Advances in the Remote Sensing of Rain at Sea

J.W. Gregory, G.D. Quartly, T.H. Guymer & K.G. Birch

Southampton Oceanography Centre,
Empress Dock, Southampton SO14 3ZH
jwg/gdq/thg/kgb@soc.soton.ac.uk

Abstract

Precipitation is one of the major factors governing ocean climate. Despite the successful use of satellite-based observations in this area of remote sensing it is necessary to obtain *in situ* measurements to observe the high temporal and spatial variation in rainfall.

One of the most promising new technologies for measuring rainfall at sea is an Acoustic Rain Gauge (ARG). Rainfall measurement is achieved through sampling of the underwater sound spectrum between 500 Hz and 50 kHz, before sound intensity algorithms are used to classify the sound source and generate estimates of wind speed and rainfall rate. These algorithms are still experimental, and extended comparison with ancillary instrumentation is necessary to develop them further.

This paper discusses trials carried out in October 2000 to validate ARG rainfall measurements against land-based rain-sensing instrumentation. At present, the sensors show good agreement with ancillary instrumentation in terms of rainfall detection, but development of the algorithms is necessary to improve acoustic classification and hence remove false rainfall detections. Quantitative rainfall rate measurement shows promising agreement with ancillary sensors.

1. Introduction

Understanding the Earth s climate requires knowledge of the global distribution of precipitation, since precipitation is a key factor in the hydrological cycle, which in turn controls the circulation of the atmosphere and ocean. Unfortunately, measuring precipitation at sea is extremely difficult. Satellites provide part of the answer, particularly with respect to global monitoring, but are limited by the relatively low frequency with which they can observe a specific region, and by their large spatial footprint. Rainfall can vary on such small spatial and temporal scales, that satellite measurements are unable to resolve the diurnal variation, and regional rainfall differences become homogenised. Thus, there is a need for *in situ* measurements not only to resolve the spatial and temporal variation in rainfall but also to provide validation data for improving satellite products.

One of the most promising *in situ* technologies for measuring precipitation at sea is through inversion of the underwater ambient sound field. Wind and rain generate noise in the frequency range 500 Hz to 50 kHz, and furthermore the spectra generated are characteristic of the source, and louder than other sources (breaking waves, biology etc.) by several orders of magnitude. Thus, it is possible to discriminate as to the source of the acoustic signal, and use algorithms to make quantitative estimates of wind speed and rain rate. This sampling is carried out by the use of a sub-surface hydrophone, which avoids working at the hostile ocean-atmosphere interface.

Although it has been proven that the underwater ambient sound field can be used to measure wind speed and precipitation in this way, the accuracy of this method has yet to be established. The classification and rain rate algorithms are still considered to be experimental. Improvement requires testing of these algorithm products by comparison with conventional land-based rain-sensing instruments. This paper discusses the results of trials carried out in order to assess how well the current systems work, with particular emphasis on detecting and quantitatively measuring rainfall.

2. Experimental Set-up

2.1 Background

The results discussed in this paper form part of a long-running project to develop a reliable acoustic rain sensing system. Acoustic systems have been deployed nearly continuously since May 2000, providing us with an exhaustive dataset of sound field sampling and ancillary weather-sensing instrumentation. This paper relates only to analysis of the second stage of this deployment, from October 2000 onwards. (Work by Quartly et al. 2001 covers May data).

2.2 Location

The experiments were carried out in Loch Etive, a relatively exposed loch on the westcoast of Scotland. The site was chosen for its high annual rainfall, (among the wettest regions in the UK) and because the deep, saline loch should closely mimic open ocean locations. An open ocean mooring site was avoided at this stage as we need to validate the rain records of the acoustic system, and all ancillary rain sensors are land-based. Whilst we were not granted access to the neighbouring land, we were able to make use of a large moored raft belonging to the local mussel farmer. The raft is extremely stable, so that the land-based rain sensing equipment can be assumed to operate as normal. Previous comparisons using data from May have shown the rain-sheltering effects of the rafts super-structure to be negligible.

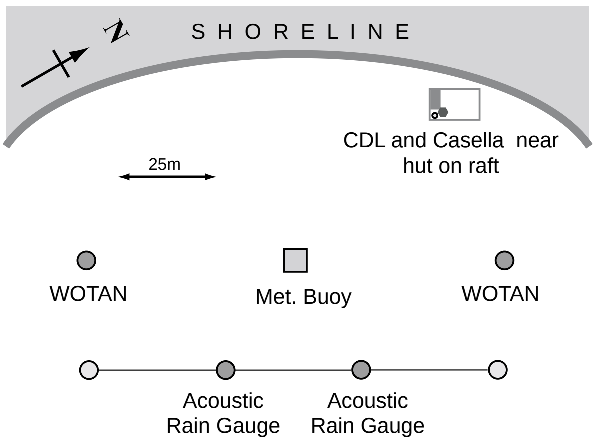

Figure 1. Mooring arrangement for deployment of weather-sensing instrumentation in Loch Etive.

