How reliable are acoustic rain sensors?

G. D. Quartly[1], J. W. Gregory[1,2], T. H. Guymer[1], K. G. Birch[1], D. W. Jones[3], S. J. Keogh[3]

[1] Southampton Oceanography Centre, Empress Dock, Southampton, Hants, SO14 3ZH, UK.
gdq@soc.soton.ac.uk, jwg@soc.soton.ac.uk, thg@soc.soton.ac.uk, kgb@soc.soton.ac.uk
[2] Also at Department of Physics, University of Surrey, Guildford, Surrey, GU2 7XH, UK.
[3] Met Office, Beaufort Park, Easthampstead, Wokingham, Berks, RG40 3DN, UK.
dwjones@meto.gov.uk, skeogh@meto.gov.uk

Abstract

Underwater acoustic sensors have been identified as one of the possible methods of monitoring rainfall in deep ocean locations outside the range of land-based radars. In this paper we assess the reliability and accuracy of such a system, using a deep Scottish loch, where many ancillary data are available. At present, the sensors have some difficulty discriminating between the sounds of wind, rain and other activity, but show good quantitative agreement in rainfall totals with other instruments when the acoustic contamination is absent.

1. Introduction

Precipitation at sea is one of the important factors controlling the freshwater budget of surface waters, and thus their density. However it is very difficult to measure. This is, first, because of its small spatial and temporal scales (necessitating fine sampling); and, second, because of the difficulty of making accurate long-term measurements. Satellites are essential for global monitoring. However there is still a need for *in situ* sensors to resolve the diurnal variation at a location, and also to provide validation of the products from satellite data. One of the most promising *in situ* technologies is via inversion of the underwater acoustic spectrum. The use of a sub-surface hydrophone removes the problem of operating at the hostile atmosphere-ocean interface, and provides a mechanism for obtaining areal averages of rainfall. Potentially, very deep hydrophones could yield average rain rates for a disc a kilometre across, which is more appropriate than point measurements to the oceanographic need.

There are many sources of underwater noise. Wind and rain are the principal environmental ones, but for particular locations shipping and various forms of sea life may be key contributors. Each source tends to have its characteristic spectrum. For example, wind produces breaking wavelets and spray, resulting in a broad-band effect, with a typical spectral decay of -19 dB/decade [1]. In contrast, the small drops in drizzle produce a peak near 14 kHz [2, 3], and the large drops in heavy rain augment the sound intensity for all the frequencies between 4 and 21 kHz [4], producing a plateau in the spectrum. This description of the environmental sources of underwater sound is complicated somewhat by winds, which are able to modulate the signature of drizzle [5]. However, it should be possible to use multi-frequency analysis to separate the contributions due to wind and/or rain from those of other sources, and hence to infer the local meteorological conditions.

2. Experimental set-up

2.1 Location

The west coast of Scotland is generally one of the wettest parts of the UK. The results presented here are from a trial of underwater acoustic sensors in May 2000 in Loch Etive. This location was chosen as it provides a deep saline environment (similar to open ocean locations, albeit with much smaller waves owing to the limited fetch), yet has reasonable opportunities for deploying ancillary instrumentation nearby. We were not granted access to neighbouring land, but were instead able to make use of a large moored raft belonging to a local mussel farmer.

2.2 Acoustic systems

The equipment we are using is an Acoustic Rain Gauge (ARG) produced by *Metocean Ltd of Nova Scotia*, which samples the ambient spectrum in 16 bands spanning from 500 Hz to 50 kHz. The system consists of a small surface package (containing temperature and pressure sensors plus logger and satellite relay), with a hydrophone suspended many metres below via a thin cable. We have adapted the system to record the spectrum on a logger every 1.5 minutes (rather than relying solely on occasional satellite data relay) and increased the hull size to incorporate more batteries (to extend the duration of deployments). The buoy was deployed in 50 m of water (with the hydrophone 20 m below the surface), using a dual-tether arrangement to avoid the hydrophone cable being wrapped around the mooring line [6]. As well as recording the sound level in the 16 channels, the onboard processor also provides parameter estimates derived from the acoustic signal. We also had two WOTAN (Weather Observation Through Ambient Noise) buoys on site, but they are not discussed here.

