



Cacophony: Building a resilient Internet of things

by John Brock and Donald J. Patterson

Abstract

The proliferation of sensors in the world has created increased opportunities for context-aware applications. However, it is often cumbersome to capitalize on these opportunities due to the difficulties inherent in collecting, fusing, and reasoning with data from a heterogeneous set of distributed sensors. The fabric that connects sensors lacks resilience and fault tolerance in the face of infrastructure intermittency. To address these difficulties, we introduce Cacophony, a network of peer-to-peer nodes (CNodes), where each node provides real-time predictions of a specified set of sensor data. The predictions from each of the Cacophony prediction nodes can be used by any application with access to the Web. Creating a new CNode involves three steps: (1) Developers and domain-knowledge experts, via a simple Web UI, specify which sensor data they care about. Possible sources of sensor data include stationary sensors, mobile sensors, and the real-time Web; (2) The CNode automatically aggregates data from the relevant sensors in real time using a JXTA-based peer-to-peer network; and, (3) The CNode uses the aggregated data to train a prediction model via the Weka machine-learning library (Hall, *et al.*, 2009). Real-time predictions made by the CNode are then made publicly available to applications that wish to use data from a CNode's particular set of sensors. The real-time predictions themselves can also be used recursively as sensor data, enabling the creation of CNodes that make predictions based on other CNodes.

Contents

- [1. Introduction](#)
- [2. Network architectures](#)
- [3. Related work](#)
- [4. User interface](#)
- [5. Monitoring a decentralized infrastructure](#)
- [6. Conclusions](#)

1. Introduction

Advancements in manufacturing efficiencies and the subsequent reduction in prices of sensors have caused their deployment to rapidly increase. From barometers on cellphones and motion sensors on light switches in offices, to seismographs in civil infrastructure and smart meters in homes, physical sensors are proliferating. In order to facilitate system maintenance and the analysis of data, many of these sensors are connected to apps and to Internet services. At the same time, software, both simple and sophisticated, is being used to create "virtual" sensors; for example, geographic sentiment analysis derived from Twitter feeds that capture the mood of a city (Kouloumpis, *et al.*, 2011), and search engine aggregations that identify disease outbreaks (Ginsberg, *et al.*, 2009). While the spread of sensors has created increased opportunities for context-aware applications and is tightly coupled with the Internet of Things (Atzori, *et al.*, 2010), there is little accommodation for redundancy in the face of sensor failure or disconnection. As a result, it is difficult to use the opportunities afforded by the scale of this trend when also faced with a future of intermittent infrastructures. There are difficulties inherent in discovering, collecting, fusing, and reasoning with data from the heterogeneous set of distributed infrastructures. Statistical machine learning offers a rich set of techniques for reasoning from large amounts of data, such as that provided by ubiquitous sensors. However, applying these techniques can be difficult because understanding the nuances of machine learning is difficult (Domingos, 2012). Domain expertise is essential to most successful machine learning projects because an expert's experience and insight is critical to choosing which algorithms to use, which data should and should not form the input to the statistical model, and how the structure of that data and

the learning algorithms interrelate. This is assuming that you can even find the relevant data in a form that is amenable to automatic processing. Unfortunately, this expertise is needed over and over again because there is not currently an easy way to make the details of successful models publicly available in a ready-to-use form. There is no large, dedicated public catalogue of trained predictive models from which application developers or researchers can pick-and-choose for their own applications, or upon which they can easily improve. This duplication of effort causes the effective use of sensors to lag far behind their deployment as duplicate teams "reinvent the wheel".

In the face of these trends and challenges, we introduce Cacophony, a network in which computational processes, or "nodes", provide value-added services for networked physical and virtual sensors. Cacophony exists on a peer-to-peer network to provide resiliency against single-point failures that would otherwise prevent access to underlying sensors and provides services to substitute for the sensor if it becomes completely unreachable. Such services include:

