LONDON, METEOROLOGICAL OFFICE.

Met.O.16 Branch Memorandum No.7.

An operational system for the remote location of lightning flashes using a VLF arrival time difference technique. By LEE,A.C.L.

London,Met.Off.,Met.O.16 Branch Mem.No.7, 1985,31cm.Pp.40,4 pls.6 Refs.Abs.p.2.

An unofficial document - restriction on first page to be observed.

Met O 16 BRANCH MEMORANDUM No 7.

AN OPERATIONAL SYSTEM FOR THE REMOTE LOCATION OF LIGHTNING FLASHES USING A VLF ARRIVAL TIME DIFFERENCE TECHNIQUE

Anthony C. L. Lee

Meteorological Office, Bracknell

1985

ABSTRACT

An operational system for the remote location of lightning flashes at ranges of thousands of kilometers is presented. The vertical electric fields of VLF radio atmospherics ("Sferics"), together with time data, are observed at a network of Outstations dispersed over the UK and Mediterranean. These data are forwarded to a Control Station which locates individual flashes in real time using Arrival Time Difference (ATD) techniques, analogous to Loran-C, and generates regular reports for automatic distribution. The Service Area covered exceeds 40°W to 40°E, 30°N to 70°N; flash location accuracy is 2-20km; and approximately 400 flashes/hour can be located.

1. Introduction

Lee(1985) describes an experimental study of the Arrival Time Difference (ATD) technique for the location of lightning flashes at ranges of thousands of kilometers, and shows how the accuracy is potentially superior to the existing Cathode Ray Direction Finding (CRDF) operational technique by an order of magnitude or more. The study equipment was controlled by Outstation operators, and used recording and processing methods incompatible with a real-time operational system. This paper describes the application of the ATD technique to a new operational system, indicating the necessary further developments and their impact on fix rate, system availability, and accuracy.

2. An ATD replacement for the CRDF system

The obsolescent manually operated CRDF system comprises a dispersed network of CRDF Outstations, each capable of measuring the bearing of individual lightning flashes from their received radio atmospherics ("Sferics") and reporting the results to a "Control Station". Here the bearings are plotted, and a lightning flash location (or "fix") obtained. The operation and limitations of this system are described by Maidens (1953) and Lee (1985).

The CRDF system is being replaced by a highly automated ATD System which has a Service Area from at least 40°W to 40°E, 30°N to 70°N. This does not measure bearings as outlined above, but instead determines the arrival time of the Sferic at each of the Outstations. From the difference between arrival times measured at one Outstation and at each of the others, the ATD System's Control Station estimates a fix for the flash.

The ATD technique can be illustrated by considering a lightning flash which radiates electromagnetic waves received as Sferics at two geographically separated Outstations. The difference in arrival times of Sferics at these Outstations defines a locus of constant ATD over the earth's surface through the flash location. If a Sferic is also received at a third Outstation, a second locus is defined by considering the third Outstation and one of the others. Unfortunately, the two loci intersect at two points, and a fourth Outstation is necessary to resolve this ambiguity and fix the flash at the common intersection area. With four or more Outstations the flash location is over-determined, allowing the consistency of the ATD data to be estimated. In practice more Outstations are operationally necessary to ensure that loci intersect at suitably large angles (for fix accuracy) over the Service Area. The techniques for estimating the fix, the ATD data consistency, and the geographical fix accuracy are described and discussed in Lee (1985). Flash location errors are essentially proportional to ATD errors, but vary with location because of unfavourable geometry at extreme ranges. For example, under ideal conditions of minimally spaced perpendicular loci, a 25μs root mean square (RMS) ATD error gives a 5.3km RMS fix error. Fig. 1, from Lee (1985), shows a possible geographical distribution of seven ATD Outstations, and also the RMS km fix accuracy if a 25μs RMS ATD accuracy is assumed; the reference gives details. This accuracy chart is regarded as a realistic operational target. The trial results of Lee (1985) show that an RMS ATD of 3.5μs is achievable, although this is operationally degraded to around 8μs to achieve economies in communications as described below. This still gives a factor of three safety margin, although this may be temporarily eroded by technical or communications loss of an Outstation, especially Cyprus or Gibralter.

3. Correlogram vs. threshold techniques

Lee(1985) demonstrates that an ATD network based on the detection of vertical electric field Sferics within the 2-23kHz frequency range, together with a correlogram (waveform-matching) method of extracting ATD, can produce an accurate long range lightning flash locating system.

The CRDF system operates at frequencies near 10kHz, where the conducting earth and ionosphere act as a waveguide spreading energy two rather than three dimensionally, giving potential for great range. The ATD system replaces the long range CRDF system, using similar frequencies.

Some commercial lightning flash locating systems operate at higher frequencies where the ionosphere is less reflective, increasing propagation attenuation, and limiting reception range. For these local area systems this ensures that signals from unwanted distant areas are removed, but is disadvantageous for a long range CRDF replacement.

At these higher frequencies the received field may have a sudden onset, allowing the leading edge to be timed for an ATD system, thus avoiding transmission of waveform data. However, at 10kHz the emitted signal is convolved with the resonant impulse response of the earth-ionosphere waveguide, making the received envelope less distinct at long ranges, especially at night. Computer simulations of various "Threshold" (waveform, or filtered waveform, feature timing) types of ATD systems were run using the experimentally gathered 2-23kHz waveform data described in Lee (1985). The results lacked the required reliability, giving substantial "tails" with errors exceeding 25μs in measured ATD distribution, due to timing different features.

Thus it was concluded that any long range system requiring accurate ATD must operate near 10kHz, and must accept waveform-matching and the implied communication of waveforms to the Control Station.

4. Sferics data processing in the ATD system

a. Overview

The operational network will eventually comprise the seven Outstations shown in Fig 1, with precise geographical details given at Lee (1985). Initially, the Outstation shown at Aughton (near Liverpool, UK) is installed at Beaufort Park (near Bracknell) for administrative reasons.

These Outstations are linked via low-speed (up to 110 baud) full duplex communications lines to the Control Station, a powerful minicomputer at Beaufort Park which is expandable to accommodate further overseas Outstations. The Outstations are fully automatic, but a Control Station operator supervises and controls the overall system.

Each microprocessor-based Outstation records the vertical electric waveform of the Sferic from an individual lightning flash. Additionally, each Outstation records the time, more strictly the instant-of-time or epoch, of a waveform sample. In general, each Outstation performs these tasks autonomously, but under the direction of Control Station messages.

Outstations report waveform and epoch data to the Control

Station, which extracts ATD values and fixes the flash. The process of recording and forwarding Sferic data, and obtaining flash fixes, continues at a rate such that most storm areas within the System's Service Area are detected and located. The Control Station combines flash data into meteorological messages for automatic international distribution.

b. Communications

The above communications links are asynchronous serial telegraph lines connected via RS-232 ports and Line Interface Units, initially via land-line to the UK stations, via satellite to Cyprus, and via HF radio link to Gibralter. Some circuits impose delays for Automatic Repeat Request action at this physical level, so that traffic and Control Station processing may run several minutes in arrears of receiving Sferics.

Communications at link level use an HDLC-like bit-oriented protocol, although the bit-stream is packed into asynchronous serial characters. The protocol includes error control through a cyclic redundancy check, so that low link-level error rates are achieved in spite of physical level errors. Error rates are further reduced by additional levels of error control, allowing high reliability in transmitting non-redundant data or software.

Each link-level frame includes an address. Most Control Station frames are for the Outstation's main microprocessor, but some are for an auxilliary link-monitoring processor (Fig 2) which facilitates microprocessor selection or reset from Bracknell, allowing recovery from software problems.

c. Outstation Broad-Band processing

Fig 2 presents a simplified block diagram of the Outstation's microprocessor and peripherals.

The vertical electric field Sferic signals are capacitively coupled to an externally mounted Sensor antenna, inducing currents into an electronic "virtual earth" kept at low impedance so that damp-induced leakage conductance has little effect. The Sensor includes an integrator to give flat frequency response, switchable gain, is designed to withstand the 10-kilovolt signals induced by near-by lightning discharges, and has surge arrestors to divert any heavy currents from actual strikes.

