density data available at this time, so we cannot estimate what the delays should be, but the figures seem reasonable, and consistent with a timing accuracy of few tens of microseconds.

Further examples of the observations now possible in BM2 using enhanced timing accuracy are shown in LaBelle et al. (2022). Figures 3 through 7 of that paper are examples of AKR spectrograms produced using the BM2 data, while Figure 8 shows the footprint of the estimated location of AKR sources determined by differential time delay measurements, for which accurate inter spacecraft timing is essential. These measurements could only be

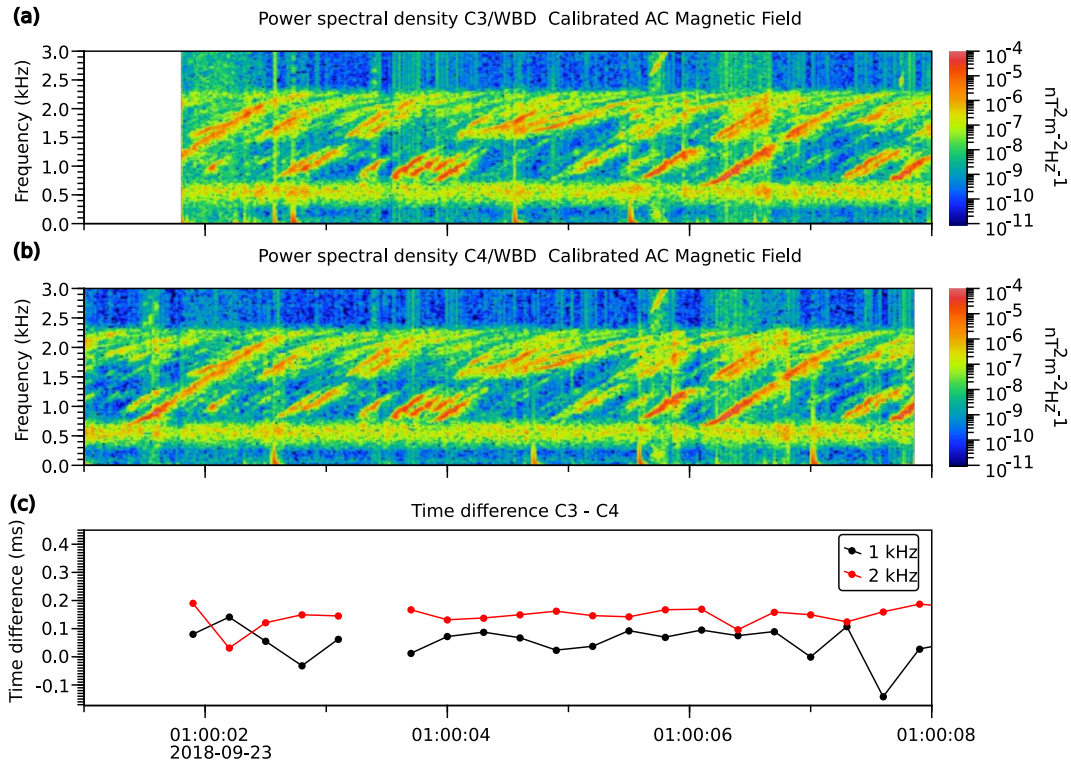

Figure 3. Spectrograms of whistler mode chorus observed on Clusters C3 (a) and C4 (b), together with the time difference determined by cross-correlation of the amplitude envelope in two bands centered on 1 and 2 kHz (c). The spacecraft were separated by 6.5 km at location R_E 6.7, MLAT 3.2°, and MLT 11.0.

performed using BM2 at this stage in the mission as there were no ground stations available to take real time data when the spacecraft was in the southern hemisphere.

6. Status/Gain Timing Accuracy

In burst mode, the WBD status information, including the gain of the WBD AGC system, is sent separately in the HK telemetry. This gain information must be recovered with sufficient time accuracy to allow the gain of each data packet to be determined. The gain changes in response to a gain update clock generated by DWP, which has a maximum frequency of 10 Hz. A maximum of 52 gain values may therefore be stored in each HK packet, which covers a period of 5.152 s, but there is no mechanism to determine which specific gain value applies to each WBD data packet (an oversight in the original design).

This is actually quite difficult to resolve. Although the frequency of the gain update clock, and the number of gains stored in each HK packet are known, there is no record of exactly when each gain change took place. However, it is possible to get a clue from the number of gains in each HK packet.