1. Predictions: Based on historical observations of sensors and concurrent observations of related sensors ("features" in machine learning parlance), future sensor values can be estimated.
2. Robustness: In the presence of sensor failure or sensor inaccessibility, Cacophony can provide current estimates of a sensor's values based on historical observations and related sensors.
3. Transparency: By asking for the configuration of a node, anyone can evaluate the basis on which predictions are made.
4. Improvement: By minimally modifying the configuration of an existing node and then using it as the basis for a new node, predictive models can accommodate new domain knowledge and new input sensors as features. This can be done with less expertise than initially creating the model may have required.
5. Homogeneity: By creating a consistent interface to sensors on the Internet of Things, Cacophony nodes support developers in the creation of applications; applications only need to implement a single connection to the Cacophony network in order to find any relevant sensor data.
6. Scaling: By virtue of using a decentralized peer-to-peer approach, Cacophony supports extremely large resilient networks of sensors.

In this paper, we give an overview of the Cacophony architecture and user interface, followed by descriptions of two context-aware applications that leverage Cacophony.

2. Network architectures

At a high-level, Cacophony is a peer-to-peer network of nodes and a directory for discovering these nodes. A Cacophony prediction node (CNode) is a predictor for some sensor value known as the "target". Each node regularly retrieves information (known as "features") from some set of sensors, and uses that information to create a learning model for predicting the target. Ground truth is obtained by reading text values from a Web page or RESTful Web service. The learning model is made available via the CNode's Web interface, both on a Web page and through a REST API. In addition to the CNodes, the Cacophony network includes at least one directory node. This directory node provides a public list of the known CNodes, along with relevant information, such as what the CNode is predicting and what sensors it is using to make predictions. This list helps users determine if a certain value is already being predicted, or if they need to make a new CNode to predict the desired value.

2.1. Creating and configuring a new CNode

Launching a CNode process creates new CNodes. A new node can be configured *de novo*, or it can be instructed to copy the configuration of an existing CNode. In either case, configuration can be performed through REST API calls to the running node or via a Web interface provided by the node.

If the CNode is being configured *de novo* via the Web interface (*i.e.*, it is not copying the configuration of an existing node), a user first finds a source containing the ground truth for whatever values she wants to predict (for example, if the user wants to predict the temperature, she must find a Web page that is regularly updated with the temperature). Then, on the CNode's configuration Web page, the user inputs the URL of the ground truth Web page. This will load the ground truth Web page inside the CNode configuration Web page. Once the ground truth page has loaded, the user simply clicks on the desired piece of data, and the CNode records the XPath of that element, *i.e.*, the element's location in the Document Object Model (DOM). The user then selects other sources of data from a list of existing CNodes provided by the Cacophony directory as features from which the model is trained.

The XPath specifies a text element, but it is possible that the actual target or feature of interest is a

substring of what is found at the given XPath. Therefore, if the XPath alone is not precise enough to extract the desired target or feature, the user can supplement it with a regular expression.

If a new CNode is not being configured from scratch, the user can initialize it with the configuration of an existing CNode. The user simply needs to specify the address of the existing node within the peer-to-peer network of CNodes, and the new node will automatically retrieve the appropriate configuration information. This would be a way to rapidly create a sensor that is similar to another existing sensor, or to incorporate changes into the design of an existing sensor in an effort to get better statistical performance.

2.2 Network protocols

Communication within the Cacophony network is accomplished via p2p4java [[1](#)], which is a modification of JXTA (an open source peer-to-peer protocol originally developed by Sun Microsystems). p2p4java offers several advantages over alternative network protocols. For example, consider a smartphone on some wireless carrier's network: we wish to make this phone's sensor readings available to a CNode. One naïve approach is to run a Web server on the phone that provides the sensor values via a REST API. However, wireless carriers typically use firewalls and network address translation (NAT), which can make it impossible to access this Web server from the greater Internet. Fortunately, JXTA offers a solution via relay peers, which enable communication with devices that are behind NATs or firewalls. p2p4java enables this functionality to be extended to Android phones so that a background application on the phone can declare itself a JXTA peer, enabling a CNode to contact it for sensor data.

CNodes can retrieve predictions from other CNodes via the p2p4java network, *i.e.*, a new CNode can use an existing CNode's predictions as features. In this way, the Cacophony nodes can form a directed graph, where outputs (predictions) from some nodes are used as inputs (features) for other nodes. JXTA and our p2p4java extensions are open-source, operate on desktop and Android platforms, and feature advanced cryptographic security.