The Sensor's analogue signal is passed to the Analogue

Processing Unit (APU) which conditions it through a 2-18kHz Broad-Band filter (detailed at Appendix A), having anti-alias properties in preparation for sampling. The APU has software-controlled variable gain amplifiers to adjust the equipment to the strength of received signals.

The Broad-Band filter's output is taken to the Analogue to Digital Convertor Board (ADC). This features a multiplexer, here set to accept the Broad-Band signal, a Sample and Hold Unit sampling every 12.8μs (giving a Nyquist limit of 39.0625kHz), followed by a 14-bit Analogue to Digital Convertor module. The digitised amplitudes are presented to the Timing Board, which formats this data for input to the microprocessor.

d. Sferic data identification and storage at the Outstation

The APU also generates an Event signal identifying a Sferic. The Sensor signal is passed through the "Threshold Gain" software controlled variable gain amplifier, and then through a Narrow Band band-pass filter centered on 9.76kHz, designed to select Sferic-like disturbances (Appendix A). The resulting narrow-band signal triggers a re-triggerable monostable at the peak of any positive half-cycle that exceeds a "held" measure of the previous highest peak, or an initial Threshold. The monostable relaxes 205μs after the last trigger, i.e. 205μs after the narrow-band signal's highest positive peak, irrespective of amplitude, provided the Threshold was exceeded. This trailing edge defines the digital Event. The aim is to generate an Event from the same waveform feature at all Outstations, although propagation distortion and added noise limit success. The "held" measure is allowed to decay back to Threshold with a time-constant of 1ms.

The Timing Board contains a 40-bit binary counter driven from an external precision 10MHz "Timekeeping Oscillator". These form a "Local Timescale", whose epoch can be latched and stored without disturbing its continued operation. The Timing Board also contains a "Delay RAM" used to retain the last 512 samples of 14-bit Amplitude from the ADC. In this case the oldest sample retained is 6.55ms old.

When a Sferic Event from the APU activates the Timing Board, it triggers a sequence that latches the Local Timescale value at the next waveform sampling epoch, and passes it to a First In First Out (FIFO) buffer holding data destined for the microprocessor, together with the oldest amplitude in the Delay RAM (which preceeded the Event by 6.55ms). If nothing further happens, the Timing Board will update the Delay RAM, and then dump its oldest Amplitude to

FIFO every sample epoch for the following 1023 samples. Thus the FIFO contains the epoch for a defined sample, and 1024 associated Amplitudes representing a 13.1ms data window, with the Event's waveform in the centre. Should another Event occur before completion, the sequence is re-started with a new epoch and Amplitude passed to FIFO, to be followed by a further 1023 Amplitudes, etc. Thus for any Sferic Event there will always be a 13.1ms window of data surrounding the associated Sferic waveform.

A minor exception is that Events within 128 samples of any previous accepted Event are ignored to avoid saturating downstream processing. Thus the Sferic may be off-centre by 1.6ms, but no significant data is lost.

The Timing Board's data formatting mirrors that of the Transient Recorder in Lee (1985), with the addition of sequence re-starting. Neither Transient Recorder nor CRDF can report all over-Threshold Sferics, so Outstations must select, and some Outstation data may be unavailable for particular flashes. This implies dramatically reduced fix quality, or substantially reduced data availability (Lee, 1985). However, sequence re-starting in this system stores data at all Outstations for all selected flashes, provided the Sferics were all above Threshold. This is encouraged by using one Outstation as a Selector with an artificially low Threshold Gain. Thus the new ATD system uses its Outstations more effectively than CRDF.

Data in the Timing Board's FIFO is removed by the microprocessor and stored in its Random Access Memory (RAM).

e. Software processing of Sferics data

As described above, each Outstation detects Sferics and records those above Threshold into RAM with their Local Epoch.

The average rate of Sferics above Threshold roughly doubles for every 4dB increase in Threshold Gain. Outstations adjust Threshold Gain to maintain their RAM buffers (slowly being emptied on Control Station command) partially full, subject to a Control Station imposed Maximum Threshold Gain. This is adequately high for most Outstations, which operate autonomously maximising the number of largest Sferics recorded. However, one Outstation is designated Selector; the Control Station gives it a low Maximum Threshold Gain, almost completely controlling its actual Threshold Gain, and making it relatively less sensitive.

The Selector Outstation reports each Sferic Event's epoch urgently to the Control Station. This selects an epoch for each Sferic, subject to a minimum intervening period; converts this Local Epoch value to an estimate of International Atomic Time (IAT), and then into Local Epoch values for each of the Outstations. It then requests Sferic data from each near these epochs.

The Control Station operator normally designates the Selector Outstation away from local storms to avoid selections not recorded at all Outstations.

Each Outstation searches its RAM for an appropriate Event and surrounding 13.1ms data window, reports them in a highly compressed form to the Control Station, and then allows this and all earlier data to be overwritten.

The Control Station waits for a time-out period from the Event epoch, to allow for communications delay, then continues processing with whatever data is available, late data being subsequently ignored. This scheme reduces system vulnerability to communications or other technical problems.

The Control Station extracts the ATD, in IAT units, between the Sferic received at the "Reference" Outstation and each of the others by matching waveforms; the degree of mis-match contributing to the estimated ATD variance (Lee, 1985; Appendix C). A further contribution comes from the Local Timescale to IAT conversion variance.

From the set of variance-weighted ATD values, the most likely fix for the flash is estimated iteratively. This technique tests the ATD values, their variances, and fix, highlighting statistically inconsistent data; and estimates location accuracy, so that poor fixes can be discarded (Lee, 1985; Appendix D).

In practice, the Control Station cannot usually fix flash data immediately on time-out because of processing limitations. It therefore keeps its buffers partially full by sending messages modifying the Selector Outstation's Maximum Threshold Gain, effectively modifying sensitivity, adjusting the rate of above-Threshold Sferics, and thus the incoming data rate. This control loop remains stable with ten-minute communications delays in the data rate response to such a message, and reacts in around an hour. Thus the system optimises its resource utilisation, maximising flash fixing rate while selecting the strongest Sferics available. This increases the probability of locating Service Area storms given numerous background Sferics of distant origin. The present processor can fix around 400-500 flashes per hour, depending on the Outstations involved, although this number is reduced if the system is simultaneously used for development or diagnostic tasks.

f. Formatting and output of Sferics messages

Flash locations gathered over an observation period are consolidated into messages distributed via the Meteorological Office automatic communications system AUTOCOM.

The WMO SFLOC message, generated for international distribution, mirrors the CRDF system, giving a 0.5° resolution (degrading accuracy) from 40°W to 40°E and 30°N to 70°N. However, a 30-minute observation window, rather than the CRDF's 10-minute window, is used to increase storm observation probability, and data is reported to AUTOCOM twice per hour within 10 minutes of the end of the observation window, rather than once per hour within 10-50 minutes. This message is archived.

In addition a "High Resolution" (HR) message is generated. This covers an unlimited Service Area (subject to adequate fix accuracy) with a resolution of 0.025°N and 0.05°E. The HR message is updated every 5 minutes, and contains two parts covering 5-minute observation windows terminating 5 and 10 minutes before the update time. The earlier version of an observation window may be less complete or accurate because of long international communications delays.

The HR message is too detailed for an unaided human, and is not encoded to be immediately legible. Instead minimum redundancy encoding with a cyclic redundancy check is used, so that suitable equipment can process and display certificated data. The message is transmitted immediately after updating, and repeated 2.5 minutes later, allowing another decode of corrupted messages. HR messages are absolutely limited to 500 characters to facilitate machine storage or buffering.

5. The economics of communications and operational complexity

A "long range" system approach mandates the use of VLF with its low propagation losses, and the use of expensive waveform-matching techniques to achieve high ATD accuracy. With both long range flash detection and high intrinsic precision, an operational network can be designed to give the long base-lines and fine resolution necessary to cover a wide Service Area with relatively few Outstations. Compared to a system based on potentially cheaper short range Outstations, this reduces capital costs by dramatically reducing the number of Outstations necessary for total coverage, and includes otherwise inaccesible areas. Running costs are reduced as fewer lines incur the excessive communications costs associated with extended distances. In addition, the overall system uses all Outstations to locate each flash (not just those nearest the flash location) so that the system can tolerate temporary Outstation losses without losing part of the Service Area.