2.3 Acoustic Systems

The equipment we are using is an Acoustic Rain Gauge (ARG) produced by *Metocean Ltd of Nova Scotia*, originally produced for use on drifting buoys but adapted by us for deployment on fixed moorings. The sensor incorporates a sub-surface hydrophone, deployed at a depth of 20 m, providing a sampling area of approximately 1200 m^2. (If absorption and refraction are neglected, the acoustic intensity is independent of depth). The hydrophone samples the underwater sound spectrum in 16 bands spanning 500 Hz to 50 kHz, and carries out on-board processing to classify the sound source and then calculate rain rate and wind speed. The system also incorporates a surface-mounted package to measure air pressure and air/sea-surface temperature. The hydrophone is suspended from a surface float by a coaxial cable, with a terminal weight suspended below the hydrophone to maintain depth. The ARG was deployed using a dual-tether mooring arrangement designed to avoid the hydrophone cable becoming wrapped around the mooring line. The rain gauge was modified to improve data collection and deployment length by fitting an internal logger, recording every 1.5-minutes, and by increasing battery storage capacity. During these trials we were able to achieve buoy deployments of approximately 40-days. In addition to recording data via the logger, the buoy transmits via ARGOS satellite relay, providing facility for real-time monitoring. Two such ARG systems were deployed simultaneously for the trials beginning in October, and are referred to by serial numbers 10392 and 11033.

2.4 Ancillary Instrumentation

Comparison data were provided by an array of ancillary instrumentation deployed in close proximity to the ARGs. Figure 1 shows the mooring arrangement. A *Casella* tipping bucket gauge was mounted onboard the nearby raft and fitted with the smallest available collection resovoir, each tip corresponding to 0.1 mm of rainfall accumulation. A second tipping bucket gauge was included as part of a Climate Data Logging system (identified as CDL in Figure 1) provided by the UK Met. Office, but this malfunctioned and no rainfall data from this sensor have been recovered for this stage of the deployments. However, analysis of earlier data has shown excellent agreement between this and the *Casella* tipping bucket, so we have confidence in the measurements made by *Casella*. Additional rainfall validation data were provided by the Met. Office s *NIMROD* rain radar network, which provides 15-minute coverage over a 5 km x 5 km area enclosing the Loch Etive mooring site.

Wind speed comparison data were provided by two WOTAN (weather observation through ambient noise) buoys, which use the same acoustic inversion algorithms for wind speed as the ARGs. However, these were near the end of their lifetime, and the data returned are heavily corrupted. Extensive quality checks are required before these data can be used, and this has not yet been undertaken. In addition, further wind-related readings are obtained from a meteorological buoy equipped with two anemometers and a wind vane on a mast approximately 2.4 m above the water. The wind measurements of the ARGs are not discussed here, due to the lack of WOTAN data and the fact that wind speed measurements from other sensors must be corrected to wind speed at 10 m above the surface, as returned from the ARG algorithms. This will form part of future data analysis.

3. Acoustic Discrimination

3.1 Weather Classification Algorithm

The *Metocean* ARG we are using is supplied with the algorithms developed by Nystuen (1996). Before estimating geophysical parameters, it is necessary to identify the source of the acoustic signal. This is possible because of the characteristic spectra associated with each sound source. Wind for example produces breaking waves and spray, hence giving rise to a broad band effect, with a typical spectral decay of —19 dB/decade (Vagle et al. 1990). Drizzle produces the same spectra under 10 kHz, but gives rise to a peak at around 14 kHz. This is due to the small diameter (0.8-1.1 mm) raindrops present in light rain, which produce sound in this frequency by the ringing of small bubbles generated by the raindrops. Medium sized raindrops (diameter 1.1-2.2 mm) are too big to generate small bubbles, and so make no significant contribution to the spectrum. Heavy rain however comprises small, medium and large bubbles, and hence generates sound by ringing and by the noise associated with large raindrop impact. Heavy rain therefore augments the sound intensity at all frequencies in the range 4-21 kHz, producing a plateau in the spectrum. Figure 2 shows typical acoustic spectra due to each source. There are several good papers regarding the sound radiated by different drop sizes (Nystuen 1996). As a result of this spectral characterisation, it is possible to separate the contributions due to wind and/or rain from other sources, and hence classify the acoustic signal as shown in Table 1.

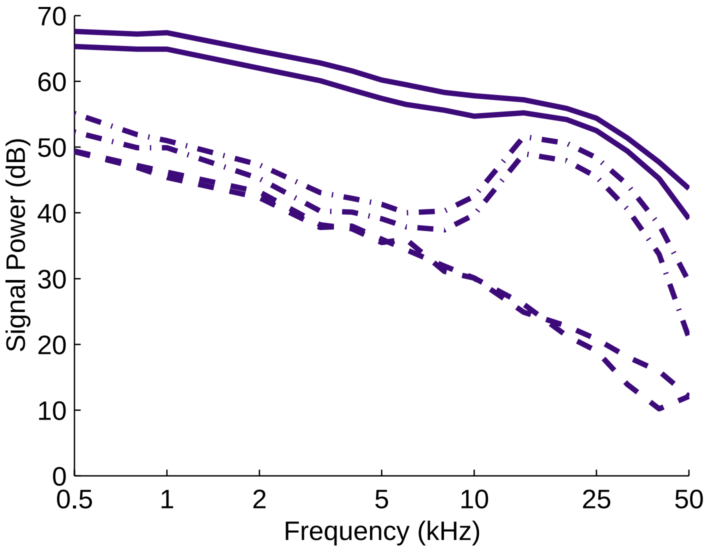

Figure 2. Acoustic spectra generated by different environmental sources: wind only (dashed), drizzle (dash-dot), and heavy rain (solid). [Signal power is in dB re 1μPa^{2}Hz^{-1}.]