Internally, Cacophony uses the peer-to-peer network for reliability, robustness, and decentralization. However, HTTP is used by the Cacophony network to retrieve sensor data from the rest of the Web and to make predictions available to applications that don't participate in the peer-to-peer network. Each node in the Cacophony network includes a lightweight Web server, which serves a Web page for making configuration changes and for displaying status information. The Web server also supports a RESTful service for supplying predictions.

2.3. Retrieving sensor data and generating predictions

On a regular basis, each Cacophony prediction node loads the ground truth's Web page to check for changes. If the ground truth's value has changed, the CNode then also retrieves the feature values. Retrieved data is stored in an SQLite database along with all of the previously retrieved values. All of this data is used to train a predictive model via the Weka machine learning library. Whenever a CNode retrieves new sensor data, that node updates its model internally.

Real-time predictions based on the updated model are publicly available to any applications with access to the Web, under the Internet of Things SOA model (Atzori, *et al.*, 2010). Notably, the predictions themselves can also be used as sensor data, enabling the creation of CNodes that make predictions based on other CNodes. In order to prevent loops from creating endless fetching of data, basic caching functions are supported that limit the rate at which features are fetched and predictions are calculated.

In addition to just creating a node using the shared configuration of another node, CNodes can be configured to also share their stored data. This is critical for newly created CNodes, since machine learning algorithms can only produce reasonable predictions once a sufficient amount of training data have been accumulated. Sharing data supports rapid bootstrapping of CNodes.

Note that since users are able to clone and modify whichever CNodes are making the best predictions, an iterative evolutionary process results: the most accurate CNode can be observed and then cloned to create multiple new CNodes, and those clones can then be modified incrementally in an attempt to create a more accurate CNode. In principle, this process can be done automatically, but this is still in the realm of future work. Even as a manual process, this enables the accuracy of CNodes to increase over time.

See [Figure 1](#) for a diagram of the Cacophony network.