Taking a typical case of 1 Hz gain update rate, there may be either 5 or 6 gains in each HK packet. When there are 5 gains, we know that the first one must have a time between 0.152 and 0.999 s from the start of the packet (else there would be room for 6). When there are 6 gains, we know the first one must be between 0.000 and 0.152 s. Each HK packet then defines a certain time range when the gain changes could have occurred. Continuing the regular gain update cycle over many successive HK packets narrows down the range to only a few milliseconds.

The WBD gain update clock is synchronized to the DWP 900 Hz master clock, so it is possible to accurately track the gain update cycle over a long period of time. The phase only has to be determined once for it to be known over the whole observation.

The first step is to determine a cumulative master clock count by summing the number of master clock cycles that have occurred during the period of each HK packet. Then, the master clock count is divided by the gain update

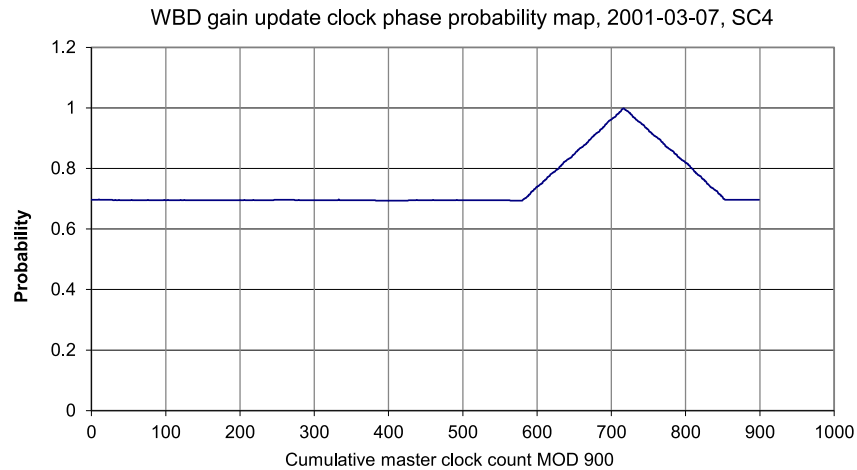

Figure 4. Cluster SC4 2001-03-07, gain update clock phase probability map.

period and the remainder is used as the index to a clock phase probability table. For each HK packet, every entry in the table is incremented when it is possible that the clock might have that phase. After processing a large number of data packets, we should find that only one or two entries will have been incremented on almost every packet.

Figure 4 shows the probability table for a typical period of BM2 data. The counts in each entry were divided by the total number of packets. We can see that in this case the gain update clock phase is 716 counts on the 900 Hz master clock scale, and this allows the time of every gain update clock pulse during the observation to be determined.

An alternative method uses a patch to the onboard software to include the status information within the science data packets, as is the case for real time data. This seems like an obvious solution but in practice it can only be applied in limited situations. Firstly, there must be sufficient spare science telemetry bit rate available to accommodate the extra data, which is not always the case. Secondly, due to severe constraints on code patch complexity, this can only be achieved in the duty cycle mode (where it is simply required to increase the size of the packets), and not the digital filter mode where a separate channel would need to be established to allow the status data to bypass the filter.

7. Conclusions

The Wideband Data instrument includes a secondary data path called Burst Mode 2 at reduced bit rate, where the data is transferred to the Digital Wave Processor instrument and then to the spacecraft Solid State Recorder. Limited resources, both for onboard hardware and within the spacecraft assembly, integration and testing program, meant the design of this backup data path was less than optimal. Work around solutions have been implemented to optimize bandwidth utilization and timing accuracy. Over the full duration of the Cluster mission and the four spacecraft, a total of around 4,000 hr of data have been collected in BM2 and are available at the Cluster Science Archive.

Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

Data Availability Statement

The WBD data, and all data from the Cluster mission, are archived at the Cluster Science Archive (https://csa.esac.esa.int/csa-web/#search). The WBD data are also available at CDAWeb (https://cdaweb.gsfc.nasa.gov/pub/data/cluster/). *Software Availability Statement*: The Matlab program adamsfir.m, used to calculate the optimized filters using the method described by Adams (1991), is available at: https://www.ece.rice.edu/dsp/software/cl2.shtml. Figure 3 was prepared using Autoplot (https://autoplot.org/).

Acknowledgments
KHY acknowledges financial support from ESA via JSOC subcontract UKRI-2630. SNW was funded by ESA contract 19023/05/NL/NR and the UK STFC ST/Y001575/1.