In principle, flash location can be achieved with just four Outstations. In practice unfavourable geometry would then limit accuracy in parts of the Service Area, and more Outstations are necessary. The reliability of the network is improved as more Outstations are added; although loss of an Outstation then reduces accuracy, the overall system continues giving graceful degradation rather than catastrophic failure.

Random addition of Outstations may not compensate adequately for particular losses. Long base-lines are provided by overseas Outstations, and to provide compensation more of these are necessary. This is not a trivial matter because of political realities, and because the potential extra communications costs can swamp equipment capital costs in discounted cash flow terms.

To reduce communications costs, each Outstation incorporates a powerful 16/32 bit microprocessor with a megabyte memory capacity. This uses signal processing and data compression techniques which reduce the communications volume of Sferic waveform data by two orders of magnitude. In combination with buffering to smooth traffic flow, this allows the use of 50-110 baud telegraph circuits rather than medium speed data lines, achieving dramatic savings in communications costs if suitable circuits are cheaply available. In practice this limits the choice of Outstations sites, so that network geography cannot be optimised for redundancy through multiplicity of Outstations, and paradoxically the need for fewer Outstations with higher intrinsic precision is emphasised.

Waveform matching, powerful computing capability, and high intrinsic precision, all imply large numbers of failure-prone components. To minimise this, self-calibration and correction techniques are employed to reduce dependence on module stability, allowing simpler hardware designs. To alleviate component failure, detectable by these techniques, controlled redundancy is employed to give graceful degradation at each Outstation, improving availability for some increase in capital cost.

The redundancy techniques are shown in Fig 3. Duplicated microprocessors, selectable from the Control Station, provide flexible control. A RAM with error detection and correction covers most memory faults, while re-locatable software gives further protection. Duplicate Sensor, APU, ADC/Timing/Calibration, and Loran analogue modules are used, with two Timekeeping Oscillators. All are termed "Top Half" or "Bottom Half" to reflect their physical position, while twelve Notch Filters (to remove man-made transmissions) can service either half. Data routes are selected by software controlled switches. Normally, independent halves cover Sferic processing and Epoch Calibration (described below), so that at the designed performance level redundancy implies little extra cost.

The externally mounted Sensors are vulnerable to failure, but either can serve both halves with some compromise in gain. One APU can also handle simultaneous signals by compromising gain and Notch Filter deployment.

Failure of an ADC or Timing Board implies loss of that facility for data acquisition, but the remaining pair can be time-shared with some data loss. The two Timekeeping Oscillators are normally inter-compared to detect clock jumps: losing one merely loses this extra check.

Many component failures thus give only marginally reduced performance, or increased risk of significant failure. Thus the urgency of maintenance response is greatly reduced, offering substantial savings in maintenance costs.

6. Further aspects of system operation

a. Introduction

Section 4 described the basics of Sferic waveform processing. This section highlights further aspects of the operational VLF/ATD system, including interference supression, and the necessary support functions of Spectral Calibration and Correction, and Epoch Calibration.

b. Narrow-bandwidth interference supression and the detection of VLF transmissions

Within the useful 2-18kHz radio-frequency spectrum, the strength of many man-made transmissions are comparable with weaker useable Sferics. Fortunately, for this system, at these low frequencies transmitter antennae are limited to a small fraction af a wavelength. They must form part of a high Q resonator to radiate useful power levels, so transmissions are limited to 50-100Hz, and may be blocked with band-stop (Notch) filters narrow enough to minimise distortion of the remaining Sferic waveform (Lee, 1985; Appendix A).

Each Outstation cascades twelve such filters with software controlled centre frequencies to accommodate the most significant transmissions, which differ with time and place.

Each Notch Filter comprises a frequency selective feedback path optionally connected via software controlled switches (Fig 4; A and B, or C and D) from the output of an inverting unity-gain amplifier within the forward path of the APU, to a summing junction at the amplifier input. With switches open, the feedback path is isolated, giving a level amplifier frequency response. With one set of switches closed, the feedback path monitors signals near the notch centre frequency, integrates these over time, and injects a voltage into the summing junction nulling out the unwanted signal.

This feedback path (Fig 4) multiplies the output signal by a synthesised centre frequency (complex product detection), reducing its frequencies by this amount to produce a complex amplitude with respect to the synthesised reference. The real and imaginary components are integrated and low-pass filtered with seperate circuits to define the Notch profile, and to determine the unwanted APU signal's complex amplitude. The result is multiplied by the synthesised reference (complex amplitude modulation), raising frequencies to the original band, and then injected into the amplifier's summing point to null unwanted signals. This technique gives high stability as the reference is synthesised from the Local Timescale oscillator, tolerance related profile errors depend only on small frequency differences from the reference, and nulling gives intense notch centre attenuation. Such intrinsic stability allows a highly selective Chebychev band-stop profile to be accuarately synthesised (Appendix A).

The Notch Filters' switches provide fail-safe redundancy. Notches are switched in only if required, and are isolated on failure as one of many alternatives can be similarly deployed to either half.

The complex amplitude of the signal removed by a Notch (Real/Imaginary Output, Fig 4) is monitored via two multiplexers on the Calibration Board (Fig 2). These pass the amplitude components through 75Hz low-pass anti-alias filters, and present the results to the ADC Board where a further multiplexer samples alternate amplitudes for the digitiser every 1.6384ms, recording 1.68s of data in 1024 alternating samples. The system may then determine whether any Notch is unnecessary, and should be re-deployed, or has failed and should be isolated. Thus monitoring and flexible control allows the Notch Filter facility to degrade gracefully with component failure, gradually losing the ability to perform the least significant tasks.

Monitoring a Notch's complex Output realises a 150Hz bandwidth VLF receiver. The phase of a Notch's reference can be measured, thus phase-coherent transmissions such as Omega Navaid may be acquired by triggering waveform collection at the correct epoch, coherently integrating to give high signal/noise ratios, and the received phase used for calibration of the Local Timescale.

c. Broad-bandwidth interference

As well as narrow-bandwidth interference, removed as described above, there is broad-band interference due to man-made noise and a background of weaker Sferics.

Sferics from flashes of distant origin are usually weaker, more easily corrupted by additive interference. For a compact network of Outstations, such as those in the UK, the Sferics from such a flash traverse similar propagation paths, and thus exhibit similar waveforms at these Outstations. Thus the ATD extraction process, based on a time-lag correlogram (Lee, 1985), estimates ATD using a near-optimum matched filter technique, minimising the effects of additive noise (Schwartz, 1970). This is particularly important for distant flashes as higher ATD precision is necessary to achieve reasonable location accuracy. For more local flashes the "matched filter" is less effective because of the different propagational paths, but is less important. One may include a propagation model in the ATD extraction process to correct the pair of waveforms involved to some intermediate range, but this requires further processing resources, and is not done here.

d. Spectral calibration and correction

To achieve high ATD accuracy, the spectral characteristics of the APU must be known accurately.

The APU contains a pseudo-random noise generator. At regular intervals this is presented to the APU instead of the Sensor (Fig 2), and the result digitised in the usual way. The Timing Board is triggered by a synchronising signal transferring the digitised filtered waveform to RAM, and eventually to the Control Station. From the filtered waveform's spectrum the Control Station calculates the spectrum of the analogue filters, and hence the "Deviation Spectrum" which describes how the analogue filters differ from their mathematical ideal.

This, combined with previous information, is used to correct Sferic waveforms subsequently received from Outstations, prior to ATD extraction, so that they appear as though received through identical Outstation filters. This technique (Lee,1985) removes the effects on ATD of the inevitable minor Outstation hardware differences. Usually hardware filters are highly stable, so average Group Delays can eventually be corrected to a few tenths of a microsecond, making their contribution to the error budget insignificant.