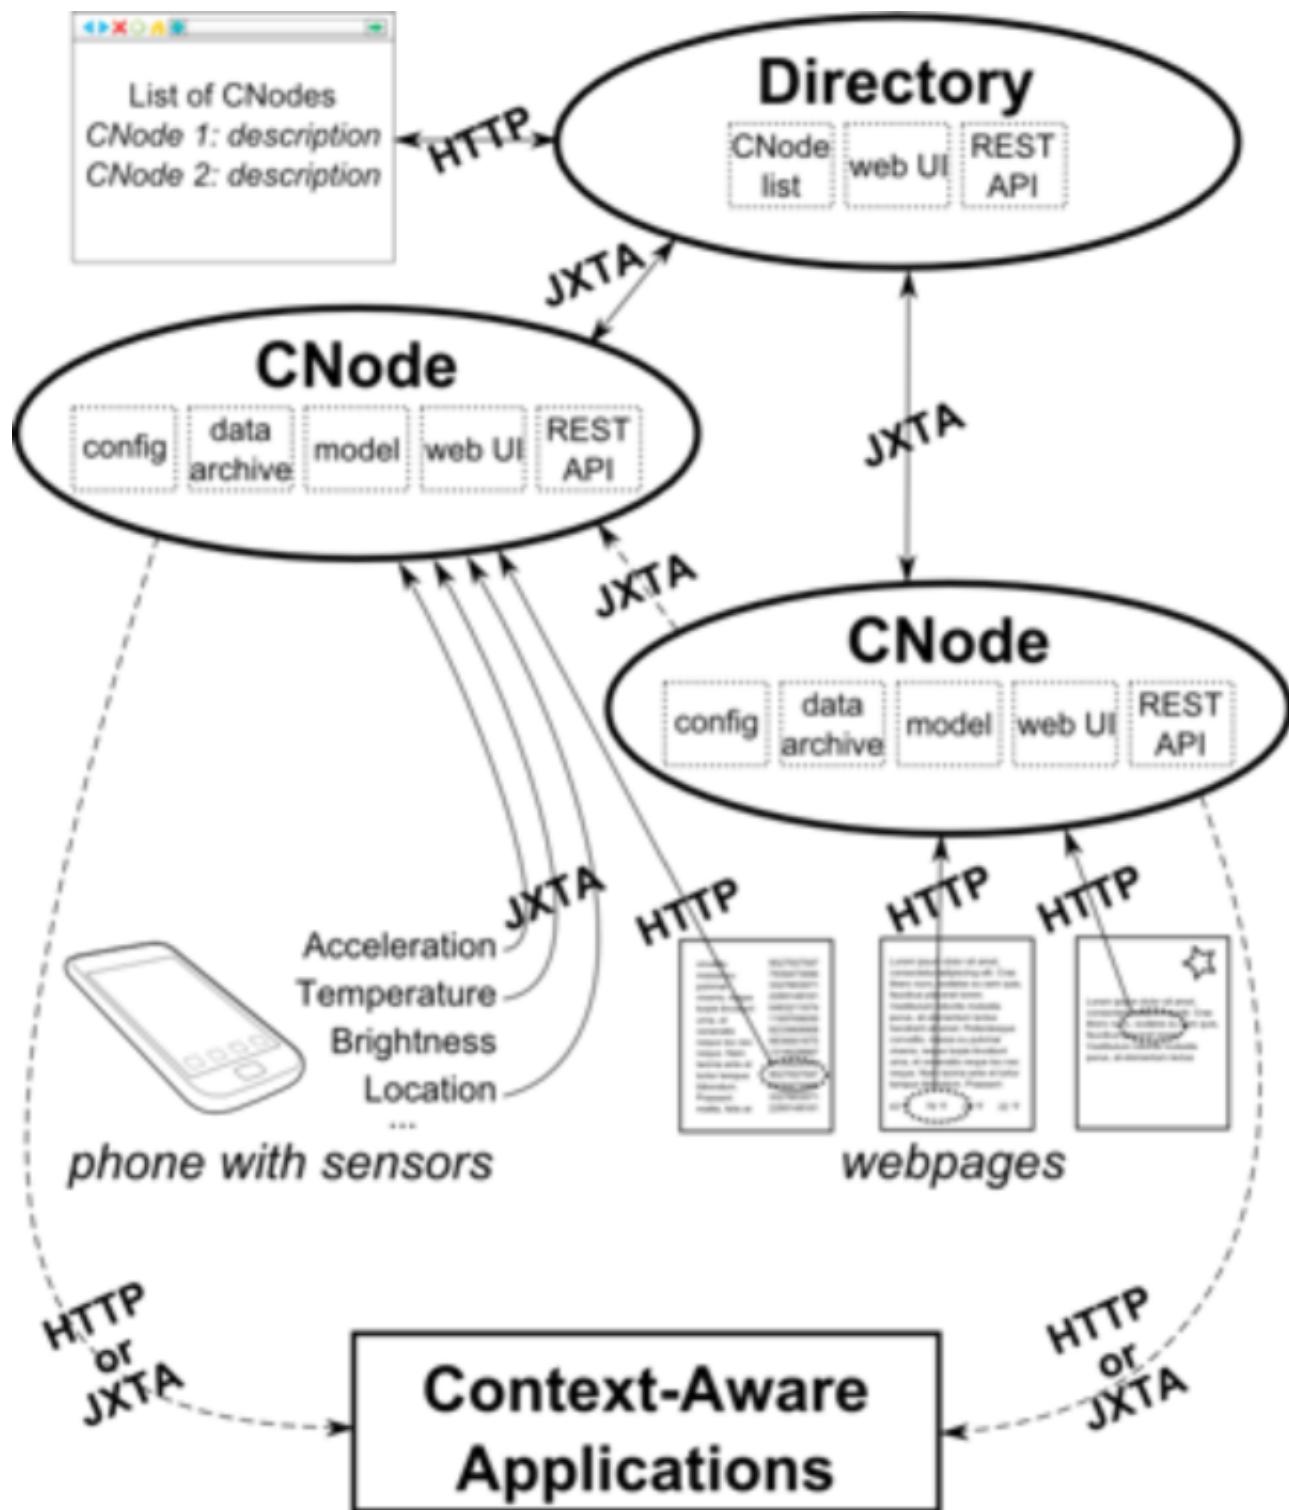


Figure 1: A diagram of the Cacophony network showing the flow of sensor data, predictions, and CNode metadata. Dashed lines are predictions being retrieved from a CNode.

