

Dependable Outlier Detection in Harsh Environments Monitoring Systems

Gonçalo Jesus¹[0000-0002-8431-3877], António Casimiro²[0000-0002-5522-5739] and Anabela Oliveira¹[0000-0002-7537-0984]

¹ Hydraulics and Environment Department, LNEC. Lisbon, Portugal

² LASIGE, Faculdade de Ciências, Universidade de Lisboa. Lisbon, Portugal

gjesus@lnec.pt

casim@ciencias.ulisboa.pt

aoliveira@lnec.pt

Abstract. Environmental monitoring systems are composed by sensor networks deployed in uncertain and harsh conditions, vulnerable to external disturbances, posing challenges to the comprehensive system characterization and modelling. When unexpected sensor measurements are produced, there is a need to detect and identify, in a timely manner, if they stem from a failure behavior or if they indeed represent some environment-related process. Existing solutions for fault detection in environmental sensor networks do not portray the required sensitivity for the differentiation of these processes or they are unable to meet the time constraints of the affected cyber-physical systems.

We have been developing a framework for dependable detection of failures in harsh environments monitoring systems, aiming to improve the overall sensor data quality. Herein we present the application of an early framework implementation to an aquatic sensor network dataset, using neural networks to model sensors' behaviors, correlated data between neighbor sensors, and a statistical technique to detect the presence of outliers in the datasets.

Keywords: Dependability, Data quality, Outlier detection, Machine learning, Neural networks, Water monitoring.

1 Introduction

Environmental monitoring systems based on wireless sensor networks consist of sensor nodes distributed over the observable areas, which are connected to sink nodes through a wireless network. The sink nodes gather the readings of the sensor nodes, perform sensor fusion and other data processing tasks, and transmit the results to decision-support systems, alert and emergency systems, or actuators in the field. Therefore, environmental monitoring systems can be considered complex cyber-physical systems.

Given the importance of the flux of the gathered data into end systems, it is crucial to recognize, categorize and mitigate sensor faults in environment monitoring networks. As sensor data comprises both correct measurements and noisy or faulty values (either outliers or continuous sets of unreliable data), this error-prone data can lead to the issuing of false warnings or backing wrong decisions.

While fault detection in sensor networks has been subject of several studies [1-3], many of the methodologies or techniques used are not appropriate in harsh environments applications. Harsh environments are characterized by extreme or abnormal conditions that can both damage sensor components and create unusual perturbations in environmental conditions, which pose challenges for ensuring the quality and correctness of monitoring. Hence, harsh environments require monitoring solutions that are resilient to these disturbances and uncertainties, in contrast to solutions for controlled and mostly stable settings. And given the criticality of applications, it becomes imperative to meet the requirements on data quality collected from the monitored environment and used in the cyber-physical system.

An important step towards the increase of data quality in sensor networks is the employment of fault detection methodologies [4], starting by a characterization of the possible faulty behaviors affecting the monitored data. However, in harsh environments this characterization is harder than in controlled environments, because unexpected changes in the environmental conditions, such as storms, earthquakes or fires, create abnormal patterns in sensor data which, however, should not be treated as faulty behaviors. For instance, a temperature sensor for weather control affected by a fire incident will output values significantly above the normal one, which, however, do not portray a sensor failure situation.

Furthermore, to improve the monitoring system it is also necessary the definition of strategies to mitigate or, ideally, correct the faulty readings, making it possible to utilize the readings with some level of confidence in the end-systems.

Considering the presented challenges and requirements for monitoring in harsh environments, we have been developing a generic monitoring framework that we briefly describe in this paper. The framework provides: 1) support for the detection of faulty readings in the sensor measurements; 2) quality information associated to every sensor measurement; and 3) a reassessment of a measurement whenever it considered faulty or with insufficient quality. Based on this framework, the main contributions of this paper are the following:

- The application of the framework to an existing sensor network in a river-estuary environment that is subject to severe weather conditions and various hazards affecting sensors;
- An overview of the designed strategy for detecting outliers with high accuracy, along with an approach for the reduction of false positives;
- Experimental results of the application of machine learning techniques in both the framework steps, customized for the scenario of the river-estuary environment using real datasets from sensors deployed in the water, measuring temperature and salinity parameters.