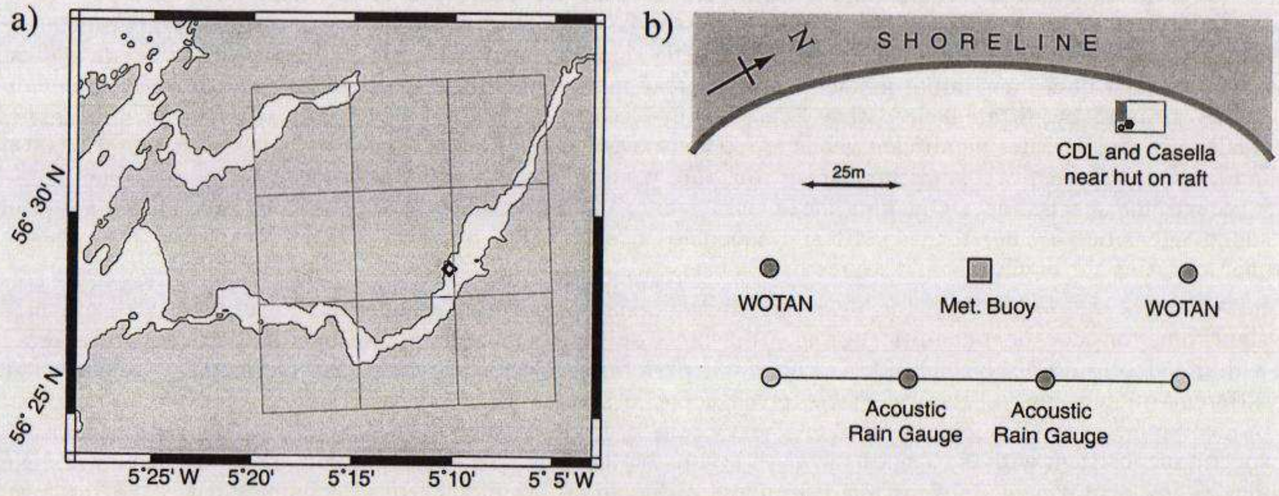

Figure 1. a) Location of Loch Etive, a saltwater loch in southwest Scotland; the diamond marks the mooring site, and the pattern of 5 km squares shows the grid of the NIMROD (rain radar) product. b) Mooring arrangement — only one of the Acoustic Rain Gauges worked in the trial in May 2000, and the WOTAN data are not discussed here. (CDL stands for Climate Data Logger and Casella is a tipping bucket gauge.)

2.3 Acoustic discrimination and parameter estimation

Figure 2a shows some of the acoustic spectra recorded during this deployment. In the majority of spectra there is a fall off in intensity with increasing frequency. The grey lines show the spectral shape for wind-only conditions (with higher wind speeds creating higher sound levels across the whole range). Vagle *et al.* [1] noted the spectral slope to be -19 dB decade^{-1}, whereas Nystuen and Selsor [7] found a value of around -14 dB decade^{-1} for their work with drifting buoys. For our data, we get a mean spectral slope of -18.5 dB decade^{-1} when the whole frequency range is considered; however for the lower frequencies (500 Hz to 10 kHz) the value is just -14 dB decade^{-1}.