3. Related work

One class of applications that Cacophony was built to support is context-aware applications. Issues of reusability and abstraction when building context-aware applications were previously addressed by Salber with the Context Toolkit (Salber, et al., 1999). Cacophony builds on their ideas by specifying

more specific support for machine learning, providing a mechanism for scaling, and developing an implementation that supports modern smartphones.

Previous work by Heer, *et al.* (2003) on the liquid project addressed ways to create a scalable, distributed system for supporting context-aware computing. Streaming database research inspired their work. Cacophony, in contrast, is designed to function within existing Web protocols. In Cacophony, creating a node that collects data from a sensor requires a user to specify how the data can be found. This requires identifying a location on the Internet where the data resides. For example, if the data is embedded in a Web page, a URL and an HTML-path query (or more generally an XPath query) are needed to retrieve the desired data; these parameters need not be entered explicitly by the user, but can be inferred from user interface actions. There is a large body of programming by demonstration (Paynter, 1996) research that is directly relevant to this process of automatically extracting queries in response to user actions. Several systems have already been built that use programming by demonstration to construct screen-scraping information agents. For example, Bauer, *et al.* built a system to extract text from a Web page based on guidance by the user (Bauer, *et al.*, 2000a; Bauer, *et al.*, 2000b). The system allowed a user to create information agents for retrieving text from Web pages, and used programming by demonstration so that the user did not need to manually specify the exact behavior of these information agents. Their system was also robust enough to handle significant changes to a Web page's layout. When the system could not recover from a change to a Web page's layout, it would alert the user, who could then provide new examples to the information agent. The Internet Scrapbook was another system using programming by demonstration to retrieve text from Web pages (Sugiura and Koseki, 1998). The user specified which part of a Web page he or she was interested in by copying the text in the desired portion to the Scrapbook page. Text from different Web pages could appear on the Scrapbook page, and would reflect updates to the source Web pages. Cacophony incorporates these ideas into its configuration (and maintenance) process.

Alba, *et al.* (2008) also describe various ways for retrieving data from the Web, such as screen scrapers and APIs. When performing screen scraping, there are several ways to refer to part of a Web page, such as the relative positions of elements, regular expressions, and XPath expressions. The Internet Scrapbook used the relative positions of elements, by analyzing the location of headers relative to other text (Sugiura and Koseki, 1998). Regular expressions are a useful way to refer to part of a Web page, but it is easy to under specify or over specify the expression. If the regular expression is too vague, parts of the Web page may be retrieved other than the portion that was desired. If the regular expression is too specific, slight changes in the Web page's text may result in nothing being retrieved. Fortunately, fairly reliable techniques exist for automatically generating regular expressions based on user input, such as blockwise grouping and alignment (Fernau, 2005). XPath expressions pose problems similar to those posed by regular expressions. XPath expressions can be too vague or else overly specific; specifying the appropriate XPath falls under an area of research known as XPath containment (Vion-Dury and Layaïda, 2003). For simplicity, we rely on XPath expressions to specify a feature on a Web page, and require that the feature be directly inside an HTML element with a unique ID attribute.

Baldauf, *et al.* (2007) when discussing various context-aware systems, make a distinction between physical sensors, virtual sensors, and logical sensors. Physical sensors are typically hardware sensors that capture simple physical measurements, such as temperature or acceleration. Virtual sensors report data based on software systems; Baldauf and Dustdar provide the example of using event location entries in an electronic calendar system to determine a person's location (as opposed to using a physical sensor such as a GPS device). Logical sensors provide data based on a combination of data from physical and virtual sensors. Cacophony can consume information from all three types of sensors as long as the sensor's output is available via an HTTP request. Predictions made by Cacophony can be further built upon as they themselves form virtual sensors and/or logical sensors.

4. User interface

The CNodes themselves, once they are configured and running, require no UI for collecting and processing data to generate predictions. Cacophony provides a "directory" user interface tool as a means for people to search for CNodes, to ensure their CNodes are up and running, and to discover CNodes whose sensor data might serve as inputs for their own CNodes' predictions. To accomplish these tasks one needs to know what CNodes exist and what data they provide.