Table1 Weather classification (Nystuen 1996).

Classification	Description
1	Wind only, low sea states
2	Wind only, high sea states (bubble clouds present)
3	Wind with drizzle
4	Heavy rain
5	Contamination

Having classified the sound source, the ARG then applies algorithms to infer wind speed and rain rate. Wind speed measurements use the algorithm developed by Vagle et al. (1990) to determine wind speed estimates from the sound intensity at 8 kHz. This algorithm is limited to wind speeds of greater than 2.1 ms^{-1}, since below this wind speed there are no breaking waves or wavelets, and hence no sound generated. This wind estimate is made for all wind conditions, that is, classifications 1 to 3. In the case of drizzle, a rain rate of 1 mmhr^{-1} is assigned, since no inversion is possible due to the effects of wind. For heavy rain, two possible rain rate algorithms can be applied. The first is a simple empirical relationship to the acoustic intensity at 5 kHz. The second uses the intensity at every frequency in the 16-channel range to infer a drop size distribution, describing the proportion of drop sizes present, and calculates a rainfall rate from this. This 16-channel algorithm is still being developed. The inversion coefficients used by the *Metocean* system were tuned for a brackish pond near Miami (Nystuen 1996) where rainfall estimates were found to show deviation from expected values, leading to the suggestion that the coefficients for this inversion are regionally specific. Analysis of the 16-channel algorithm products using May data have shown a tendency to underestimate rainfall rate relative to other sensors (Quartly et al. 2001). The coefficients of this 16-channel algorithm have been modified since the manufacture of the *Metocean* buoy, and we plan to re-calculate rain rates from our acoustic data using these. As a result, products from the 16-channel algorithm obtained with the old coefficients are only briefly discussed.

The classifications returned by the two ARGs were compared for a 30-day period during which both buoys were operational and no recovery work (generating unwanted noise) took place. Time series showed excellent agreement between the buoys, with over 85% of classifications matching. The remaining 15% are assumed to be due to either acoustic activity limited to the sampling area of one buoy, or to small timing differences between the buoys, resulting in adjacent records of different classification being compared. Due to the integer quantisation of this parameter, a time series does not best illustrate the agreement between the two buoys. Instead, we consider the percentage of records at each classification. The results are summarised in Table 2. Although we see excellent agreement between the two buoys, the large proportion of data classified as contaminated is surprising. Data recovered from May deployment showed only 8.5% flagged as contaminated. The results were checked for the entire dataset, but the results were the same, with 10392 recording 19% and 11033 recording 22% contaminated, hence it is not due to a spurious signal limited to this period of data. Understanding why records are classified as contaminated is an important part of analysing the acoustic record, hence we looked more closely at the classification algorithm used by the buoys.

Table 2 ARG classification comparison

Classification	ARG 10392 (%)	ARG 11033 (%)
1	56	53
2	4	4
3	11	11
4	10	11
5	19	21

3.2 Rainfall Detection

In order to estimate how well the classification is working, we first compared the results from ARG 11033 with the *Casella* tipping bucket gauge. (We used 11033 because it provides an uninterrupted response for the period of best *Casella* data). Data were examined over a period of 30-days during November/December 2000. The results are encouraging. Figure 3 shows the results for day-328, which is representative of the entire sample.

There is good agreement in terms of both timing and qualitative measurement. The *Casella* gauge tips more frequently during heavy rain and only records onto its data logger when a tip occurs. Hence, the density of *Casella* tips (represented here by red crosses) is a measure of the rainfall rate. We can see that the two sensors compare extremely well in terms of detecting heavy rain. However, the *Casella* seems less responsive to the occurrence of what the ARG classifies as drizzle . This has been noted in previous trials, and is due to the fact that the *Casella* has a limited resolution of 0.1 mm. Below this, rainfall is insufficient to cause the bucket to tip, preventing the *Casella* from detecting light rain. Acoustic rainfall measurements however have been shown to be extremely sensitive to these low rainfall rates, another advantage of this type of sensor over conventional rain sensing equipment.

The data also show a number of strong acoustic signals, classified as heavy rain but which are undetected by the *Casella* rain gauge. This is unexpected, given the excellent agreement in terms of heavy rainfall agreement elsewhere in the time series. One such feature is shown in Figure 3 occurring at around 15:00. These spurious readings were noted both in our earlier analysis of data from November 1999 (Quartly et al. 2000) and from May 2000 (Quartly et al. 2001), and warrant closer inspection.