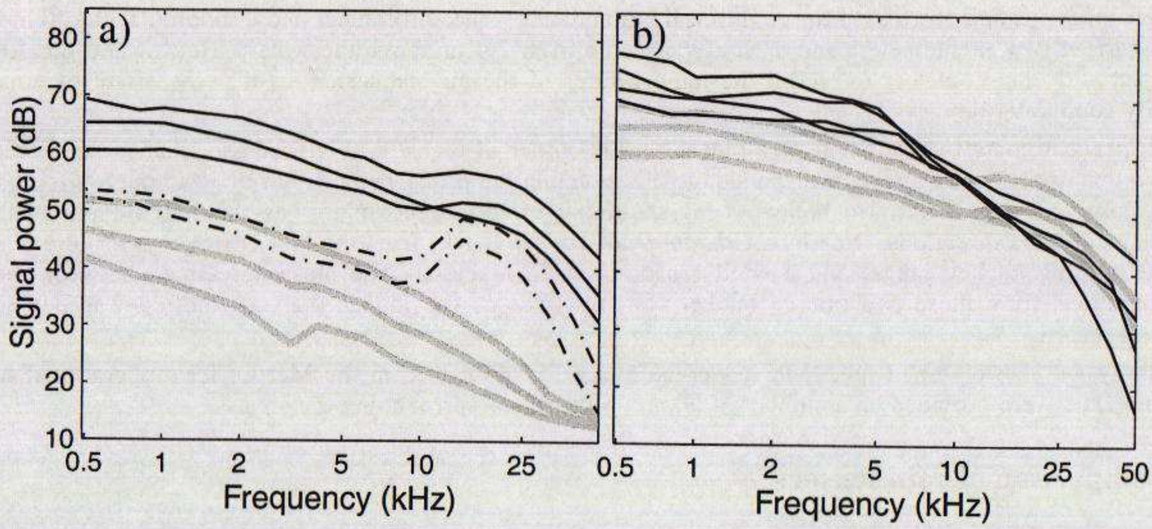

Figure 2. a) Spectral levels caused by different environmental sources: wind-only (grey), drizzle (black dash-dot) and heavy rain (solid black lines). b) Spectra of false detections of rain (black) compared to genuine signals of heavy rain (grey). [Signal power is in dB re $1\mu Pa^2$ / Hz.]

The dash-dotted lines show typical spectra for drizzle, where the spectral shape below 10 kHz is as for wind-only, but there are pronounced peaks near 14 kHz due to the ringing of small bubbles produced by the 0.8-1.1 mm diameter drops present in drizzle. (Although wind causes wave-breaking, and the generation of droplets of spray, its acoustic signature is very different from rain because the drops are at nowhere near terminal velocity.) Medium-sized drops (1.1-2.2 mm in diameter) are too large to generate bubbles in the same manner as the small drops do, and consequently produce minimal acoustic signal. However, larger drops than these create significant sound through both their initial impacts and the creation and ringing of many large bubbles. (Examples of the energy radiated by different drop sizes is shown in Nystuen [8].) Thus heavy rain, containing both large and small raindrops, creates significant acoustic intensity over all the frequencies observed (see the solid black lines in Figure 2a). Heavy rain is distinguished from pure wind in that, not only does it produce much louder noise, but it reduces the spectral decay, making the spectra 'whiter'. Nystuen and Selsor [7] note that winds cannot produce sound intensities greater than 45 dB at frequencies over 20 kHz, and this is used as part of the process for characterizing the environmental source of an observed signal.

Although the technology for detecting and measuring underwater acoustic spectra is well-proven, the algorithms for acoustical classification are still under development. The *Metocean* ARG that we use is provided with the algorithms of Nystuen [8]. In short, this performs various comparisons of the intensities observed at the different frequencies, to classify the received acoustic spectra as 'wind only', 'high seas' (that is, high winds producing many small bubbles), 'drizzle', 'heavy rain' or 'contaminated'. This last classification encompasses all spectra inconsistent with the shapes shown in Figure 2a; these may be due to shipping or biological noise. During the 30-day deployment 8.5% of the 1.5-minute observations were flagged as contaminated. The instances of contaminated data are clumped together (probably indicating some extra noise source, such as shipping activity): when analysis is in 15-minute periods, 4.8% lack any valid observations.