The paper is structured as follows. In Section 2 we present an overview of the monitoring framework, with the description of its components and procedures. Section 3 is focused on an application to a river-estuary case study and on the machine learning solution conceived for the detection of outliers. Extensive results comparing the solutions under different customizations are explored in Section 4, while Section 5 closes the paper with some conclusions and future work considerations.

2 Dependable Monitoring Framework

The dependability of sensor data is directly related with the quality of measurements. Any given framework or fault-tolerant strategy must be able to characterize the possible sources of faults affecting the sensor network, related failure behaviors and corresponding fault models. We presented in [4-5] a brief overview of the commonly observed failure behaviors, and the strategies based on machine learning technologies, which are proven to be adequate to model the observed system. Furthermore, we proposed and described a generic dependable monitoring framework for harsh environments in [6].

The framework, which we briefly overview, is based on data fusion techniques using machine learning that model each sensor behavior according to all sensors past measurements. Considering that there must be an history of measurements previous to the application of the framework procedures, there is a preliminary step consisting on the creation of such models. For fault-detection purposes, each sensor must be represented by at least two models exploring temporal, spatial and value correlations between either target sensor past measurements or a combination of the sensors existent in the sensor network. Additionally, the framework is composed by the following four main components that are executed in the sink node whenever a new measurement from any sensor is received:

- *Prediction (P)* – in a harsh environment the truthfulness of a measurement can only be assessed by knowing the involved impacting conditions, what is seldomly feasible. Alternatively, to substantiate the truthfulness of a measurement, an attainable method is to employ several prediction methods to estimate the measurement;
- *Failure Detection (FD)* – the objective is to detect and identify failure behaviors affecting measurements of the target sensors. This involves the execution of procedures not only to characterize abnormal measurements but also to distinguish sensor faults from environment-related events. The predictions obtained in the previous component are of great importance in this step;
- *Quality Evaluation (QE)* – the outcome of the previous components enables the quantification of the confidence level on the new received measurement. The component output is a quality coefficient associated to the measurement. When facing a failure behavior, this value is set to 0;
- *Measurement Reassessment (MR)* – if a measurement is faulty it is important for future predictions and for further systems to have access to an estimate of the expected measurement, even if it is not fully reliable (confidence level below standard). This component enables the mitigation of the effects of the failure behavior by producing an estimation of measurement as if it was not affected by any fault.

The flow of information processing through these components is depicted in Fig. 1. Before a measurement can be fully processed, a warm-up stage is necessary to build the histories of past measurements for each sensor. Only then it is possible to employ the prediction methods to obtain an estimate of the expected measurement. Additionally, besides the past measurements it is possible to use other information correlated with the target sensor, such as correlated sensor data or validated numerical models that

simulate the dynamics of the monitored system, which work as virtual correlated sensors providing estimated readings.

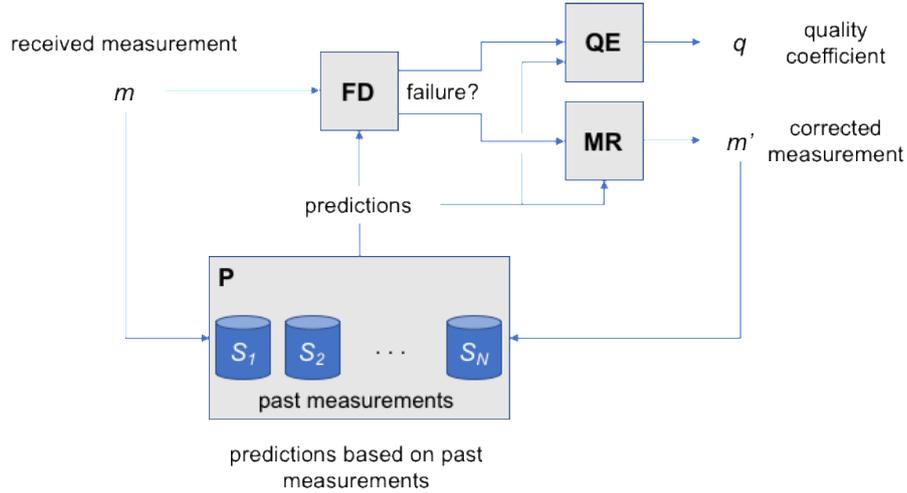


Fig. 1. Flow diagram of the dependable monitoring framework.