4.1. Methodology

The directory is itself a peer on the p2p4java network. The directory collects information, both passively and actively, about CNodes in the network. As a passive collector, the directory receives information from CNodes that optionally and periodically announce their existence. Each CNode has a default Cacophony directory with which it can register, but in general, anyone can operate his or her own directory. The directory compares the name of any registering CNode against a list of known

CNodes. If no match is found, the directory queues the CNode name in order to retrieve additional information when resources are available. Similar to a Web crawler, the directory visits each CNode in this queue and sends a message requesting the CNode's configuration. The directory parses and indexes all relevant information from this configuration file. Of special note are any CNodes that serve as features, as these extend the frontier of CNodes known by the directory. Any unknown newly discovered CNodes are added to a queue for future crawling. Since CNodes are not required to register themselves and any attempts to register are "best effort", this active crawling serves as a compliment to passively waiting for registrations. Finally, known CNodes are revisited periodically to retrieve any updates to their configuration or dependencies.

4.2. Collection of CNode properties

We categorize the CNodes based on the configuration data retrieved from the CNode registration and crawling. This includes its name on the JXTA network, its features, and metadata about the CNode. Predictions have an associated Java data type (from Weka) and an optional qualitative description. Additionally, the metadata includes machine learning algorithm selections, associated options, guidance on polling frequency, database size and information about node health. Predictions can represent a wide variety of data for example: network throughput, in bytes, for a border router at 137.164.24.49; a barometric reading, in mbars, from 33.684167 to -117.7925; blood sugar levels, in mmol/L, of John Smith; or the NASDAQ Composite index in dollars. The relationships between CNodes are recorded as the flows of predictions from one CNode to another.

4.3. Visualizing the Cacophony network

Once the directory node has collected information about CNodes, it can provide a visualization showing interrelationships between CNodes and the classification of each CNode. The CNode network is visualized as a directed graph, where a node represents each CNode, and an information flow from one CNode to another is represented as a directed edge. The visualization enables one to quickly understand CNode relationships, to gauge the complexity of the network, and to recognize potential configuration errors.

4.4. Finding CNodes

A CNode may be found through a variety of means. The CNode representations can be directly manipulated via the UI or searched for by address structure or attribute. The address structure of a CNode is a URL in the form of p2p://<device_name>/<semanticpath>/<cnode>, entering the device name will highlight all CNodes contained therein or entering the full address will display the CNode itself. Any of the attributes described previously can be searched in a parametric fashion. All CNodes which match the query are highlighted.

4.5. Interacting with CNodes

When a CNode is selected, various actions are possible: retrieving the current prediction or making future predictions, cloning the CNode, observing current node operational health (such as load and fetching frequency), and displaying configuration attributes. Retrieving the current prediction returns the selected CNode's most recent prediction. Cloning a CNode outputs a configuration file and the necessary data for making a duplicate CNode. Displaying configuration attributes returns the extended properties not already shown in the UI. Additional features may include retrieving past predictions or retrieving the values used to produce the current prediction.

5. Monitoring decentralized infrastructure

One proposed application of Cacophony is to use it to monitor alternative decentralized infrastructures (ADIs). Small-scale attempts to create alternative infrastructures (food, energy, water, etc.) enjoy few of the economies of scale of centralized infrastructure (e.g., rooftop gardening vs. agribusiness), but larger efforts typically require more inputs (such as fertilizer and pesticides) that eventually create expensive externalities. Decentralized efforts, however, currently tend to be isolated, inefficient activities that do not address regional, community-based needs. Cacophony could be used to help these ADIs scale up and be operated with intelligent, computer-based management. Resilient, open-source software can be used to help merge these small-scale efforts into a more integrated, regional forms.

Initial prototypes of this software have been developed but substantial work is still needed in order to allow it to monitor real-world alternative infrastructures. What follows are some examples of how Cacophony could support alternative infrastructures.