In cases where only wind is believed to be present, an estimate of the wind speed is made from the intensity in the 8 kHz channel, using the algorithm of Vagle *et al.* [1]. The minimum wind speed detectable is 2 ms^{-1}, because lower wind speeds produce no wave-breaking capable of producing acoustic signals. If the signal is due to 'drizzle', then a rain rate of 1 mm hr^{-1} is assigned, because no quantitative inversion is possible because of the effects of wind [5]. For 'heavy rain', two algorithms are proposed. The first uses only the signal level at 5 kHz, whilst the second takes advantage of the extra information present in all the other frequencies. This latter method attempts to infer the drop size distribution (DSD) i.e. the relative proportions of 'small', 'medium', 'large' and 'very large' drops, and calculate a rain rate from this. The coefficients for such an inversion have been tuned empirically from observations in a large brackish pond near Miami [8], where it was noted that medium-sized raindrops produced more sound than expected from laboratory experiments. This complete inversion using all 16 channels is still at the experimental stage, but its estimates of DSD have shown good agreement with those from mechanical disdrometers.

2.4 Ancillary instrumentation

Four other systems were used to provide validation data. A meteorological buoy, moored about 30 m from the ARG, carried two anemometers and a wind vane on a mast 2.4 m above the water. Both it and the ARG were equipped with thermometers to record the temperatures of the air and water, so that the effect of atmospheric stability could later be investigated.

A *Casella* tipping bucket gauge was placed on the corner of the raft. We used the smallest available bucket size, so that each tipping event corresponded to an accumulation of 0.1 mm of rainfall. Although this was located quite close to a small hut, we believe wind-sheltering effects to be minor, as the hut did not obstruct the predominant wind directions. Nearby a *Climate Data Logger* (CDL) from the Met Office was installed. This had another tipping bucket gauge, and a wind sensor on a 10 m mast. The close agreement between the hourly accumulations from these two tipping bucket systems gives us confidence that raft motion was not producing spurious records.

Finally we also obtained rain radar coverage every 15 minutes from the Met Office's operational *NIMROD* system. Data were provided on a number of 5 km x 5 km pixels surrounding the area.

3. Comparison of rain sensors

3.1 Problems in acoustic classification

A simple comparison of the ARG's acoustic classification with tips of the *Casella* shows there is often a good correlation (left-hand side of Figure 3). However there are also periods of no agreement (right-hand side). The first result confirms that the timing on the two systems is synchronized, that there is a good correlation between rain events at the ARG location and at the raft (60 m away), and that the classification algorithm is sensitive to light drizzle. The second result shows that there are strong acoustic features, apparently not associated with rain, yet classified as though they were. This occurs during a period when many of the other records are flagged as containing noticeable acoustic contamination. Inspection of the full time series shows that most of these spurious events occur between 09:00 and 15:00 GMT. This has been noted in previous deployments in November 1999 [9].

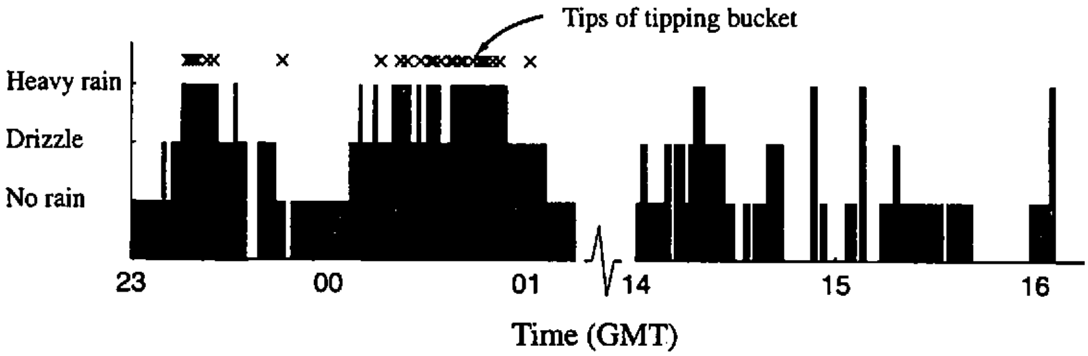

Figure 3. Comparison of acoustic classification by ARG (1.5-minute resolution) with events detected by the Casella (accumulations of 0.1 mm rainfall). Two periods from 26th May are shown, with bars missing when the acoustic spectrum was classified as 'contaminated'.