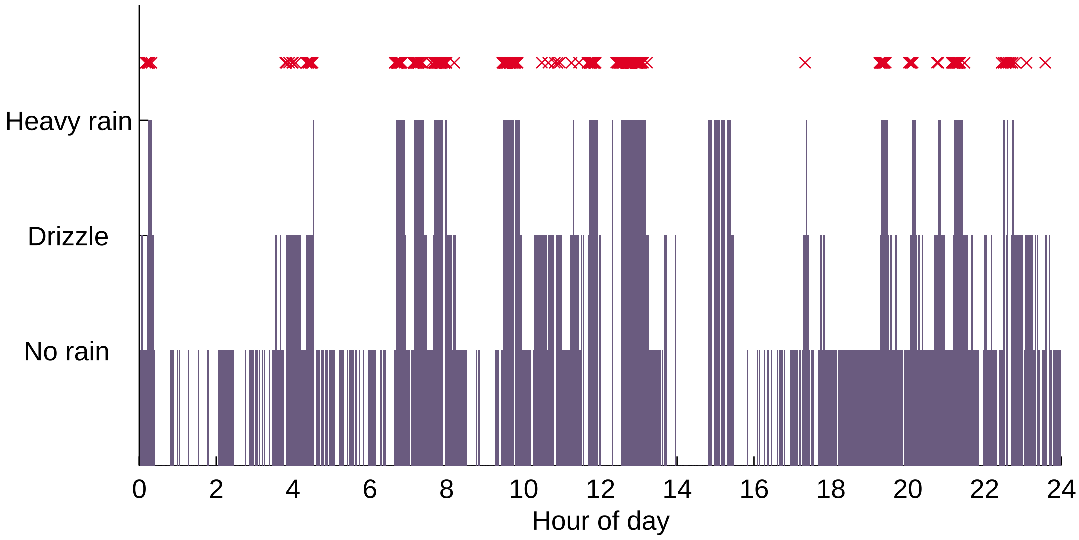

Figure 3. A comparison between acoustic classification by ARG (1.5-minute resolution) with rain events detected by *Casella* (accummulations of 0.1 mm) for day-328. Note the spurious heavy rain detection by ARG at approximately 15:00. Gaps in the data refer to records classified by ARG as contaminated .

3.3 Problems in Acoustic Classification

Examination of a period of 4 days looking closely at each observed anomaly shows they occur clustered together around midday, with no occurrence of these spurious readings outside this time. Figure 4 shows the spectra associated with each spurious signal cluster obtained for each day. (Confirmed heavy rain spectra are included for comparison — shown in red). There appear to be two different types of spectra. The first type of spectra, referred to as type-1 , appear very similar to the heavy rain spectra, and occur primarily on day-316 and day-317. These show 8/7 spurious readings respectively, each closely matched to the typical spectra for heavy rain. They reveal a steeper spectral decay however, with a faster fall-off above 10 kHz, again as was previously observed (Quartly et al. 2001). It was this type of spectra that prompted proposal of an improvement to the classification algorithm. It was suggested that records be classified as contaminated if they fulfil Equation 1. Figure 5a shows a cluster plot of SPL_5 and SPL_{25} (sound pressure levels at 5 and 25 kHz respectively) of day-316 data, with spurious records highlighted in red.

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

Analysis of this plot shows that the spurious readings are all grouped in the region through which our proposed adjustment passes — that is, the adjustment will index around 40% of the spurious records. Assuming the other acoustic signals not validated by the *Casella* lie in the same region, it would seem safe to propose an adjustment to the line, raising it slightly to index all of the so-far observed spurious data. This was avoided previously since the lack of heavy rainfall meant the impact of the adjustment on removing valid rainfall records was unknown.

In summary, these acoustic signals consistently appear in our data around midday, and in the absence of other explanation are attributed to either activity associated with the nearby mussel farm, or the local seal colony. As such, they represent a contaminant sound source that is being indexed by the exisiting classification (falsely) as heavy rain . The proposed adjustment to the classification algorithm succeeds in only indexing around half of the spurious records, but could be adjusted to detect nearly all of these data by raising the line 2-3 dB (Equation 2). The effect of this will need to be investigated by comparison with other records, to ensure this does not index valid heavy rain records.