Organization of existing sensors: This project does not seek to deploy new sensors, but to integrate existing sensors into a network. In the pairing process of a Cacophony node with an

underlying sensor, the sensor is exposed to the rest of the network with a consistent data interface (e.g., API). The node/sensor pairs become searchable through a directory and the complexity of gathering data from now-exposed sensors is greatly reduced. Examples of sensors that we have targeted include smart meters on homes, electric current sensors on solar panels, irrigation system operational status, micro-weather stations, and smartphone/wearable sensors.

Intermittency tolerance: The CNodes communicate using a peer-to-peer networking protocol that is itself resilient to intermittency, but also, as a result, supports continued monitoring in the face of partial infrastructure outages.

Scaling: By using a decentralized peer-to-peer network, the CNodes support extremely large networks of sensors. We are designing based on incorporating approximately 10,000 sensor values into our deployment, many focused on geographic specializations.

Analytics: By incorporating statistical modeling of the sensor values, CNodes can be given predictive analytics that enable them to function on par with utility-grade infrastructures.

Operator reflection: Consistent, reliable access to decentralized sensor systems coupled with high quality user interfaces will enable alternative infrastructure owner/operators insight into how to manage their ADIs more effectively.

Research data: Through monitoring of sensors, we will be able to collect empirical data that can influence future designs.

Applications that are built on top of Cacophony will be able to harvest information from a network of sensors, including physical stationary sensors, mobile and worn sensors, and virtual or physical social sensors and will provide richer insights and potentially support novel use cases.

6. Conclusions

In this paper we have introduced the architecture and vision for Cacophony, a resilient distributed machine-learning layer for the Internet of Things. Cacophony provides value-added services to existing sensors that are accessible on the Web. These services include system architecture enhancements such as sensor robustness, interface homogeneity, and scalability. They also include data-oriented services such as historical sensor logging, statistical analysis of data, and predictive estimates of future readings.

An important administrative tool that Cacophony supports is a directory that functions like a Web search engine for Cacophony-managed sensors. Through this directory, people and applications can discover and monitor what sensor data is available on the Cacophony network. With this knowledge, new services can be created and existing services can be incrementally improved.

We developed Cacophony in response to a practical need to support open source applications that needed to leverage wider portfolios of sensor data than were previously available. For example, previous research supported the automatic inference of context descriptions from the sensors on a single device. Cacophony has allowed us to write applications that determine the context of an infrastructure from a wide range of sensors that span many devices, administrative domains, and intended uses.

In the future we hope to provide more support to users who would like to create their own CNodes in the Cacophony network. This includes automatic inferring of features for target predictions, user interface enhancements for launching CNodes, and broader deployments in support of a variety of applications. Network effects make Cacophony increasingly valuable as the number of sensors available on the Cacophony network increases. 

About the authors

John Brock is a graduate student in the Department of Computer Science at the University of California, Irvine.

Donald J. Patterson is Associate Professor in the Department of Informatics at the Donald Bren School of Information and Computer Sciences at the University of California, Irvine and an Associate Professor at Westmont College in Santa Barbara.

Direct comments to: djp3 [at] ics [dot] uci [dot] edu

Acknowledgments

Portions of this work were funded by grants from the University of California, Irvine's Center for Organizational Research.

Note

1. <https://github.com/djp3/p2p4java>.