The prediction methods consist on the data fusion techniques that correlate sensors past information to build a faithful model of the target sensor. There are no constraints on the used techniques, but some may be more suitable to the conditions and specifics of the sensor network and the environment. Regarding the quantity of predictions, the only requirement is to have at least two predictions, being one of them only based on past measurements of the target sensor. This ensures that either faulty measurements or specific events of other sensors do not propagate to at least one prediction.

In the Failure Detection component, processing is performed to compare the predictions to the newly received measurement. If there are no faults affecting the sensors, or other environmental events, all predictions will be consistent with the measurement. On the other hand, if some prediction is significantly different from the received measurement, this might indicate the presence of faults. It is the comparison strategy that enables the detection of faults and their distinction from real environmental events. As described in the Prediction component, the comparison selection strategy is also part of the application process since there is no specific method to assess the significance of the differences. The method depends on the type of failures, the process dynamics and the type of correlations existent in the predictions. In Section 3 we exploit the distribution of square differences to detect outliers, but other statistical or machine learning methods can be also considered.

In the Quality Evaluation component, the confidence level in the received measurement will be quantified as a quality coefficient, between 0 and 1, being 0 the lowest and

1 the highest possible quality. This confidence value can be extracted through an evaluation method that, for example, calculates the significance of the differences among measurement and the predictions.

The Measurement Reassessment component considers the same information, measurement and predictions, and produces an estimated measurement that is expectedly free from the effect of faults. For instance, if an offset failure is affecting the received measurement, the offset value should be removed from the received measurement, and if the received measurement is considered an outlier, the corrected value should be estimated using only the predictions and discarding the faulty measurement.

3 Application

Our case study addresses the problem of abnormal measurements in aquatic monitoring environments. We present in the first subsection the case study, that fully represents the harsh environment definition, with high variability in water parameters, such as temperature and salinity, accompanied with the unpredictability of weather scenarios and possible human interference. For this effect, in the second subsection, we applied the framework to the available data sets with the goal of detecting outlier situations in the temperature data. The application process is briefly described, where we employed artificial neural networks (ANNs) to model the sensor behaviors and differentiate sensor faulty measurements from environment event-related abnormal behaviors affecting sensor measurements.

3.1 Case Study

Although the purpose of the presented framework is to be applied in real-time monitoring, it is also possible to employ the procedures for the analysis of pre-collected data sets. In fact, the results presented in Section 4 were obtained using existing data sets from a monitoring network in the Columbia River estuary, situated in the northwest coast of the USA. The sensor network is owned and maintained by CMOP's Science and Technology University Research Network monitoring network (denoted SATURN [7]), and it is comprised of various sensors deployed along the estuary, measuring many water parameters such as water levels, salinity and temperature or biogeochemistry on a 24/7 basis.

From the 27 available stations, each one with many sensors, we selected a few for our experimental evaluations. Jetty A, Lower Sand Island light, Desdemona Sands light and Tansy Point were selected, due to the sensor (temperature, elevation and salinity) and monitoring periods similarity.

In the preliminary stage of the framework application, and given the environment specific characteristics, it was observed that the weather seasons have a clear (and expectable) impact on parameters such as temperature. Another important observation is related to the specific characteristics of the considered aquatic environment, an estuary, where the water parameters are significantly impacted by the tides throughout the day.

This influence is represented by tidal harmonic constituents [8], where the largest referring to the principal lunar semi-diurnal is a cycle of 12 hours and 25.2 minutes.

3.2 Prediction Models

The river-estuary system is a good example of the non-linearity present in natural and harsh environments. To successfully model these characteristics and be able to capture the existing correlations in the system, we had to consider machine learning techniques. Among various available techniques, we selected artificial neural networks (ANNs) for their capability to learn patterns using training data and the ability to control the inputs and consequent outputs.

In the preliminary phase, we used the specific knowledge of environment being monitored to better prepare a feature selection step and to define the structure of the input vectors of the ANNs. The objective is to predict the next measurement of a particular sensor based on historic knowledge (past measurements) of that sensor and the neighbors, or a combination of this information. Thus, the feature selection step involves a careful selection of which sensors and which data is used in each model.