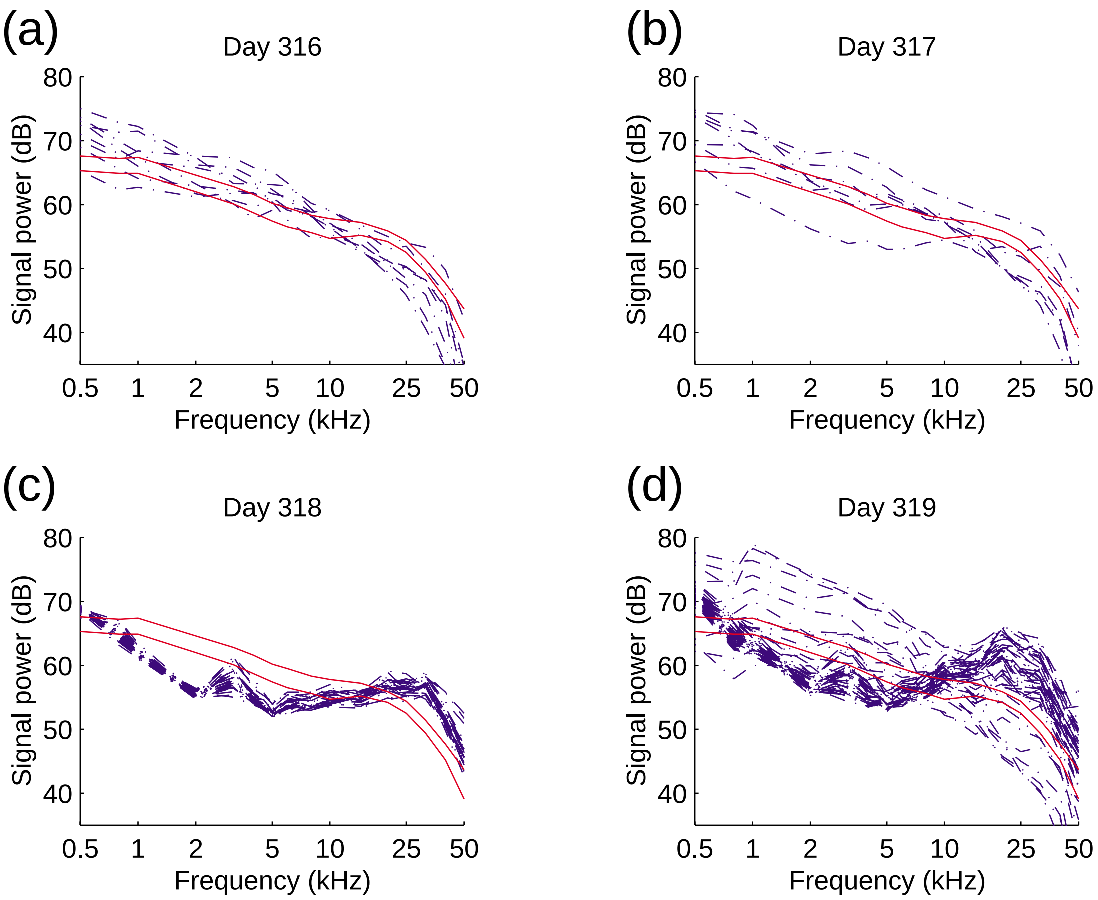

Figure 4. Spectral plots for 4-days showing events classified by ARG as heavy rain , not detected by *Casella*. Two spectral types are evident. [Signal power is in dB re $1\mu Pa^2Hz^{-1}$.]

$$1.25\ SPL_5 > SPL_{25} + 25.35 \tag{2}$$

Having accounted for the type-1 spectra observed for the first two days, we considered the spectra evident for day-318 and day-319, referred to here as type-2 spectra. These are shown in Figure 4c,d. These spectra have a very characteristic shape, to the eye easily discernible from that due to heavy rain, but are being classified as such by the ARG s internal algorithm. The spectra show a linear decline in intensity over 500 Hz to 2 kHz, before peaking sharply at 3.15 kHz. This then falls sharply at 5 KHz, before a steady climb to a second peak at around 25 kHz, followed by a sharp fall-off. We devised a simple algorithm to detect these signals and hence determine how often they occur throughout our chosen 30-day time period. The nature of the spectra suggest they may be caused by activity associated with the nearby mussel farm. Over the 30-day period, we found evidence that these type-2 spectra occur on 10 separate days, over a period of 1-2 hours, near midday. Further investigation shows that the spectra occur for Monday 13th November to Friday 17th November, again for Tuesday 21st November to Friday 24th November, and then again from Monday 27th November. Hence, we conclude that the source of these signals is anthropogenic, either in terms of noise generated by the mussel farmer s boat, or by machinery located on the nearby raft.

Of the 192 records indexed as exhibiting these type-2 spectra, 79 (41%) were correctly indexed as contaminated by the exisiting classification algorithm. However, 97 (51%) were falsely classified as heavy rain , and the remaining 16 (8%) as drizzle . These may account for some of those instances where the ARG detects drizzle and the *Casella* fails to detect rain, other than those due to the *Casella* s inability to resolve fine rainfall (Figure 3). We would expect any noise from the mussel farmer s boats to be classified as shipping, and indeed, a spectral plot of data from day-318 (see Figure 5b, spurious signals highlighted in red) show these data are well removed from the zone identified as contamination due to shipping. As a result, we conclude these spurious signals are probably due to machinery located onboard the raft that is operated around midday for most of the working week. Since these signals are peculiar to Loch Etive, we do not suggest any measures for their removal from acoustic spectra. In any case, these signals are located in a region which has for day-316 and day-317 (see Figure 5a) returned valid heavy rain readings, so extraction of these data would require consideration of alternative frequencies.