3.2 Potential improvement to flagging of data

We have examined the spurious inferences of rain in more detail. Neither the Casella nor NIMROD showed any evidence of rain in Loch Etive during the first half of May 2000. Neither was there any acoustic signal of rain during those nights, but the daytime results showed 40 individual 1.5-minute spectra classified as 'heavy rain'. In fact the signal levels were so high that the inferred rain rates were of the order 100-200 mm hr^{-1}, thus each spurious signal was contributing 3 to 5 mm of rainfall to the hourly total. Four of these strange spectra are contrasted in Figure 2b with the signatures of genuine 'heavy rain' events.

The spurious spectra all have a rapid fall off above 10 kHz, which could possibly be associated with intense bubble clouds in high seas. However the independent wind speed measurements from the meteorological buoy only showed values in the range 0 to 8 ms^{-1}. Also the air temperatures were greater than the sea temperatures by upto 6°C, whereas 'heavy rain' and 'high seas' tend to be associated with negligible air-sea temperature difference, because of the intense mixing occurring. 'Distant shipping' can be ruled out as there are no large vessels on this stretch of water, and attenuation of the signal from distant vessels leads to much more reddening of the spectrum. Thus we tentatively infer that the cause must be either local boat activity or the nearby colony of seals.

We compared our data to the tests advocated by Nystuen and Selsor [7]. As the spurious spectra have a very fast spectral decay compared to genuine 'heavy rain' events, we perform a cluster analysis of the intensities at 5 and 25 kHz (see Figure 4). Nearly all (38 out of 40) of the spurious signals lie to the right of the dashed line, whereas most of the 'heavy rain' observations in the second half of the month (when rain was recorded by the other sensors) lie to the left. On this basis we offer a further flagging test for the data: they should be regarded as 'contaminated' if

$$1.25\, SPL_5 > SPL_{25} + 27.85 \qquad (1)$$

where SPL_5 and SPL_{25} are the sound pressure levels at 5 and 25 kHz respectively, and the constant is given to a fraction of a dB simply to force the line through the previous junction point on the display.

At present we only know of these anomalous spectra in our data, and are cautious of advocating a test which may be for conditions specific to this site. We had no real instances of heavy rain e.g. values of 100 mm hr^{-1} as found in the tropics, and so do not know whether the above test would discard genuine observations of significant rain. The main collection of wind-only points (grey circles in Figure 4) lie lower on the diagram than noted by Nystuen and Selsor [7]. This is not surprising given the different spectral slopes they observed; in fact, they found evidence of different slopes for different ocean basins. The acoustic classification scheme is a very important part of the suite of algorithms provided with the buoy. Clearly care must be taken to ensure improvements are not specific to one individual location.

