FEATURE

More Effective Analysis of Environmental Sensor Data

By **Mark Patrick** (Mouser Electronics)

Through Automated Cloud-Based Analysis

For the Internet of Things (IoT), tens of billions of nodes are expected to go into operation over the next few years. One of the major uses of this technology will be the monitoring of different key environmental parameters. Finding an efficient way to carry out in-depth analysis on the vast quantities of data continuously being acquired and initiating an appropriate reaction is not that simple. To maximize efficiency and keep response times as short as possible, it will be necessary to put an automated decision-making process in place.

The Internet of Things (IoT) is now starting to see significant implementation, with tens of billions of nodes expected to go into operation over the next few years. As widespread roll-out begins in earnest, one of the major uses of this technology will be the monitoring of different key environmental parameters – such as temperature, air humidity, carbon monoxide levels, etc. Via this captured data, it will be possible to make decisions and identify appropriate actions to take. For example, let's consider a chemical processing plant. If the gas content output from a chimney was to exceed a predetermined threshold, then it might be necessary to reduce the level of processing activity at that time. Alternatively, it may be a sign of a longer-term trend that will need addressing — perhaps the processing equipment involved requires maintenance, or has reached the end of its working life and should be replaced altogether.

Finding an efficient way to carry out in-depth analysis on the vast quantities of data continuously being acquired through the sensor devices incorporated into an expansive network of IoT nodes, then initiating an appropriate reaction is not that simple. Though it may not present a challenge when there is just data coming in from a small number of nodes, coping with the weight of numbers that an industrial control or environmental monitoring system will have is something entirely different. This is equally applicable to a broad spectrum of potential contexts, such as smart homes/buildings, factory automation, agriculture, smart cities, etc.

To maximize efficiency and keep response times as short as possible, it will be necessary to put an automated decision-making process in place – and this will call for the employment of advanced algorithms. Just storing a set of rules on a database then cross-referencing them against the data received isn't effective when you take into account the sheer number of IoT nodes that are likely to be involved. Such an arrangement could quickly become snowed under. A more sophisticated, yet streamlined, approach is therefore needed.

The Rete Algorithm

Some have concluded that the Rete algorithm is the key to this. Already applied to a vast array of business cases, this was first developed back in the late 1970s. The core objective of the Rete algorithm is to provide a pattern matching mechanism through which the vast quantities of pattern data (such as those that IoT networks are now starting to generate) can be rapidly compared against a database containing many different objects. The need for data iteration is avoided, with the status of a particular pattern being kept in a memory reserve throughout the process. The upshot of this is that there is no need to revert to rules that have been applied previously — thereby accelerating the whole procedure and making it much more time efficient than alternative methodologies.