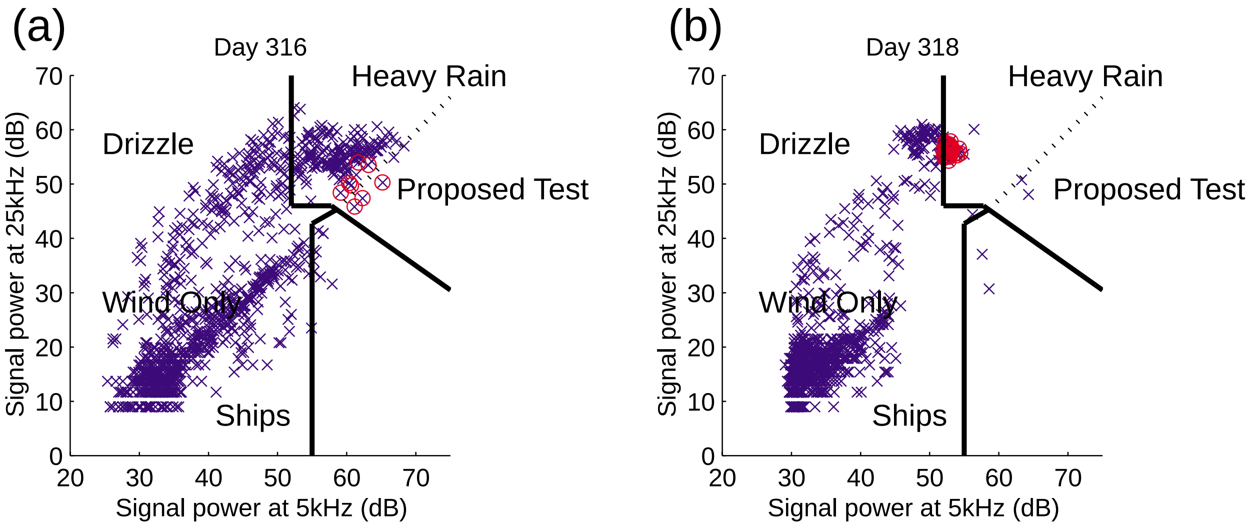

Figure 5. Classification of acoustic signals from day-316 (a) and day-318 (b) data, by comparison of acoustic intensities at 5 kHz and 25 kHz, after approach used by Nystuen and Selsor (1997). Type-1 spectra are highlighted in red in figure (a), and type-2 spectra are highlighted in figure (b). Also shown is the test previously proposed (Quartly et al. 2001) for indexing type-1 spectra (prior to the modification suggested in this paper). [Signal power is in dB re $1\mu Pa^2 Hz^{-1}$.]

4. Quantitative Rainfall Measurement

To quantitatively compare rainfall measurements from different rain sensors it was necessary to grid to the same time intervals. Here we compare data from the ARG, for both 1-channel and 16-channel algorithms, as well as the *Casella* tipping bucket gauge and the *NIMROD* rain system, part of the UK Met. Office s gridded rain radar network. These sensors are quite different in terms of spatial and temporal resolution. The ARG calculates rainfall rate from data averaged over 1.5-minute periods, whereas the *Casella* records an accumulation to a precision of 1-minute. Further, the *Casella* is subject to a minimum accumulation of 0.1 mm, giving rise to a quantisation in any inferred rainfall rate. The rain radar system provides instantaneous observations every 15-minutes, with the rainfall rate returned corresponding to a 5 km square, encompassing Loch Etive. The *Casella* is effectively a point measurement, whereas the ARG sampling area is of the order of 1200 m^2.

Figure 6 shows a simple comparison of these data for a typical day. It should be noted that the results of the 16-channel algorithm are included for illustration only. The rainfall rate recorded by the logger is given to a precision of 1 $mmhr^{-1}$, so for comparison we have recalculated the 1-channel rain rate from the transmitted intensities. This has not been done for the 16-channel rain rate estimate, since we defer rigorous analysis of this algorithm until we can apply the latest coefficients to the data already collected. Despite this, we can draw some preliminary conclusions from the data presented here.

Analysis of the data collected during May showed the 16-channel inversion to be underestimating rainfall rate relative to the other sensors. This does not appear to be the case for these data. There is good agreement between the ARG rainfall rate estimates, as we would expect given that the rain rate is determined by default for all but classification 4 data (occuring 10% of the time). However, the 16-channel algorithm is generating much higher rain rate estimates than previously. We had hoped the 16-channel algorithm would return better estimates in heavier rain conditions, but these higher estimates are being returned for periods where rainfall rates from other instruments do not exceed those recorded in May. We are unable to account for this improvement in the 16-channel estimate without rigorous analysis of the data, but it may be that the rainfall type varies between May and October. That is to say that the proportions of drop sizes in the rain varies, hence changing the components of the spectral contribution, thus giving rise to a better rain rate estimate.