References

Adams, J. W. (1991). FIR digital filters with least-squares stopbands subject to peak-gain constraints. *IEEE Transactions on Circuits and Systems*, *38*(4), 376–388. https://doi.org/10.1109/31.75395

Cornilleau-Wehrlin, N., Chauveau, P., Louis, S., Meyer, A., Nappa, J. M., Perraut, S., et al. (1997). The Cluster spatio-temporal analysis of field fluctuations (STAFF) experiment. *Space Science Reviews*, *79*(1–2), 107–136. https://doi.org/10.1023/A:1004979209565

Décréau, P. M. E., Fergeau, P., Krannosels'kikh, V., Leveque, M., Martin, P. H., Randriamboarison, O., et al. (1997). WHISPER, a resonance sounder and wave analyser: Performances and perspectives for the Cluster mission. *Space Science Reviews*, *79*, 157–193. https://doi.org/10.1023/A:1004931326404

Dunford, C. M., Thompson, J. A., & Yearby, K. H. (1991). A transputer-based instrument for the ESA/NASA Cluster mission. *Concurrency: Practice and Experience*, *3*(4), 293–302. https://doi.org/10.1002/cpe.4330030407

Gurnett, D. A., Huff, R. L., & Kirchner, D. L. (1997). The wide-band plasma wave investigation. *Space Science Reviews*, *79*(1–2), 195–208. https://doi.org/10.1023/A:1004966823678

Gustafsson, G., Boström, R., Holback, B., Holmgren, G., Lundgren, A., Stasiewicz, K., et al. (1997). The electric field and wave experiment for the Cluster mission. *Space Science Reviews*, *79*(1–2), 137–156. https://doi.org/10.1023/A:1004975108657

Laakso, H., Perry, C., McCaffrey, S., Herment, D., Allen, A. J., Harvey, C. C., et al. (2010). Cluster active archive: Overview. In H. Laakso, M. Taylor, & C. Escoubet (Eds.), *The Cluster Active Archive. Astrophysics and Space Science Proceedings*. Springer. https://doi.org/10.1007/978-90-481-3499-1_1

LaBelle, J., Yearby, K., & Pickett, J. S. (2022). South Pole Station ground-based and Cluster satellite measurements of leaked and escaping Auroral Kilometric Radiation. *Journal of Geophysical Research: Space Physics*, *127*(2), 021JA029399. https://doi.org/10.1029/2021JA029399

Pedersen, A., Cornilleau-Wehrlin, N., De la Porte, B., Roux, A., Bouabdellah, A., Decreau, P. M. E., et al. (1997). The Wave Experiment Consortium (WEC). *Space Science Reviews*, *79*(1–2), 93–106. https://doi.org/10.1023/A:1004927225495

Pickett, J. S., Seeberger, J. M., Christopher, I. W., Santolík, O., & Sigsbee, K. M. (2010). Cluster wideband data products in the Cluster Active Archive. In H. Laakso, M. Taylor, & C. Escoubet (Eds.), *The cluster active archive. Astrophysics and space science proceedings*. Springer. https://doi.org/10.1007/978-90-481-3499-1_11

Woolliscroft, L. J. C., Alleyne, H. S., Dunford, C. M., Sumner, A., Thompson, J. A., Walker, S. N., et al. (1993). Implementation of the digital wave-processing experiment. In W. R. Burke (Ed.), *Cluster: Mission, payload and supporting activities* (Vol. 1159, p. 81). European Space Agency, ESA-SP.

Woolliscroft, L. J. C., Alleyne, H. S. C., Dunford, C. M., A. Sumner, A., Thompson, J. A., Walker, S. N., et al. (1997). The digital wave-processing experiment on Cluster. *Space Science Reviews*, *79*(1–2), 209–231. https://doi.org/10.1023/A:1004914211866

Yearby, K. H., Alleyne, H. S. C., Walker, S. N., Bates, I., Gough, M. P., Buckley, A., et al. (2010). Digital wave processor products in the Cluster Active Archive. In H. Laakso, M. Taylor, & C. Escoubet (Eds.), *The Cluster Active Archive. Astrophysics and Space Science Proceedings*. Springer. https://doi.org/10.1007/978-90-481-3499-1_4

Yearby, K. H., Balikhin, M., & Walker, S. N. (2014). Single-event upsets in the Cluster and Double Star Digital Wave Processor instruments. *Space Weather*, *12*(1), 24–28. https://doi.org/10.1002/2013SW000985

Yearby, K. H., Walker, S. N., & Balikhin, M. A. (2013). Enhanced timing accuracy for Cluster data. *Geoscientific Instrumentation, Methods and Data Systems*, *2*, 323–328. https://doi.org/10.5194/gi-2-323-2013