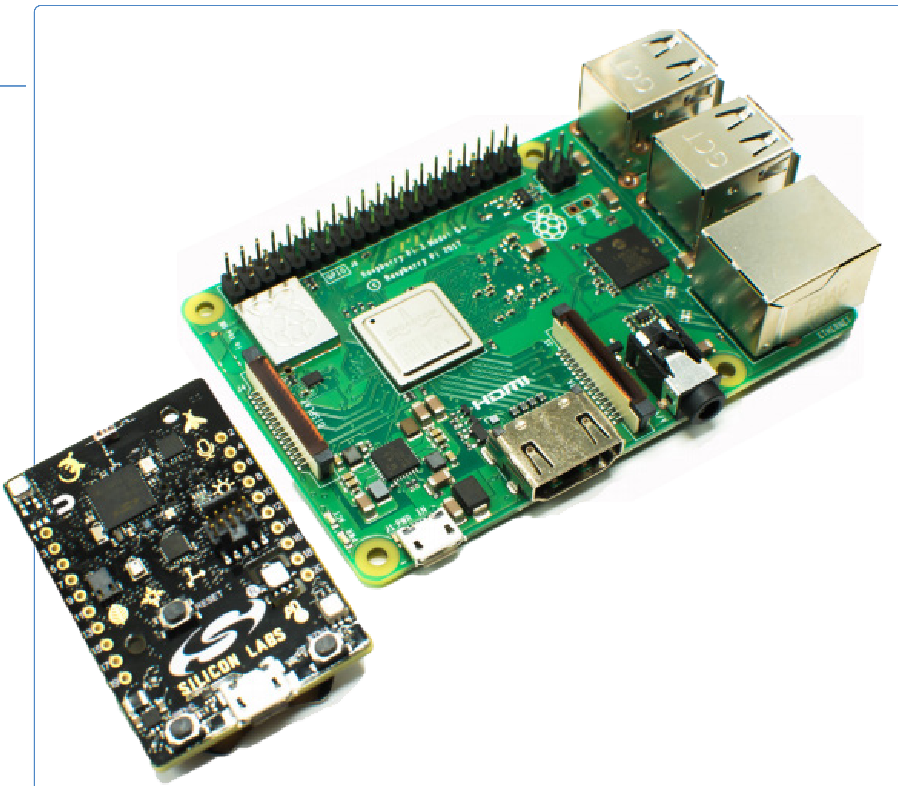

Figure 1: UrsaLeo IoT hardware — comprising a Thunderboard 2 module and a Raspberry Pi 3B+.

About the Author

As Mouser Electronics' Technical Marketing Manager for EMEA, Mark Patrick is responsible for the creation and circulation of technical content within the region — content that is key to Mouser's strategy to support, inform and inspire its engineering audience. Prior to leading the Technical Marketing team, Patrick was part of the EMEA Supplier Marketing team and played a vital role in establishing and developing relationships with key manufacturing partners. In addition to a variety of technical and marketing positions, Patrick's previous roles include eight years at Texas Instruments in Applications Support and Technical Sales.

Operating in conjunction with Silicon Labs' Thunderboard 2 sensor module and the Raspberry Pi 3B+ (which is pre-loaded with Yocto Linux), UrsaLeo's cloud-based analysis software follows this tactic. It enables captured data to be actioned by applying a tree-based rule structure for pattern matching, rather than having to refer time and again to a lookup table. This ensures that better-informed decisions are made without adding unwanted latency to the system. Customized dashboards displaying the compiled datasets can be accessed and triggers set up to alert operatives if certain events occur for which human intervention will be required.

By using the intuitive visually-based editor tool, a set of rules can be established that fully align with the particular application requirements. A trigger can be set every time a message is received, or conversely if a message has not been received within a specified period. Alternatively, a trigger can be initiated at regular intervals (after an hour, a day or a week has elapsed), for long-term monitoring purposes. A geo-fencing rule can also be defined — triggering when a mobile node moves outside/into a given area. This could be used to keep forklift trucks within predetermined confines on the factory floor, or in fleet management and livestock tracking applications. Furthermore, if a captured data value is unusual (e.g., it remains the same constantly or is permanently at zero), then it may suggest that the sensor is not operating correctly, and this can then be flagged so that an engineer can be sent out into the field to undertake any necessary repairs.

A Wealth of Different Sensors

The Thunderboard 2 module can acquire environmental data from the wealth of different sensors that have been incorporated — covering ambient light, air quality, barometric pressure, relative humidity, and temperature, as well as having scope for adding in gas detection functions. A 6-axis inertial sensor (for space orientation) and a Hall effect sensor (for geographical orientation) are also encompassed. The accompanying Raspberry Pi board serves as an IoT gateway (interfacing directly with the Google Cloud platform). Through it, all the compiled data can be transferred back to the cloud via either its wireless (Wi-Fi) or wireline (Ethernet) connectivity options. Depending on the application scenario, the hardware may be powered by USB or using a Li-Ion coin cell battery.

Even with a medium-sized server, the system can process 500,000 messages per second.

The tree-based methodology employed means that this system is fully scalable, with the ability to attend to any prospective number of connected IoT nodes within a network, rather than being limited to a certain quantity. As a consequence, it will be possible for hundreds of thousands of events/alerts to be dealt with every second. Even with a medium-sized server, the system can process 500,000 messages per second, and this could be pushed up to 1-2 million with a high capacity server.

Through the combination of highly optimized cloud-enabled hardware and superior algorithms, the way that IoT data is handled can be made considerably more time efficient than it is today. This will facilitate the ramp-up of IoT deployments to a point where there are hundreds of thousands of connected nodes involved — allowing the true benefits of this technology to be realized across all manner of different industry sectors.

21036-01