The Deviation Spectra are displayed to the Control Station operator, alerting him to any serious APU problems.

Similarly, pseudo-noise signals can be applied through the Sensor (Fig 2), with an optional 100 ohm resistor test load from antenna to earth, and a deviation spectrum displayed; or they may be used with other multiplexers on the Calibration Board to calibrate or check all other Outstation analogue filters. Any apparent failure prompts the operator to re-route processing.

e. Epoch calibration

A second important support function is Epoch Calibration: maintenance of the estimated Phase Offset (time error) between the Local Timescale and IAT. Outstation facilities are provided for comparing the Local Timescale with various external IAT estimates, each having some uncertainty.

Phase-stabilised VLF transmissions, such as radio-navigational Omega, are used to monitor Local Timescale drift by measuring their apparent carrier phase against a nominal frequency on the Local Timescale. The 1.68-second data window gathered while monitoring the signal blocked by a Notch Filter easily covers the one-second signal received every ten seconds for any particular Omega transmission. This may be selected by suitable choice of frequency and microprocessor pre-set Event epoch, reducing dependence on any one transmission. In addition, the receiver bandwidth is wide enough to accommodate carrier supressed signals, eg MSK, from which a carrier may be constructed by suitable processing, making further phase-stabilised carriers available. However, interference between ground-wave and multiple sky-wave reception of the transmitted sine-wave makes these uncertain measurements.

The apparent phase-shift of the third zero-crossing of a repetitive pulsed Loran-C navigational signal is measured against a nominal frequency on the Local Timescale. This early portion of the pulse arrives ahead of its sky-waves, so a stable ground-wave signal is measured, although complex signal processing is required to achieve adequate signal/noise ratios. A Loran Board (Fig 2) monitors the Sensor signal early in the APU processing, isolates an 88-112kHz bandwidth, adjusts its amplitude, and multiplies the narrow-band result by sine and cosine of a 100kHz Reference kept in nearly constant phase with IAT by the microprocessor. The two product signals are separately filtered by 11.16kHz anti-alias low-pass filters to produce a complex representation of the Loran signal with respect to the Reference. The two amplitudes are presented to the ADC Board and sampled alternately every 12.8μs. The microprocessor generates Events at pre-set epochs, to select appropriate 1024-sample (13.1ms) data windows each covering the pattern of Loran-C pulses generated by a single transmitter every Loran-C frame, and pass them into RAM. The phase of the 100kHz Reference is monitored by direct sampling every 12.8μs, and phase measurement on the resulting

21.875kHz alias. The processor performs coherent integration of the complex Loran signals to build up signal/noise ratio, and periodically reports the integrated signals to the Control Station, together with the current phase of the Reference. The Control Station reconstructs the Loran waveform, identifies the third zero-crossing, and calculates its phase against a nominal frame repetition frequency on the Local Timescale.

A third type of IAT estimate is based on measuring IAT phase-locked pulses as the opportunity occurs, and is termed "Opportunity Epoch". Examples occur during calibration trips, and might be the nominal 1Hz output of a travelling Atomic Clock, or Loran-C locked square-wave signals from a commercial timekeeping Loran-C receiver. The Local Epochs of such digital signals input as Events are measured, and reported directly to the Control Station. Here the phase of these signals are measured against the phase of their nominal pulse repetition frequency on the Local Timescale.

All these methods report phase as a function of time to the Control Station. These phases are converted to modulo time and are displayed as a function of time to the Control Station operator. The display used can remove rates of time-shift due to assumed errors in the Timekeeping Oscillator's frequency, allowing the operator to view several days' modulo time data with a resolution of the order of 0.1μs. Thus he can interpret these results in terms of the known problems and ambiguities of each type of measurement, and so estimate Phase Offsets and their likely errors at certain well-defined Local Epoch values.

At the Control Station a model is kept for each Outstation Timekeeping Oscillator, valid for some Local Epoch, estimating its Phase Offset, and two derivatives: the Phase Error Rate (frequency error) and Ageing Rate (rate of change of frequency). An estimate of the current uncertainty of each of these quantities is also maintained. The oscillator's frequency is a function of its environment (voltage, temperature, pressure etc), and the coefficients relating these measured and reported variables to oscillator frequency are also estimated, together with their uncertainties. Thus from an initial state at some epoch, knowing the timekeeping oscillator's environment, the Phase Offset can be predicted by numerical integration with increasing uncertainty as time passes without any new measurements being incorporated.

At intervals the Control Station operator will obtain some useful Phase Offset measurements. These are presented to a Kalman Filter algorithm which combines the measurements with predicted Phase Offset, and uses the difference between the two to make small corrections to the various estimated parameters and their uncertainties. Over a long time a series of such updates improves the model, reducing the uncertainties, and enabling the Phase Offset to be accurately predicted in the absence of measurements for extended periods. Thus Local Epoch values can be accurately and robustly converted to IAT, giving accurate ATD values.

Using high quality Loran measurements, together with careful travelling clock site calibration, IAT conversion can be held to a few tenths of a microsecond with minimum operator maintenance, making its contribution to the error budget insignificant.

f. System development

Both Control Station and Outstation designs facilitate further development. Modified Outstation software can be downloaded, while the Control Station supports development and parallel testing of alternative Control Station software.

A Diagnostic Archive facility periodically dumps the Control Station's software state, followed by subsequent external transactions. This is stored on magnetic tape for several weeks. A Replay facility can replay selected portions of the Diagnostic Archive through the current software suite, or a modified update. These facilities allow examination of unexpected results, and testing of system modifications.

7. Sferics data compression

a. Introduction

Sferics are captured as 1024 14-bit samples. In this form they are unsuitable for telegraph transmission, so data compression is applied. The details are a compromise between Outstation processing resources, available communications bandwidth, required flash rate, and the acceptable extent of accuracy degredation. This compromise is subject to operational "fine tuning", but an initial implementation is described below.

b. Bandwidth reduction

Sferic signals exhibit mean spectra which peak near 10kHz. Data can be reduced by selecting a reduced frequency band here, and sampling less frequently in the time domain.

As bandwidth is reduced, ATD accuracy is initially increased (indicated by reduced time-dimensioned RESIDUAL; Lee,1985) as regions of low signal/noise are removed. Further bandwidth reduction makes ATD vulnerable to effects near skip-zones where the received narrow-band group delay may be a non-uniform function of range.

The signal is expressed in band-limited complex form by complex Reference multiplication, and low-pass filtering the resulting real and imaginary component amplitudes. The Reference used has exactly eight 12.8μs samples per cycle, an approximate 9.76kHz. The implied 2048 multiplications are by 0, $\pm 2^{1/2}$, or ± 2, minimising computation.

The ideal low-pass filter would transfer an absolutely limited bandwidth, so that the subsequent sampling rate could be reduced to twice this Nyquist limit. Its impulse response would exhibit minimal time-domain ripple to avoid extending the signal envelope, and increasing the number of significant samples. Suitable optimum "Windows", with time and frequency domains exchanged, are discussed by Harris(1978). Unfortunately these imply computationally expensive processing, so a Bessel filter was recursively implemented in cascaded stages (Lee,1981) as an initial compromise. This has no sharp cutoff, but can have adequate anti-alias properties provided that reduced 3dB bandwidths are acceptable. For a total half-bandwidth of 4.88kHz, allowing complex sampling every 102.4μs, the 3dB half-bandwidth is then limited to approximately 1.8kHz, centered on 9.76kHz. This represents a 13.1ms data window in 128 complex samples.

c. Waveform description

A Sferic's shape contains the information for ATD extraction; amplitude is irrelevant. Thus Sferics are normalised and represented in fewer bits: 14-bit digitisation merely enhances dynamic range. When ATD is extracted, corresponding features of similar waveforms are multiplied together and summed in a correlogram peak. Thus each contribution is essentially weighted according to its amplitude, suggesting use of unequal digitisation levels (denser at the peaks) as noise in the lower amplitudes has reduced effect. Finally, individual disturbances often have short duration within the data window, although pre-cursor electrical effects from charge movement in clouds a few ms before actual Return Strokes must be represented, as must multiple Sferics. The waveform encoding technique utilises this sporadic activity.