Considering the tidal influence (specific to this environment) we designed the input vectors to cover the sensors behavior during the principal lunar semi-diurnal cycle. In the training process, given the weather seasonal patterns, we selected a full year of data from all the sensors involved to form the inputs and target vectors for the models.

3.3 Outlier Detection Strategy

In our outlier detection strategy, we verify the difference between the outputs of prediction models and the received measurements by employing a statistical method of the distribution of square differences. This comparison strategy allows the assessment of the difference with a level of significance based on historical data. This historical knowledge was supported by the fitting of the distributions of the differences between each model (ANN) and the corresponding targets for a period of two months, not included in the training data set of the preliminary phase.

With the fitted distributions, in the Failure Detection component the process is to calculate the probability that the difference between a newly received measurement and the expected value is significant, considering a confidence level. In Section 4 we study the impact of selecting different values for the confidence level.

For outlier detection, we verified if all the predictions are significantly different from the received measurement, in which case there is very high probability that it is faulty measurement. As mentioned in the previous section, it is important to use at least one model based only on the target sensor past measurements and another involving just the neighbor sensors. This step verifies if the measurement is part of an expected system behavior or if it is in fact a singular event on the target sensor, due to the existing correlations in the environment.

4 Results

In this section, we present the results obtained when applying the framework to the case study scenario, namely considering the employed prediction models (ANNs) and the outlier detection strategy. We also explore some strategies to implement the Measurement Reassessment component based on the current application of the framework.

For the results herein, we considered the measurements of temperature sensors in the four selected stations (see Fig. 2). We used a three months data set for this validation processing, different from both training and fitting data sets (a few years after both).

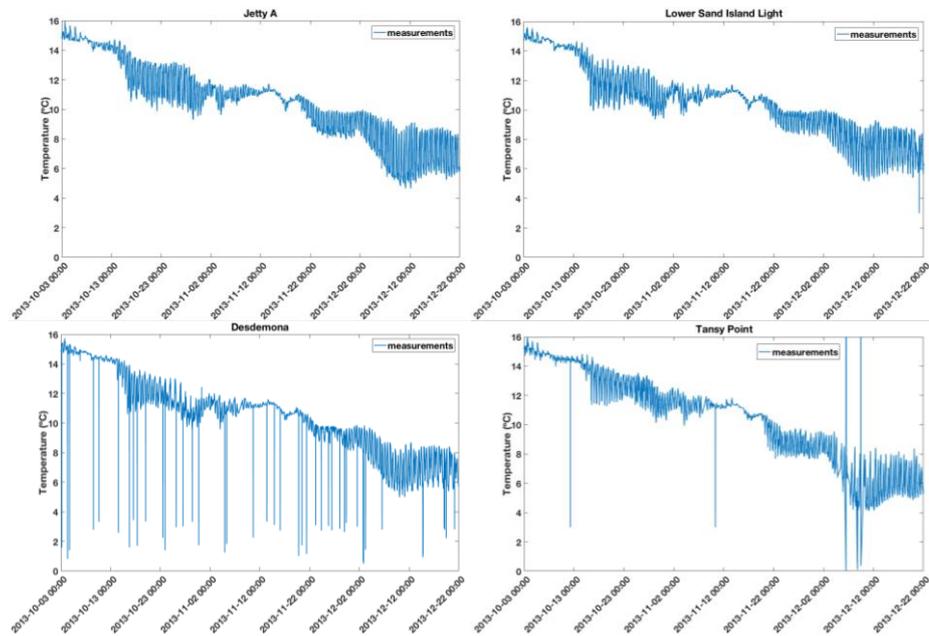


Fig. 2. Sensors measurements for the validation set.

In Fig. 2 it is possible to observe the gradual decrease of the temperature values measured by four different stations from October 2013 to December 2013, and the periodic nature of the “signal”, due to the tidal influence. Additionally, we can state from third-party expert validation that in Jetty A (top left graph) there is no presence of outliers, that in Lower Sand Island Light (top right graph) there is 1 outlier, that in Desdemona (bottom left) 43 outliers were identified and, finally, that in Tansy Point (bottom right) 11 outliers were identified, some of which beyond the represented scale.