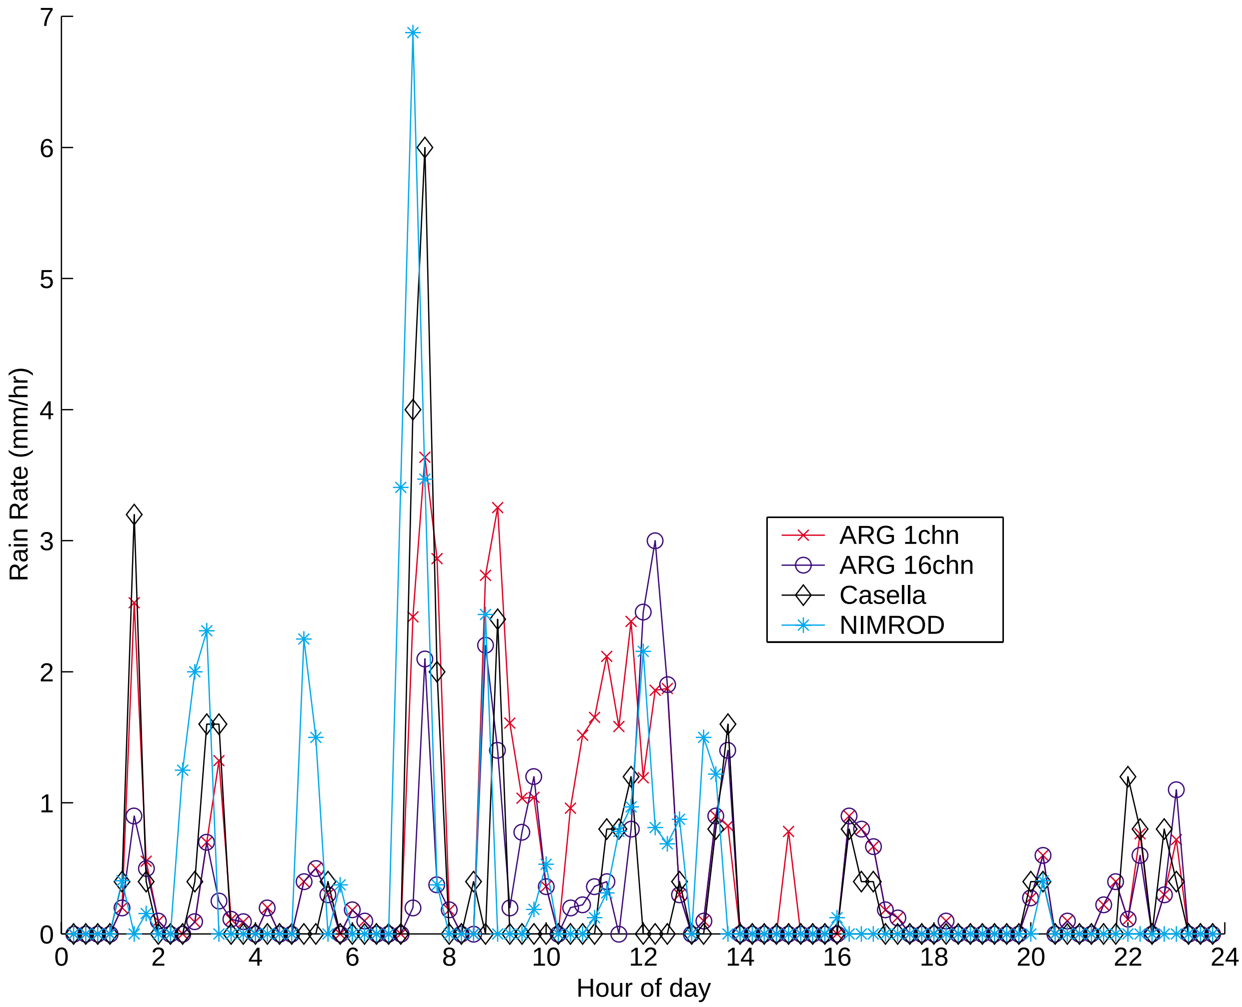

Figure 6. Inferred rainfall rates for ARG, *Casella* and *NIMROD* sensors. (Data are at 15-minute intervals from 17th November 2000).

The time-series shows that there is generally good agreement between the sensors as to when it is raining, and as to the approximate rainfall rate. However, the estimates of rainfall rate show some variation, which is a product of the different spatial/temporal resolutions. The *NIMROD* data is less useful for this form of comparison since it corresponds to a much larger area than the ARG or *Casella*. We perform scatter plots to measure the agreement in the data, and we find results are generally similar to these gained through analysis of the May data. For each, we determine the slope and intercept of a best-fit line, and calculate the correlation in the data (Figure 7).

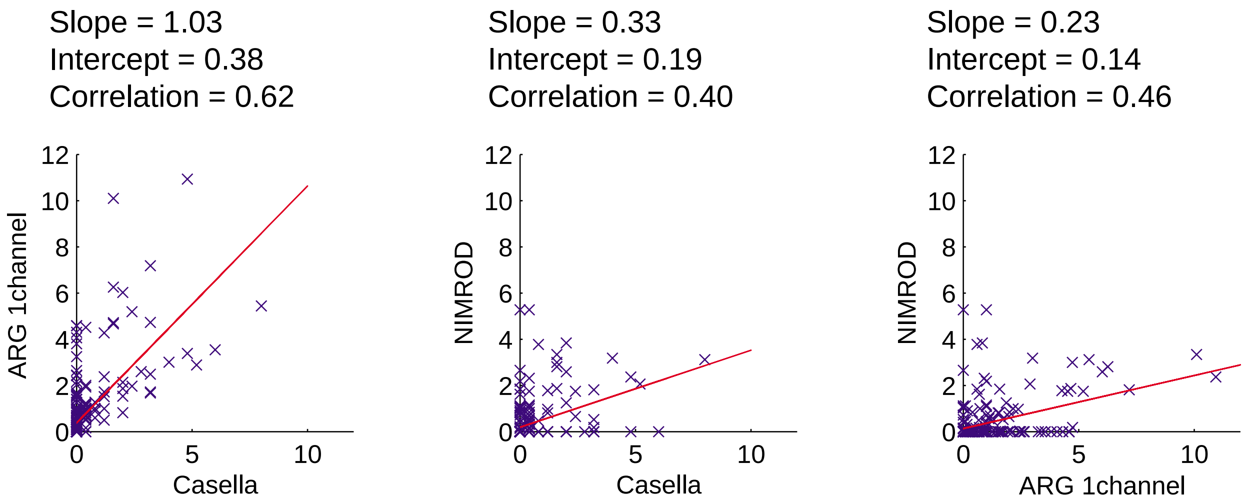

Figure 7. Scatter plots of independent 15-minure average rain rates (mmhr-1) for 30-day period. Note that only every tenth point has been plotted, but that fitted line and slope/intercept/correlation relate to entire 30-day dataset.