The complex samples are normalised to define the maximum absolute component value, then conventionally digitised to 6 bits. Lower amplitude levels are then combined in groups of 2, 4, and one group of 7 (about the mean level), reducing the number of amplitude descriptors by 30%. This coarse digitisation removes most low-amplitude residual transmitter or system noise. Finally, isolated single-level spikes above "zero" are suppressed.

The above techniques frequently produce activity isolated by "zero" sequences. To utilise this, the complex waveform is shifted left to eliminate leading "zero"s (epoch is adjusted accordingly); and run-length descriptors for binary power runs are introduced.

d. Huffman encoding

To quantify the above processing, a training set of "typical" data is required, and an algorithm to encode waveforms into descriptors. ATD accuracy is affected by bandwidth reduction and coarse digitisation, while the descriptor sequence is also affected by the alphabet of descriptors and the encoding algorithm.

Sferics Trial data (Lee,1985) was experimentally compressed by variations of these techniques to minimise the mean entropy of descriptor-encoded Sferics (in an Information Theory sense), while retaining acceptable ATD accuracy. The ATD error introduced by data compression is difficult to quantify precisely from this flawed experimental data, but was approximately $8\mu s$ for the techniques initially implemented. For a realistically calibrated operational system, this dominates the experimentally measured (Lee,1985) contribution for uncorrected propagation effects or epoch/spectral calibration errors, while remaining well within the 25μs "target" accuracy for a CRDF system replacement. Actual ATD accuracy and flash location error can be monitored operationally, and the encoding technique adjusted if necessary.

Each descriptor has its probability determined, and is assigned a minimum-redundancy code according to Huffman's(1952) technique, the most frequent having shorter bit-pattern codes. On average, 75-100 bits describe a 13.1ms data window in sufficient detail to allow ATD extraction to 8μs: a useful reduction from 14336 bits. This allows adequate flash fixing rates to be achieved with 50-baud communications.

8. Conclusions

This paper describes an operational system based on the Trial techniques of Lee(1985) with the developments necessary for robust operation.

The potential ATD accuracy is higher than the Trial results as spectral and epoch calibration techniques have been refined to reduce their error contributions to negligable levels. However, the Sferics encoding technique degrades accuracy to reduce communications costs, giving an overall ATD error of around 8μs. For flashes reported from all Outstations, this still improves Fig 1's location accuracy "target" three-fold. If necessary, some operational "fine tuning" may allow further improvement.

The Service Area reported internationally every half-hour covers 40°W to 40°E, 30°N to 70°N, with a 0.5° resolution. A wider Service Area with higher temporal and spatial resolution is also available.

The system is capable of locating flashes at rates of around 400 flashes/hour, ensuring that most storm areas within the limited Service Area will be located.

This new system should replace the current operational CRDF system formally by the second quarter of 1986.

APPENDIX A

Outstation Filter Profiles

The spectral profiles of the Outstation analogue filters were chosen to give adequate anti-alias and bandwidth properties, and sufficient stability with simple active implementation.

The Analogue Processing Unit (APU) contains high and low-pass filters, the separate cascaded stages being interleaved. The high-pass is a 6-Pole Butterworth filter with 3dB high-pass limit at 2kHz. This is cascaded with a 0.28dB ripple low-pass Chebychev filter, with 17.75kHz ripple limit; giving 60dB attenuation at 39.0625kHz, the Nyquist limit for 12.8μs sampling.

The APU also contains a Narrow Band filter for selecting Sferics. This is a 3-Pole Bessel prototype, geometrically transformed to a band-pass centered on 9.76kHz, with 3dB full-bandwidth of 2.5kHz.

The Notch Filters each have an arithmetically symmetrical profile about their centre frequency, based on a 5-Pole 0.1dB Chebychev curve (Lee, 1985; Appendix A). Either side of the center frequency, this exhibits a nominal 0.1dB ripple in the pass-band outside 96.45Hz; while the stop-band attenuation is 3.0dB at 85Hz, and at least 33dB within 50Hz.

The Calibration Board contains two 6-Pole low-pass 2dB Chebychev filters with 75.0Hz ripple limit, used for monitoring the signal blocked by Notch Filters. At 152.6Hz, the Nyquist limit when sampling each filter every 3.28ms, their attenuation exceeds 60dB.

The Loran Board contains two filter types. The first is a 3-Pole 2dB low-pass Chebychev prototype, geometrically transformed to a band-pass filter having a 2dB band-pass limit from 88-112kHz. This band-limiting allows subsequent limiter techniques to control the in-bandwidth signal level. The second is a pair of 6-Pole 1.25dB Chebychev low-pass filters, with 11.16kHz ripple limit, used for complex band-limiting. Group Delay effects in both types are removed by software techniques.

REFERENCES

Harris, F. J., 1978: On the use of windows for harmonic analysis with the discrete fourier transform, Proc. Inst. Elect. Electron. Eng., 66, 51-83.

Huffman, D. A., 1952: A method for the construction of minimum-redundancy codes, Proc. Inst. Radio Eng., 40, 1098-1101

Lee, A. C. L., 1981: Smoothing and filtering of meteorological data, Lond. Meteor. Mag., 110, 115-132.

Lee A. C. L., 1985: An experimental study of the remote location of lightning flashes using a VLF arrival time difference technique, Lond. Quart. J. Roy. Meteor. Soc., (submitted for publication).

Maidens, A. L., 1953: Methods of synchronising the observations of a 'sferics' network, Lond. Meteor. Mag., 82, 267-270.

Schwartz, M., 1970: Information transmission, modulation, and noise, McGraw-Hill Inc.