In the selected outlier detection strategy, for the verification of the conditions, there is a threshold regarding a confidence level that needs to be defined beforehand, associated with the probability that a new measurement is abnormal. Therefore, it is a higher limit threshold that we verified for accuracy by testing four thresholds: 0.997, 0.998, 0.9985 and 0.999. The results for each threshold are presented in Table 1, where it is

possible to observe how important it is the selection of the confidence threshold for the employed comparison strategy.

Table 1. Number of detected outliers and number of expert-identified outliers (Real).

Threshold	Jetty A	Lower Sd	Desdemona	Tansy
0.997	2	7	71	24
0.998	2	7	55	24
0.9985	1	4	41	21
0.999	1	2	28	15
Real	0	1	43	11

We can observe that while increasing the threshold, the number of detected outliers decreases. For instance, when using a 0.997 threshold, there is a large amount of false positive outliers, contrasting to when setting it to 0.999, in which case number of outliers is lower. In Fig. 3 we highlight the existing trade-off for the Desdemona data set, where it is clear the existence of false positives on the two top graphs and some outliers not detected on the two bottom graphs.

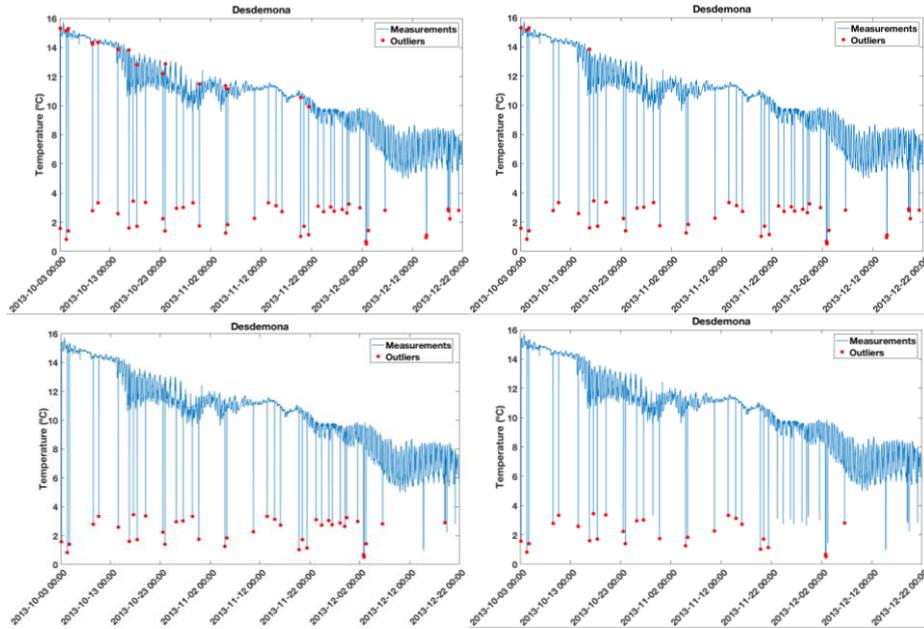


Fig. 3. Outlier detection for Desdemona temperature sensor, with confidence threshold of 0.997, 0.998, 0.9985 and 0.999, from left to right and top to bottom.

Considering both Table 1 and Fig. 3, the best candidate threshold for this application setup is 0.9985, as it provides an overall better outlier detection completeness and at the same time a lower number of false positives.

In addition to failure detection (in this case, outlier detection), there is another important aspect of the framework, namely the Measurement Reassessment (MR). We propose a simple strategy consisting on an ensemble of all the available predictions for the specific case of outliers. This procedure, which involves the replacement of the detected outliers with an average of all the predictions provided by the Prediction component, was used to achieve the results presented in Table 1.

To validate this simple strategy, we tested it against two other strategies using the same candidate threshold (0.9985) and we present the results in Table 2.

Table 2. Number of detected outliers for all sensors according to the MR strategy.

Strategy	Jetty A	Lower Sd	Desdemona	Tansy
No MR	1	3	73	29
Subset Average	16	18	41	24
All Average	1	4	41	21

The first strategy (denoted “No MR”), consists in keeping the original data set intact, allowing the use of erroneous data by the prediction models. The second strategy (denoted “Subset Average”), also involves an ensemble of predictions, but using a subset of the models. The results show that the “All Average” strategy used for MR is the best one, thus confirming that it is important to use every correlation available versus just a subset including the temporal correlations (model based on the target past measurements) and spatial (model based on the other sensors past measurements).