All the fitted lines pass close to the origin since around half of the readings for each sensor correspond to no rainfall having been detected. We find the best agreement is between the 1-channel algorithm and the *Casella* rain gauge. A correlation of 0.62 is significantly higher than attained in the other comparisons. This is somewhat disappointing though compared with the correlation of 0.90 achieved for the May data. This is partly due to the data used for these comparisons. In analysis of the May data, only nighttime records were considered, in order to remove all spurious records occurring during the daytime. Removal of the spurious rainfall detections in the ARG record here is likely to improve the correlation, and this will be investigated. Furthermore, for the May data, 87% of observations showed no rain in any sensor. For October, we find only 56% of observations showed no rain. This means that a greater proportion of the data are rainfall estimates, which compare less well given the quantisation in the *Casella* data to 0.4 $mmhr^{-1}$, and also the occurrence of heavy rain signals in the ARG not found in the *Casella* data. However, the main difference seems to be in detecting light rain, since the mean rainfall rate for the 22% of observations where the ARG detects rain not measured by the *Casella* is 1.05 $mmhr^{-1}$. This is compared with the 1% of observations for which the *Casella* detects rain not recorded by the ARG. The slope is a much better indicator in this example of the comparison, with a value of 1.03, close to unity and to the previously obtained value of 0.95. The fit may be improved still further by removal of the spurious heavy rain signals. These certainly appear to be contributing to the much higher mean rainfall rate returned by the ARG, which measures 0.71 $mmhr^{-1}$, in comparison with the *Casella* (0.32 $mmhr^{-1}$) and *NIMROD* (0.30 $mmhr^{-1}$). The agreement in these latter two does not suggest this value is true, since both the *Casella* and the *NIMROD* are known to be poor at resolving light rain, which makes up a large component of the annual UK rainfall. Without removal of identified false detections and compensating for quantisation and threshold limits, it is difficult to carry out a quantitative comparison of rainfall rate measurements.

5. Summary and Conclusions

The Acoustic Rain Gauge we have been testing has been deployed nearly continuously in Loch

Etive from May 2000 to March 2001. We have accumulated an exhaustive dataset for this period comprising 1.5-minute observations of the underwater ambient sound field, as well as data from ancillary rain and wind sensing instrumentation. These have been compared for the period October 2000 to January 2001 to determine how well the current ARG detects rain, with particular regard to classification of the sound source.

Comparison of rainfall detection between the Acoustic Rain Gauge and a tipping bucket located on a nearby raft shows that there is good agreement between the sensors over whether or not it is raining, with 77% of observations agreeing. In particular, the acoustic system seems good at resolving fine rain below the collection threshold of ancillary instrumentation. However, of the 23% disagreeing, 22% are due to rain detection by the ARG, not validated by the *Casella*. A significant proportion of these data are likely to be drizzle undetectable by the *Casella*, but will also be due to false detections by the ARG. Examination has shown these to be caused by two different spectra, occurring clustered around midday. Type-1 spectra are thought to resemble those observed in analysis of May data, and an improvement to the flagging of these spectra on the previous test has been proposed. Type-2 spectra are believed to be produced by equipment associated with the nearby mussel farm, and so no test for removal of these spectra from the data has been suggested. Quantitative comparison of the rain rate estimates suggests promising agreement between theAcoustic Rain Gauge and other sensors, with the 16-channel rain rate algorithm yielding good comparison values, despite its integer resolution. Further analysis of these data is proposed, using the latest algorithm coefficients and improved flagging and removal of known-spurious readings.

The data from this deployment are still being recovered and processed, but as the data collection stage of the project comes to an end we can begin in-depth processing and interpretation. This will include looking more closely at the rain rate estimates, with particular reference to the latest coefficients for the 16-channel rain rate algorithm. In addition, we hope to carry out similar studies on the acoustic measurement of wind speed, and to determine what effect wind has on rain detection, particularly in terms of masking the spectral contribution of light rain. We then plan to move to a more open location off the southwest coast of Wales, where the buoys will be exposed to the south westerly winds and associated Atlantic weather fronts, following hardware modifications to the current rain buoys to improve battery storage and data capacity.

6. Acknowledgements

We are grateful to John Smithers, Keith Goy and Rob Bonner for their work in development of the ARGs, and in the deployment and data recovery of sensors in Loch Etive. We are also grateful to *Metocean* for supplying the acoustic systems, and to the UK Met. Office for supply of rain radar data, and for continued assistance in data analysis. These deployments were achieved with the assistance of John Holmyard (the mussel farmer), who watched over our equipment, and the hospitality of the Polfern Hotel.

7. References

NYSTUEN, J.A., 1996, Acoustic rainfall analysis: Rainfall drop size distribution using the underwater sound field. *Journal of Atmospheric and Oceanic Technology*, **13**, 74-84.

NYSTUEN, J.A., and SELSOR, H.D., 1997, Weather classification using passive acoustic drifters. *Journal of Atmospheric and Oceanic Technology*, **14**, 656-666.

NYSTUEN, J.A., and MCPHADEN, M.J., 2001, The beginnings of operational marine weather observations using underwater ambient sound. In *Proceedings of the Institute of Acoustics*, **23 Part 2**, pp. 135-141.

QUARTLY, G.D., GUYMER, T.H., BIRCH, K.G., SMITHERS, J., GOY, K., and

WADDINGTON, I., 2000, Listening for rain: theory and practice. In *Proceedings of the Fifth European Conference on Underwater Acoustics*, **1**, pp. 723-728.

QUARTLY, G.D., GREGORY, J.W., GUYMER, T.H., BIRCH, K.G., JONES, D.W., and KEOGH, S.J., 2001, How reliable are acoustic rain sensors? In *Proceedings of the Institute of Acoustics*, **23 Part 2**, pp. 142-148.

VAGLE, S., LARGE, W.G., and FARMER, D.M., 1990, An evaluation of the WOTAN technique of inferring oceanic winds from underwater ambient sound. *Journal of Atmospheric and Oceanic Technology*, **7**, 576-595.