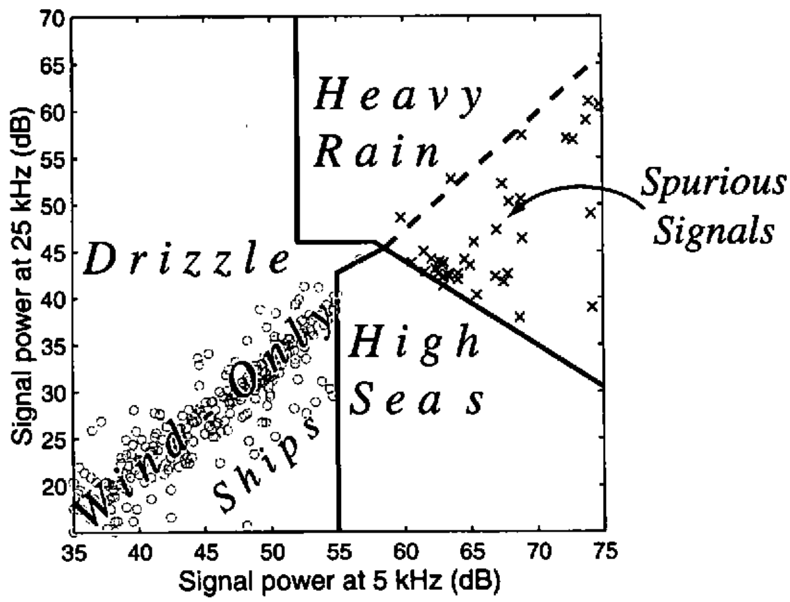

Figure 4. Classification of spectra according to the power intensities at 5 kHz and 25 kHz (after the approach shown in Figure A2 of Nystuen and Selsor [7]). Wind-only estimates are denoted by the grey circles and the false detections of rain by black crosses. The dashed line indicates the extra check on validity suggested by this paper. (Signal powers are in dB re 1µPa² / Hz.)

3.3 Quantitative analysis

The temporal and spatial resolutions of the various rain sensors are quite different. The rain radar provides instantaneous observations every 15 minutes over a square 5 km on each side. The *Casella* records the timing of tipping events to a precision of 1 minute, and the ARG calculates rainfall rate from data averaged over a 1.5-minute interval. For both, the average rainfall rate was calculated using 15-minute periods centred on the times of the radar scans. This is an appropriate timescale, not only because it is the interval between radar observations, but also because a rain system moving at 10 ms^{-1} (a typical speed for mid-latitude fronts) would take about 8 minutes to traverse 5 km.

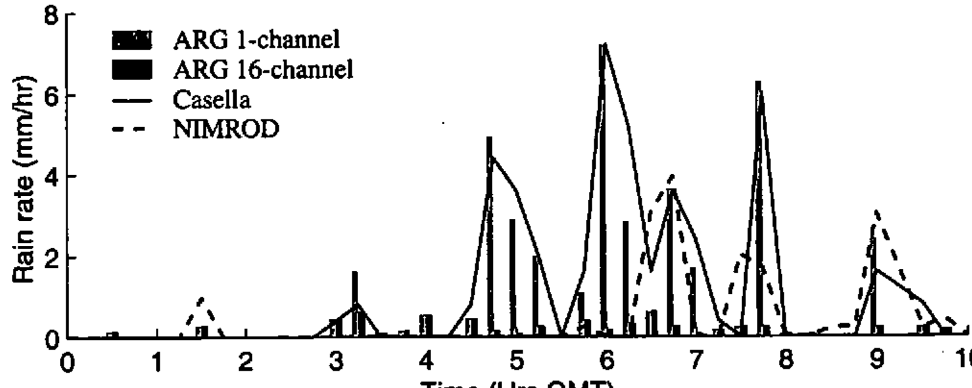

Figure 5. Inferred rainfall rates for the 2 algorithms for the ARG, the Casella tipping bucket and the NIMROD radar coverage. (Data are at 15-minute intervals from 24th May 2000.)

Figure 5 shows a comparison for the night and early morning of 24th May. All four series show similar timing of events. The time series for the *NIMROD* scans is the most different, since it corresponds to a much larger area than the others. It shows a complete absence of rain between 02:00 and 05:30 unlike the other datasets. The two series for the ARG show identical timing of events, since they are both governed by the same classification scheme as regards whether rain is present in the signal; however the magnitudes for the 16-channel algorithm are considerably less than for the other estimates. Its estimates are also much more poorly correlated with the *Casella* values than is the case for the single-channel algorithm. One would initially expect a multi-channel algorithm to yield a better performance than one using a single channel; we conclude that the poor results shown here must be due to the 16-channel algorithm being tuned for an environment somewhat different from that present in Loch Etive.