References

- A. Alba, V. Bhagwan, and T. Grandison, 2008. "Accessing the deep Web: When good ideas go bad," *OOPSLA Companion '08: Companion to the 23rd ACM SIGPLAN Conference on Object-Oriented Programming Systems Languages and Applications*, pp. 815–818.
doi: <http://dx.doi.org/10.1145/1449814.1449871>, accessed 23 July 2015.
- L. Atzori, A. Iera, and G. Morabito, 2010. "The Internet of Things: A survey," *Computer Networks*, volume 54, number 15, pp. 2,787–2,805.
doi: <http://dx.doi.org/10.1016/j.comnet.2010.05.010>, accessed 23 July 2015.
- M. Baldauf, S. Dustdar, and F. Rosenberg, 2007. "A survey on context-aware systems," *International Journal of Ad Hoc and Ubiquitous Computing*, volume 2, number 4, pp. 263–277.
doi: <http://dx.doi.org/10.1504/IJAHUC.2007.014070>, accessed 23 July 2015.
- M. Bauer, D. Dengler, and G. Paul, 2000a. "Instructible information agents for Web mining," *IUI '00: Proceedings of the Fifth International Conference on Intelligent User Interfaces*, pp. 21–28.
doi: <http://dx.doi.org/10.1145/325737.325758>, accessed 23 July 2015.
- M. Bauer, D. Dengler, G. Paul, and M. Meyer, 2000b. "Programming by demonstration for information agents," *Communications of the ACM*, volume 43, number 3, pp. 98–103.
doi: <http://dx.doi.org/10.1145/330534.330547>, accessed 23 July 2015.
- P. Domingos, 2012. "A few useful things to know about machine learning," *Communications of the ACM* volume 55, number 10, pp. 78–87.
doi: <http://dx.doi.org/10.1145/2347736.2347755>, accessed 23 July 2015.
- H. Fernau, 2005. "Algorithms for learning regular expressions," In: S. Jain, H.U. Simon, and E. Tomita (editors). *Algorithmic learning theory. Lecture Notes in Computer Science*, volume 3734. New York: Springer, pp. 297–311.
doi: http://dx.doi.org/10.1007/11564089_24, accessed 23 July 2015.
- J. Ginsberg, M.H. Mohebbi, R.S. Patel, L. Brammer, M.S. Smolinski, and L. Brilliant, 2009. "Detecting influenza epidemics using search engine query data," *Nature*, volume 457, number 7232 (19 February), pp. 1,012–1,014.
doi: <http://dx.doi.org/10.1038/nature07634>, accessed 23 July 2015.
- M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I.H. Witten, 2009. "The WEKA data mining software: An update," *ACM SIGKDD Explorations Newsletter*, volume 11, number 1, pp. 10–18.
doi: <http://dx.doi.org/10.1145/1656274.1656278>, accessed 23 July 2015.
- J. Heer, A. Newberger, C. Beckmann, and J.I. Hong, 2003. "liquid: Context-aware distributed queries," In: A.K. Dey, A. Schmidt, and J.F. McCarthy (editors). *UbiComp 2003: Ubiquitous computing. Lecture Notes in Computer Science*, volume 2864. New York: Springer, pp. 140–148.
doi: http://dx.doi.org/10.1007/978-3-540-39653-6_11, accessed 23 July 2015.
- E. Kouloumpis, T. Wilson, and J Moore, 2011. "Twitter sentiment analysis: The good the bad and the OMG!" *Fifth International AAAI Conference on Weblogs and Social Media*, pp. 538–541, and at <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM11/paper/view/2857>, accessed 23 July 2015.
- G.W. Paynter, 1996. "Generalising programming by demonstration," *Sixth Australian Conference on Computer-Human Interaction. Proceedings*, pp. 344–345.
doi: <http://dx.doi.org/10.1109/OZCHI.1996.560161>, accessed 23 July 2015.
- D. Salber, A.K. Dey, and G.D. Abowd, 1999. "The context toolkit: Aiding the development of context-enabled applications," *CHI '99: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 434–441.
doi: <http://dx.doi.org/10.1145/302979.303126>, accessed 23 July 2015.

A. Sugiura and Y. Koseki, 1998. "Internet scrapbook: Automating Web browsing tasks by demonstration," *UIST '98: Proceedings of the 11th Annual ACM symposium on User Interface Software and Technology*, pp. 9–18.

doi: <http://dx.doi.org/10.1145/288392.288395>, accessed 23 July 2015.

J-Y. Vion-Dury and N. Layaïda, 2003. "Containment of XPath expressions: An inference and rewriting based approach," *Proceedings of the Extreme Markup Languages*, at <https://hal.inria.fr/inria-00423413/document>, accessed 23 July 2015.

Editorial history

Received 15 July 2015; accepted 23 July 2015.



This paper is licensed under a [Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License](#).

Cacophony: Building a resilient Internet of things
by John Brock and Donald J. Patterson.

First Monday, Volume 20, Number 8 - 3 August 2015
<http://firstmonday.org/ojs/index.php/fm/rt/printerFriendly/6130/4849>
doi: <http://dx.doi.org/10.5210/fm.v20i8>.