In Fig. 4 we can observe the impact of the MR strategy on the outlier detection for the Tansy Point temperature sensor.

It is possible to observe that when no MR strategy is defined (left graph), many of the false positive situations are preceded by outliers, which is due to their use in past measurements for the prediction models. This problem is avoided when using corrected measurements in the history of past measurement, instead of the actual (incorrect) measurements. This corroborates the importance of the use of a corrected value in post-failure detection phase.

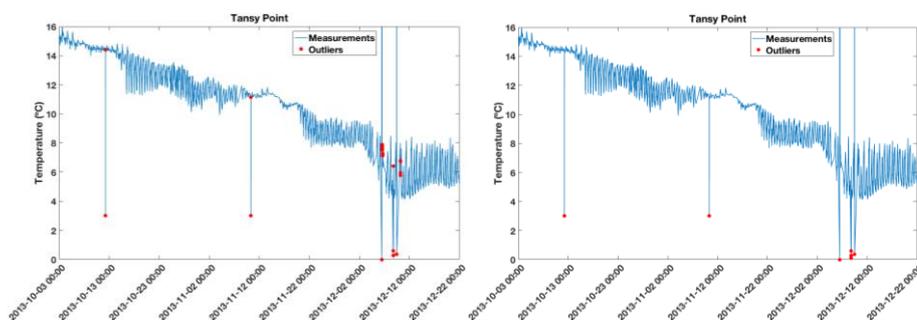


Fig. 4. Outlier detection for Tansy Point temperature sensor, using “No MR” strategy (left graph) and “All Average” (right graph).

5 Conclusions and Future Work Considerations

This paper presents on-going research for the design of strategies for a dependable monitoring framework for harsh environments. Herein, we introduced an outlier failure detection based on two levels of machine learning to provide effective prediction models that explore correlations within the sensor network, and a comparison strategy to ensure the verification of outlier failure conditions.

We also presented the results of experiments with the application of the framework to an aquatic case study, showing the effectiveness of the designed strategies for improving the quality of measurements even though the monitored environment is highly dynamic and impacted by external factors (weather conditions for instance).

Future work will include the design of strategies to detect continuous failure types such as offset and drift (either constant or other type of deviation), considering the framework used in this paper and the same type of prediction techniques (ANNs) to explore the different correlations that exist in the collected data. These strategies will enrich the framework, improving its monitoring capabilities in harsh environments.

Finally, and along with the definition of the new detection strategies, we will define the corresponding procedures performed by Measurement Reassessment and Quality Evaluation components.

Acknowledgements

The authors thank António Baptista and the CMOP SATURN team for their support in the Columbia river analysis. This work was partially supported by the FCT, through the LASIGE Research Unit, Ref. UID/CEC/00408/2013, PhD Grant SFRH/BD/82489/2011 and by H2020 WADI—EC Grant Agreement No. 689239.

References

1. Zhang, Y., Meratnia, N., & Havinga, P.: Outlier detection techniques for wireless sensor networks: A survey. *IEEE Communications Surveys & Tutorials*, 12(2), 159-170 (2010).
2. Zimek, A., Schubert, E., & Kriegel, H. P.: A survey on unsupervised outlier detection in high-dimensional numerical data. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, 5(5), 363-387 (2012).
3. Gupta, M., Gao, J., Aggarwal, C. C., & Han, J.: Outlier detection for temporal data: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 26(9), 2250-2267 (2014).
4. Jesus, G., Casimiro, A., & Oliveira, A.: A Survey on Data Quality for Dependable Monitoring in Wireless Sensor Networks. *Sensors*, 17(9), (2017).
5. Jesus, G., Oliveira, A., Azevedo, A., & Casimiro, A. Improving sensor-fusion with environmental models. *SENSORS*, 2015 IEEE pp 1-4 (2015).
6. Jesus, G., Casimiro, A., & Oliveira, A.: Using machine learning for dependable outlier detection in environmental monitoring systems. (submitted).
7. Baptista, A., Howe, B., Freire, J., Maier, D., & Silva, C. T.: Scientific exploration in the era of ocean observatories. *Computing in Science & Engineering*, 10(3), 53-58 (2008).
8. Pugh, D. T.: *Tides, surges and mean sea-level* (reprinted with corrections). John Wiley & Sons Ltd (1996).