To quantify the agreement between the sensors, we perform scatter plots of the series, determine the slope of the best-fit line and calculate the correlation. To avoid the problem of errors in the acoustic classification, we here only consider data between 21:00 and 06:00 GMT. The first half of the month was dry, but ten of the nights in the second half contained some rain. The results are shown in Figure 6.

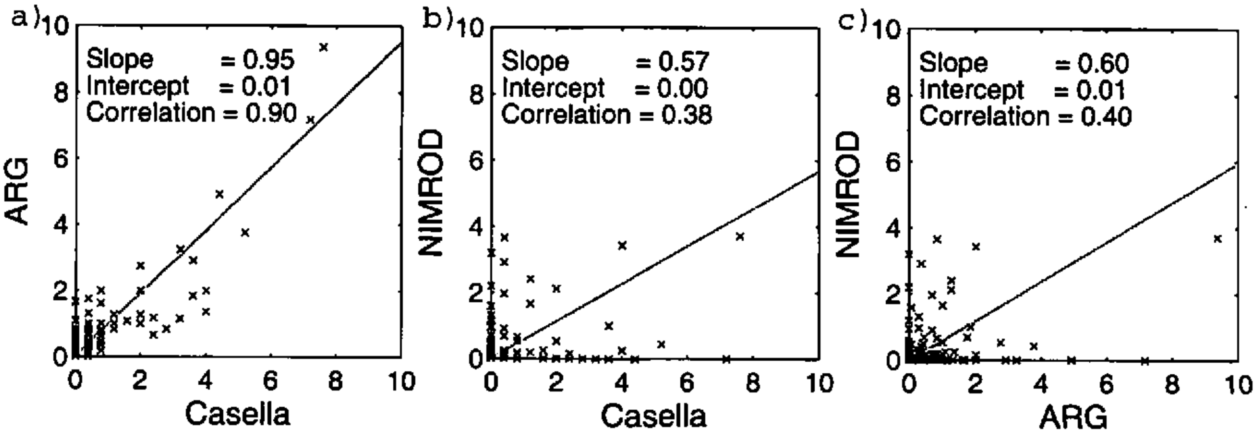

Figure 6. Scatter plots of independent 15-minute average rain rates (mm hr^{-1}) for all the nocturnal observations in May 2000.

All the fitted lines pass close to the origin, because 87% of observations showed no rain in any sensor. The agreement between the *Casella* and the simple 1-channel acoustics algorithm (see Figure 6a) is very encouraging: a slope of near unity, and a correlation of 0.90. This is much better than was achieved in the comparison between the two proven technologies viz. the tipping bucket and rain radar (Figure 6b). The low correlation achieved for that comparison is an artefact of the great dissimilarity in areas (and time spans) being compared. The correlation between the ARG and *NIMROD* (Figure 6c) is slightly better than for the *Casella.*

Of the sixty-two 15-minute periods when nocturnal rain was detected by the *Casella*, sixty of them also had rain detected by the ARG. This good detection rate is not surprising: if rain is falling at the ARG location, then there will be a spectral contribution which differs from that for wind-only conditions, and such events should be characterised as 'rain'. The problem is in accounting for the large number of times (fifty-three 15-minute periods) when the ARG suggests there is rain present, when there is no signal from the *Casella*. Although some may be due to failure of the ARG to detect 'contaminated' spectra (as in the daytime example in Figure 3), there are many instances on the edge of heavy downpours (as in the nocturnal example in Figure 3), where it is probable that the ARG is detecting drizzle too fine for the resolution of the *Casella* (which records accumulations of 0.1 mm at a time).