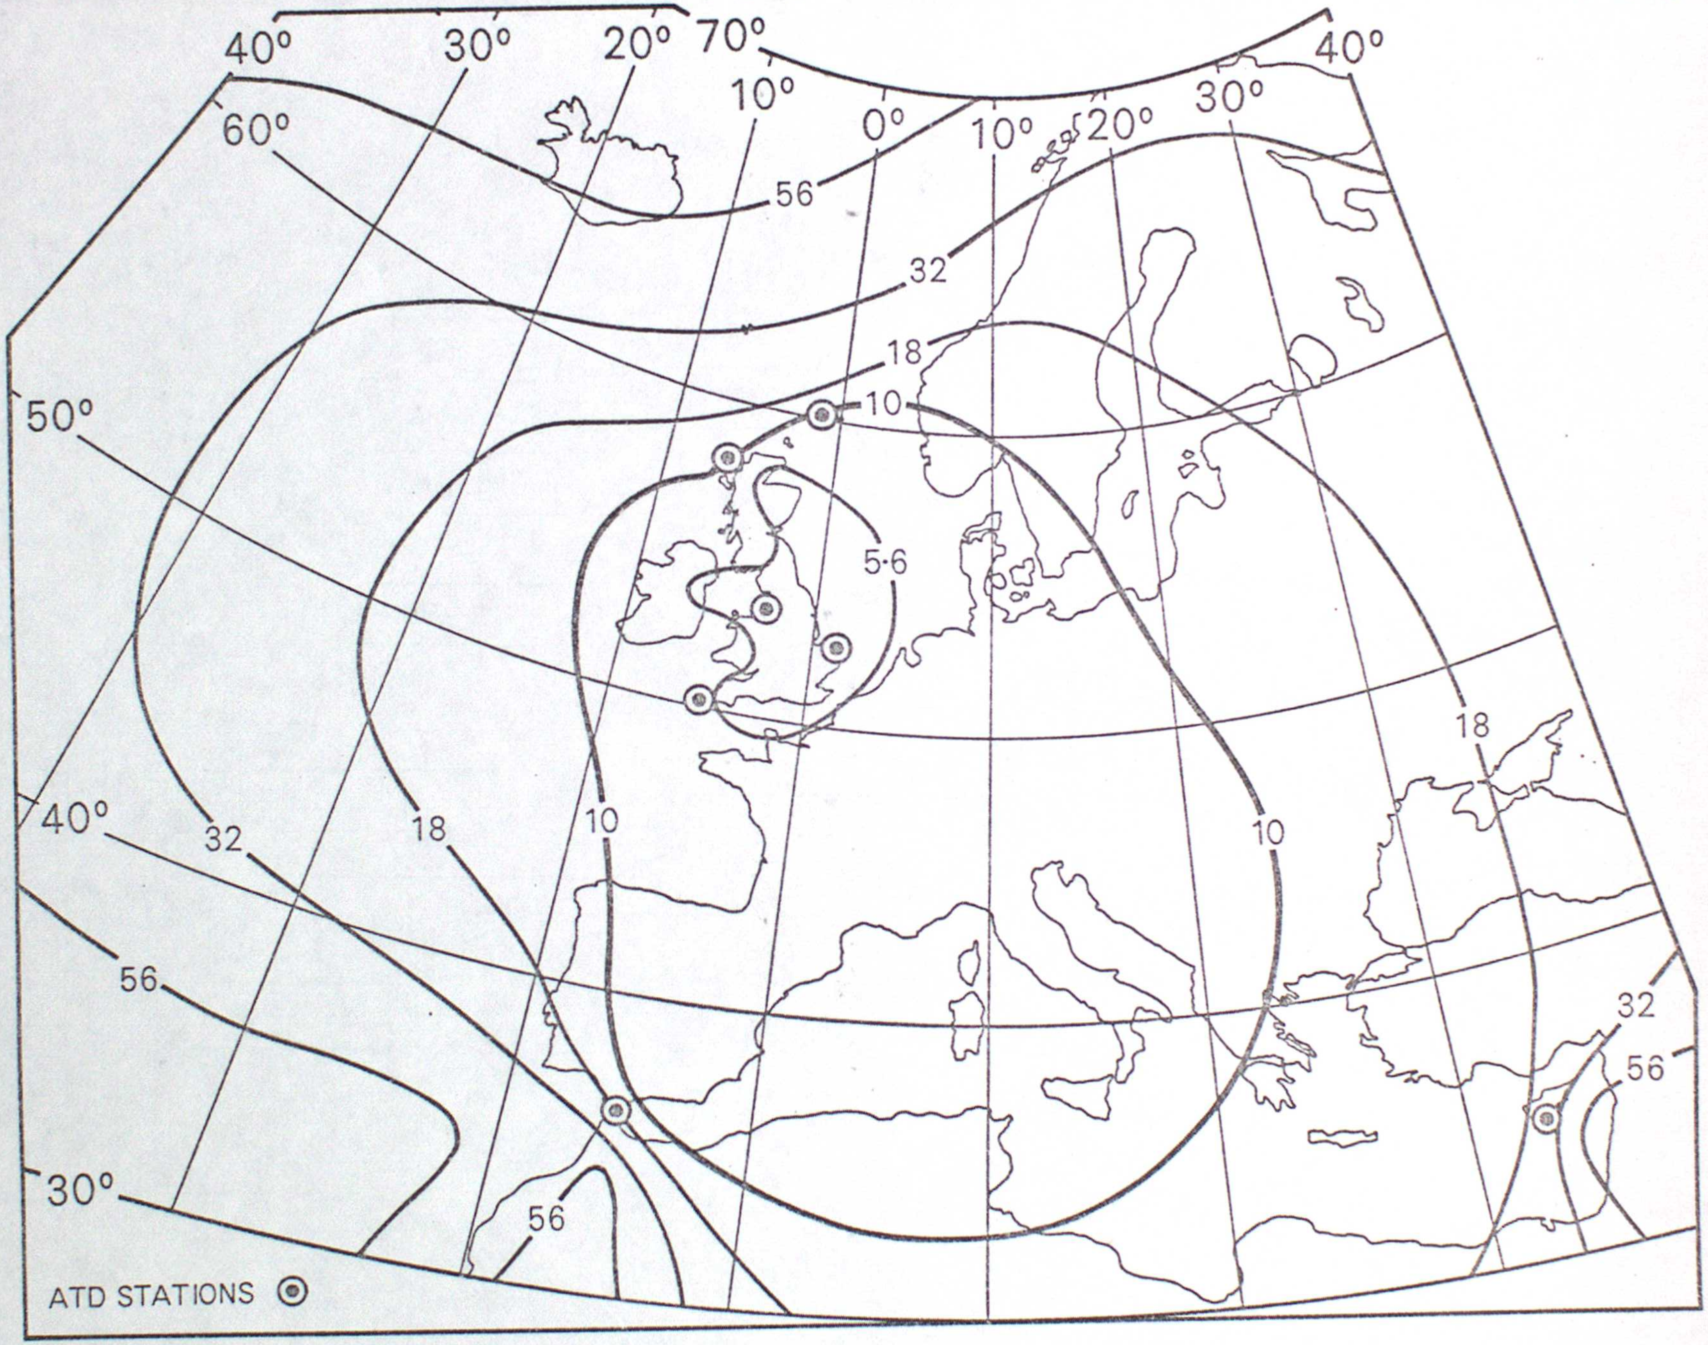

Fig 1: The accuracy in RMS km of an operational ATD network for an RMS ATD error of 25µs.

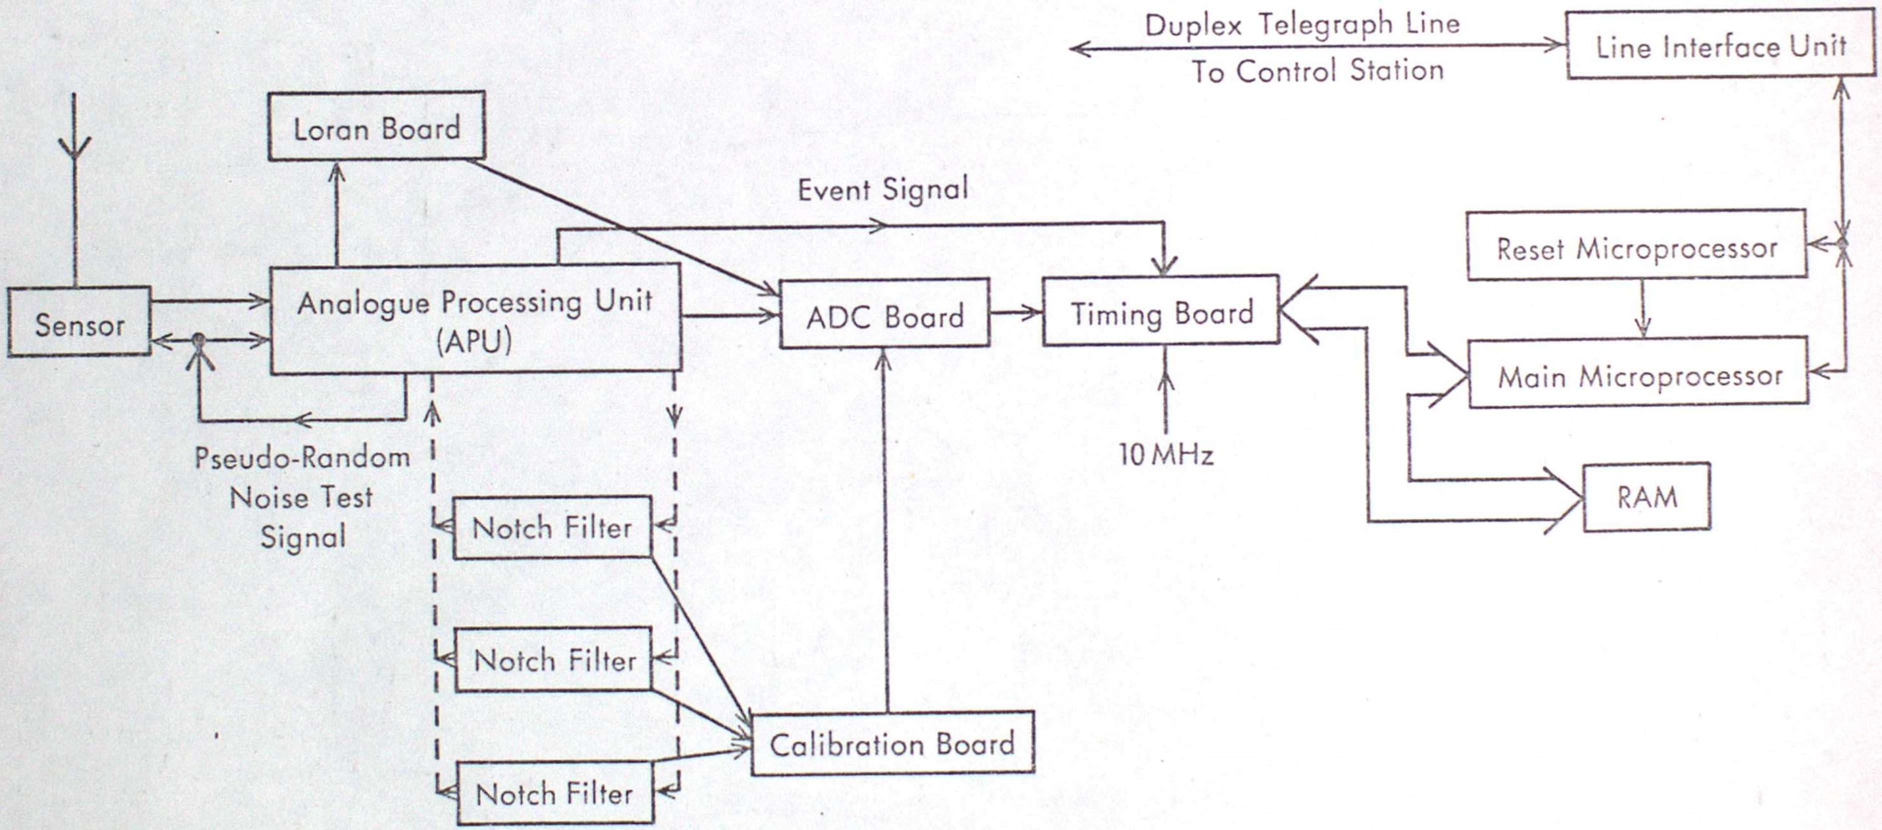

Fig 2: Simplified Outstation block diagram.

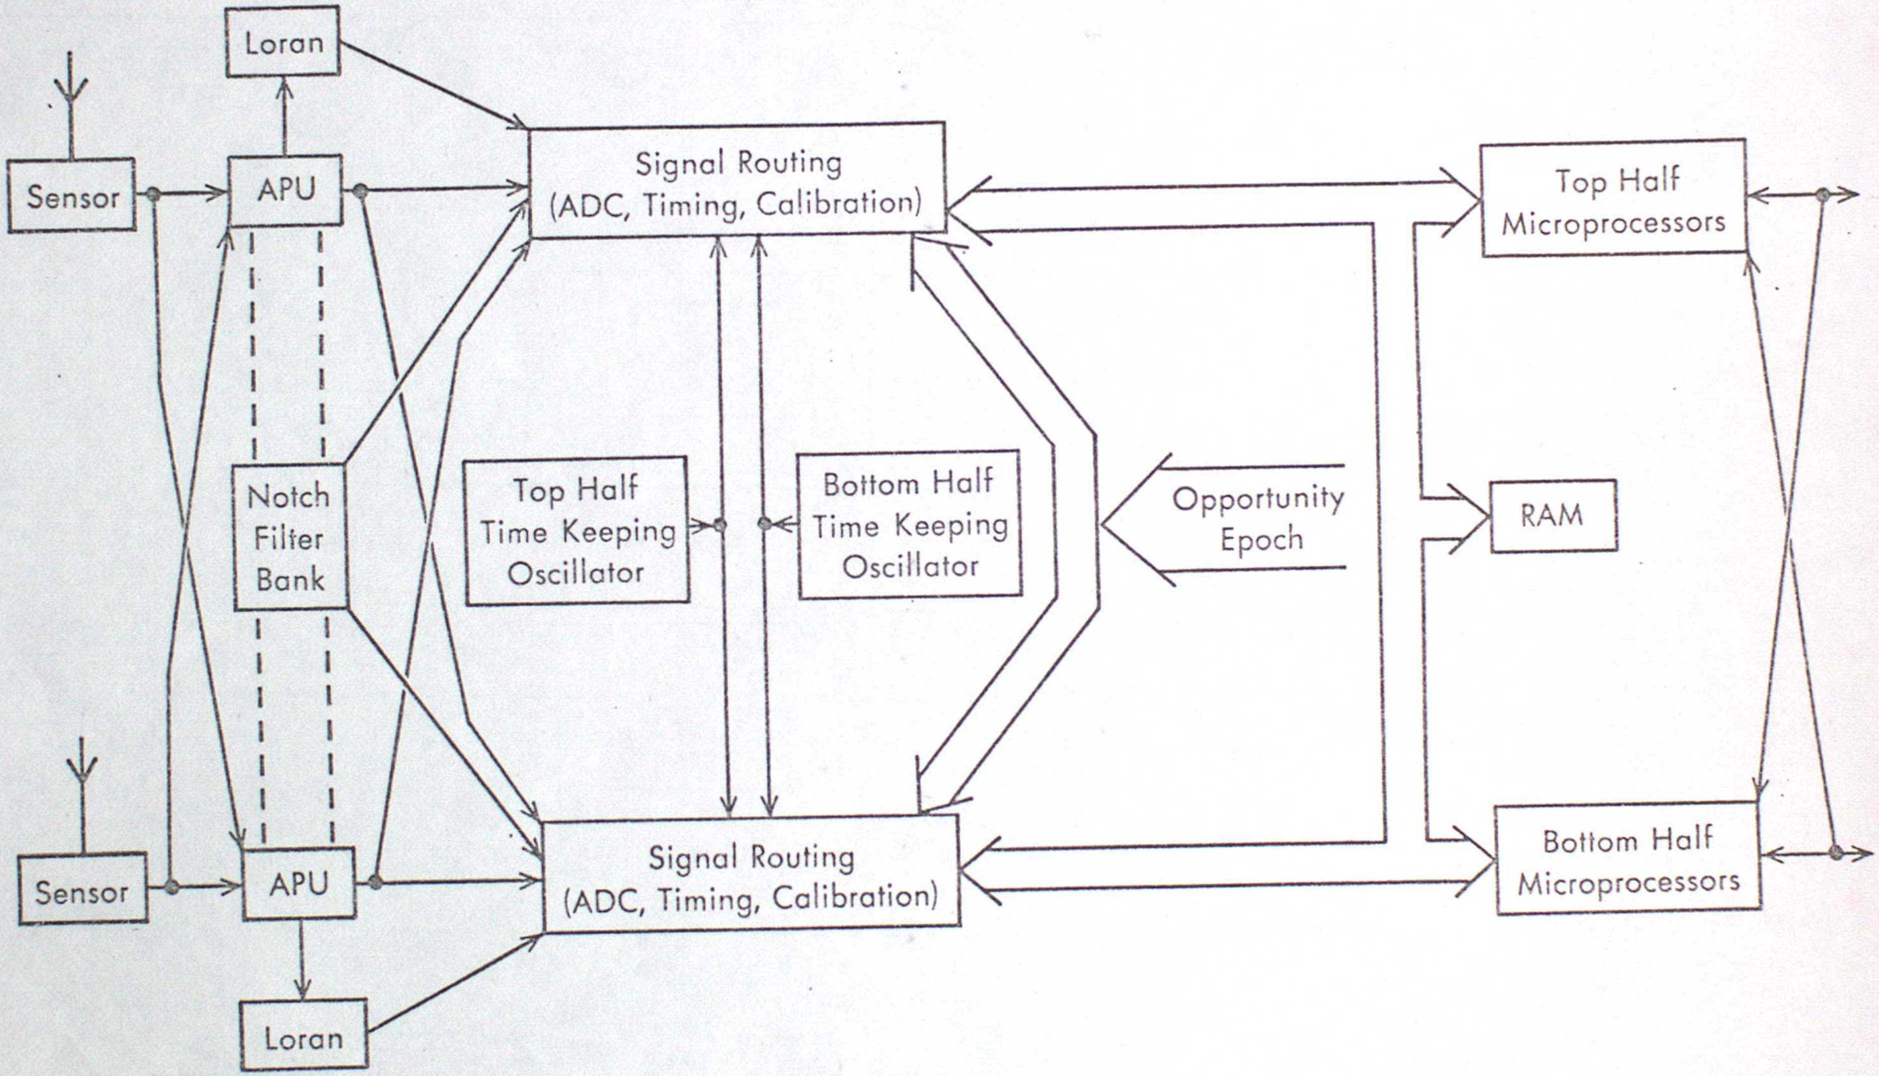

Fig 3: Outstation block diagram: highlighting redundancy.