4. Summary and conclusions

The Acoustic Rain Gauge that we have been testing has made several month-long recordings of the underwater acoustic spectrum at a temporal resolution of 1.5 minutes. The data from a deployment in Loch Etive during May 2000, in combination with many other datasets, have enabled us to assess the usefulness of underwater acoustic data for detecting rain rates in an operational context. A key part of this is the acoustic classification as to whether the spectrum is due to wind, rain or other effects. In this location, which does not have much shipping activity, only 8.5% of data is flagged as 'contaminated'. However, there do appear to be other records with enhanced acoustic levels which give rise to apparently increased winds or large rain rates, when other *in situ* sensors indicate this not to be the case. An extra flagging criterion is proposed, which works well with these data, but whose wider applicability is untested. When there is rain, the rainfall rates inferred from the full 16-channel inversion are low and poorly correlated with those from the *Casella* tipping bucket gauge. This is likely to be because the observed spectra in this location differ from those used in developing the algorithms.

However, provided that the acoustic classification is correct (as is usually the case at night), the rainfall rates inferred from the 5 kHz channel are in good quantitative agreement with the *Casella*. One caveat is that the rain rates encountered during this trial never exceeded 10 mm hr^{-1}. We are currently collecting further months of simultaneous rain observations, with the hope of recording higher rainfall events, and with greater associated wind speeds. We then plan to move to a more exposed location off the southwest coast of Wales, where the longer fetch will lead to much higher wind-waves, and to interaction with longer period swell waves. Testing under such conditions is important, because as the amount of wave-breaking increases so will the underwater noise, in a way which may not be entirely governed by the local instantaneous wind speed. Wave-breaking can also alter the spectra due to rain, through the attenuation of the higher frequencies by sub-surface bubbles.

Acknowledgements

We are grateful to John Smithers, Keith Goy and Rob Bonner for their work on preparing and deploying all the equipment, and to Helen Tanner and Dawn Harrison for extraction of the *NIMROD* data. *Metocean* helped us with the prompt delivery of replacement parts and assistance on decoding the data relayed via *ARGOS*. These deployments were achieved with the assistance of John Holmyard (the mussel farmer) and the hospitality of the Polfern Hotel.

References

[1] Vagle S, Large WG and Farmer DM. An evaluation of the WOTAN technique of inferring oceanic winds from underwater ambient sound. *Journal of Atmospheric and Oceanic Technology* 1990; **7**: 576-595

[2] Scrimger JA, Evans DJ, McBean GA, Farmer DM and Kerman BR. Underwater noise due to rain, hail and snow. *Journal of the Acoustical Society of America* 1987; **81**: 79-86

[3] Pumphrey HC, Crum LA and Bjørnø L. Underwater sound produced by individual drop impacts and rainfall. *Journal of the Acoustical Society of America* 1989; **85**: 1518-1526

[4] Nystuen JA, McGlothin CC and Cook MS. The underwater sound generated by heavy rainfall. *Journal of the Acoustical Society of America* 1993; **93**: 3169-3177

[5] Medwin H, Kurgan A and Nystuen JA. Impact and bubble sound from raindrops at normal and oblique incidence. *Journal of the Acoustical Society of America* 1990; **88**: 413-418

[6] Quartly GD, Guymer TH, Birch KG, Smithers J, Keast A and Nystuen JA. Acoustics — a sound method of measuring rain?, in *Proceedings of Oceanology International '98*, Vol. 2, 239-246

[7] Nystuen JA and Selsor HD. Weather classification using passive acoustic drifters. *Journal of Atmospheric and Oceanic Technology* 1997; **14**: 656-666

[8] Nystuen JA. Acoustic rainfall analysis: Rainfall drop size distribution using the underwater sound field. *Journal of Atmospheric and Oceanic Technology* 1996; **13**: 74-84

[9] Quartly GD, Guymer TH, Birch KG, Smithers J, Goy K and Waddington I. Listening for rain: Theory and practice, in *Proceedings of Fifth European Conference on Underwater Acoustics*, Vol. I, 723-728