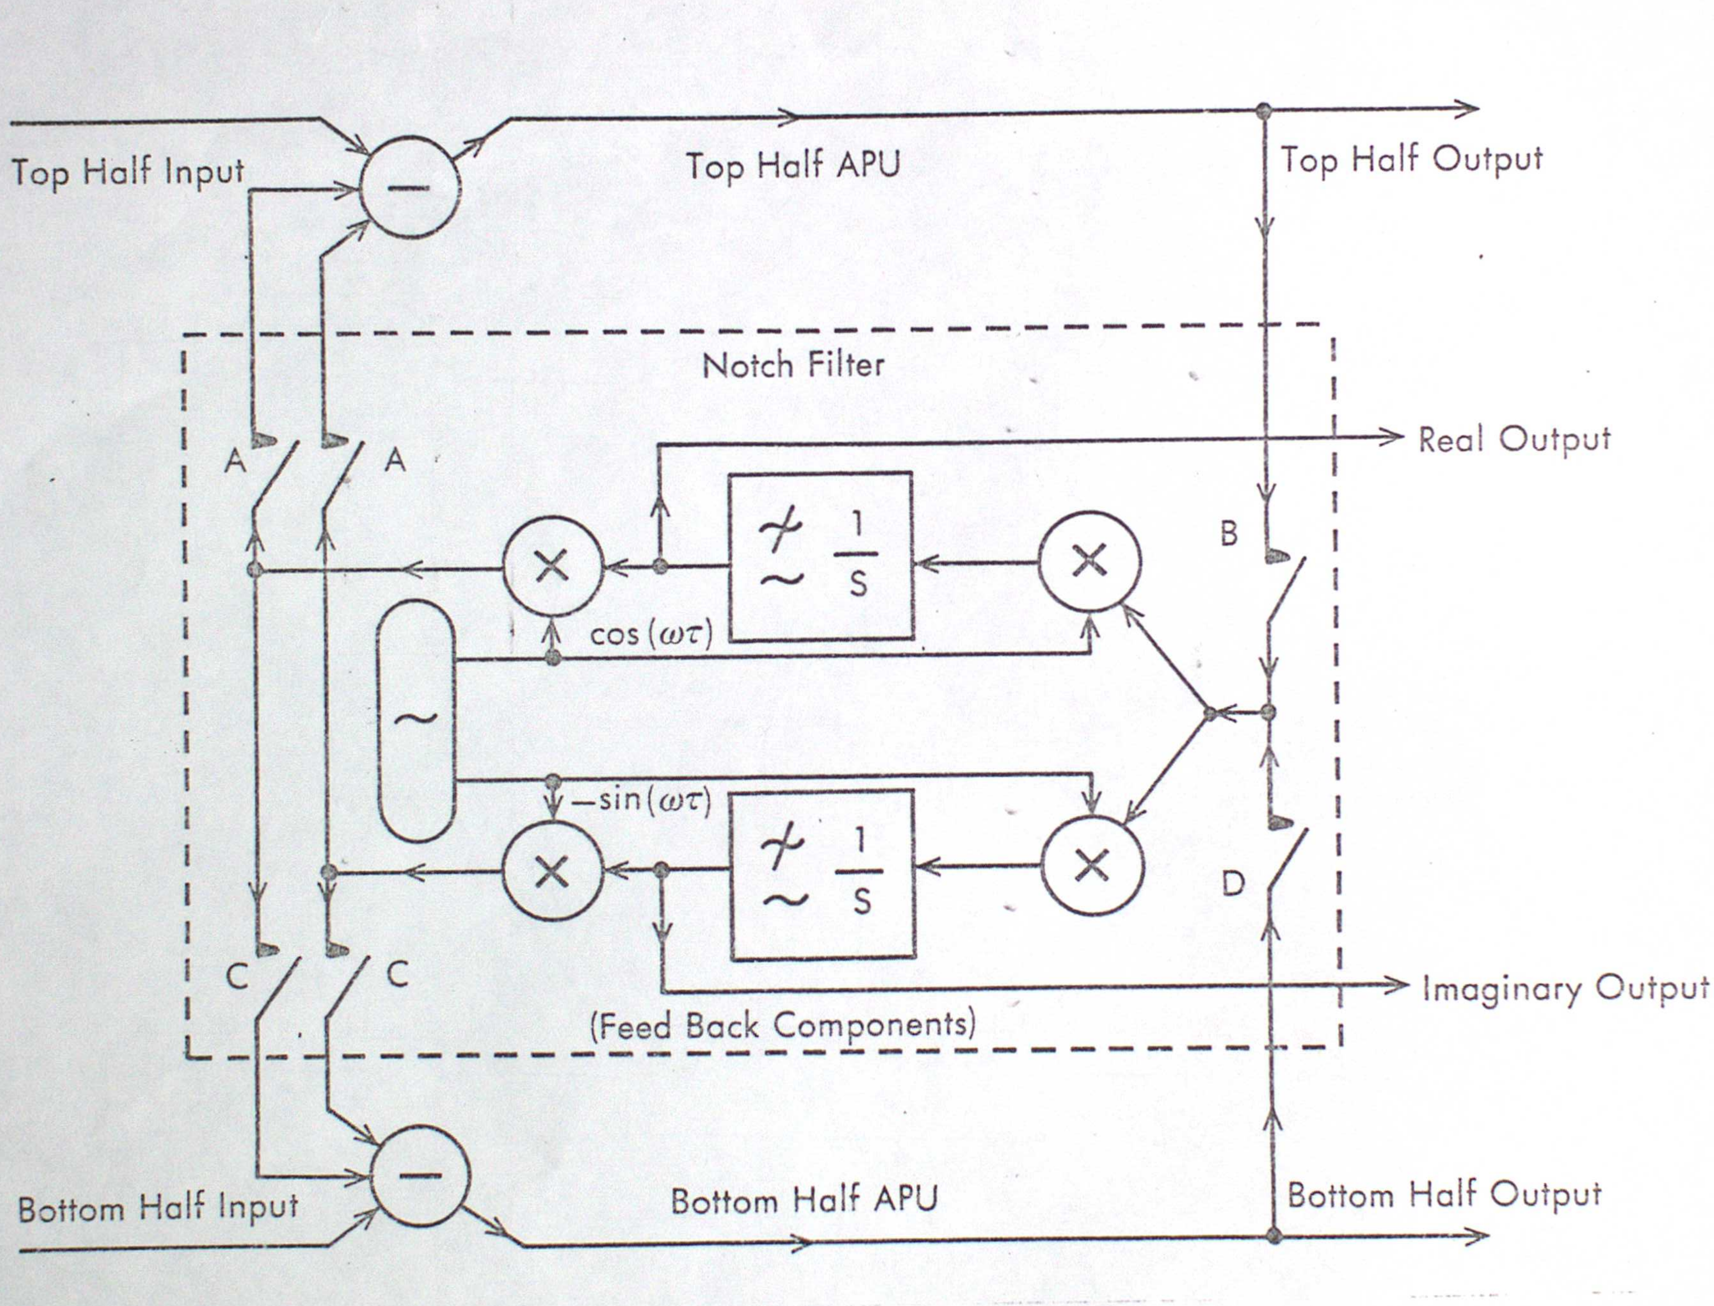

Fig 4: "Notch Filter" feedback components, and their integration into the Analogue Processing Units.