

Human in the Loop: Distributed Data Streams for Immersive Cyber-Physical Systems

Anthony D. Wood, John A. Stankovic
Department of Computer Science, University of Virginia
wood@cs.virginia.edu, stankovic@cs.virginia.edu

Abstract

Emerging research in cyber-physical systems (CPSs) often leaves out a key component—the ordinary user. We present requirements for Immersive CPSs, in which people interact with their local environments, and describe a design for distributed data stream creation and sharing.

1. Introduction

Wireless sensor network (WSN) design is converging toward the encompassing vision of cyber-physical systems, characterized by the seamless integration of computation, communication, sensing, and control. To date, WSNs have been dominated by scientific and engineering applications, with a paradigm of 1) deploying sensors, 2) collecting data, and 3) analyzing it with application-specific interfaces [1].

These finely controlled, restricted-access WSNs for personal or institutional use are important but limited first steps toward the future. Just as networking before the Internet came of age was limited mostly to business LANs (for inter-office file and printer sharing) and subscriber-only dial-up portals (for structured content), these systems are disconnected realms with privileged users who already know the extent and capabilities of the network. Data for each is collected in its own relational database, which provides the abstraction of sensors as infinite-length tables to be processed in largely traditional ways [2].

As sensor networks evolve and become ubiquitous, new opportunities will arise for ordinary people to richly interact with their environments through cyber-physical systems. Current data management patterns mediated by centralized databases will be dwarfed by user-directed access to ambient sensors via embedded interfaces. Users will be immersed not in a virtual reality, but in their actual physical world exposed to view in unprecedented detail.

Immersive Cyber-Physical Systems (ICPS) require novel research in architectures and protocols that are:

- **Distributed:** access to data is as widely distributed as is its production by sensors. Many point to point data streams overlay the network, coexisting with tree-based aggregation and collection in the background.
- **Open:** networks and services are open for use by residents, visitors, workers, and passers-by, without centrally controlled authorization.

- **Interactive:** data acquisition, processing, delivery, and visualization are driven by people—users, not administrators or scientists—for ad hoc social interactions and personal environmental investigations.
- **Discoverable:** sensor data streams and services are self-advertising and self-describing for discovery and re-use by nearby users.
- **Stochastic:** demand for sensor data and actuation is user-driven and unpredictable, straining energy-efficiency designs that assume strong periodicity.
- **Heterogeneous:** applications, user interfaces, sensor and actuator modalities, data flows, and usage patterns are diverse and evolving.
- **Localized:** the system advertises and delivers data preferentially to local users, reflecting the physical constraints of interactions with the environment.

The health-care domain presents many promising applications for Immersive CPSs, not just for disease management but also for personal health and vitality enhancement.

For example, a body network provides health enhancement by monitoring and interacting with the wearer, while discovering sensors in the local environment. Current vital signs, recent activity levels, and learned mood reactions are fused with local data on environmental pollutants, noise, or nearby visual attractions to suggest alternate walking routes for leisure or urban commute, or to find a restaurant with available outdoor seating and healthy food options.

Or perhaps an elder in assisted-living receives a package of sensors to wear or place in her apartment, based on a doctor's assessment of her health. The doctor synthesizes a virtual sensor for fall risk and shares it with the resident's wrist-worn display. Using data from mobility speed, transfers, postural hypotension, unevenness in gait, and location (near stairs or obstacles), it alerts the wearer and nearby health-care staff of an elevated risk of falling.

SenQ is a data streaming service for WSNs designed to support immersive, user-driven applications [3]. It provides snapshot and streaming access to data, runtime sensor discovery, and is extensible to new sensor and processing modules. Virtual sensors use hierarchical composition to hide custom logic and processing behind SenQ's query abstraction, and may be shared with other users' embedded UIs. It is compact enough to coexist with an embedded GUI library and data stream display application on a MicaZ mote.

2. SenQ Data Streams


Four design principles underlie the design of SenQ and enable next generation cyber-physical sensor networks.

Human in the Loop

A WSN typically performs background data collection to support automatic control and expert users' needs. Periodic reports may be hierarchically aggregated and stored in the back-end, with control logic sending actuation directives as needed. But these mechanisms are not well-suited for embedded interactivity with ordinary people.

In SenQ, devices find surrounding sensors by broadcasting discovery requests. These include filter criteria for device types or IDs and sensor modes, to reduce unwanted or repetitive replies. Rather than require directory services, the discovery is rooted in a k -hop neighborhood around the requester. Local storage is used to remember sensors that have been queried, along with their locations.

Users (or their cyber-agents) initiate queries for data streams from discovered or previously known sensors using embedded user interfaces. SenQ's implementation is lightweight and efficient enough to fit on low-capacity motes, and on other mobile devices such as cell phones and PDAs. The SeeQuery graphical application is shown at right.



Autonomous Actuation

Keeping data and control logic in-network, close to the concerned devices, saves energy and preserves scalability. SenQ supports autonomous actuation with a parametric software API for querying local or remote sensors and processing the resulting data streams using in-network computation. Application-specific embedded code controls actuation, ideally without back-end input.

Spatial aggregation is supported using coordinated queries, where one sensor is designated as the coordinator or aggregator. It re-broadcasts the query up to k -hops, and collects low-level samples from its eligible neighbors. These feed directly into the coordinator's local processing/filtering chain before being streamed to the originator.

Autonomous in-network query issue allows fine-grained self-monitoring (of battery level, sensing irregularities, etc.), that improves robustness of the CPS through proactive notification of pending failure or data quality degradation.

Composable Virtual Sensors

To foster distributed, community use and improvement of ICPSs, users must be able to tag interesting content, and fuse existing streams together into new, virtual sensors.

Standard interfaces wrap sensor types and expose metadata, easing extension to new devices and applications. This lets sensor drivers recursively use SenQ to issue local or remote queries in complex ways, transparently to other users. For example, a doctor may create a virtual Fall sensor for a body network, that combines limb acceleration, body position, and pulse rate changes to detect whether the wearer has fallen, and gauge the severity of the event.



Figure 1. Sample of supported sensors.

Virtual sensors may be bound to the devices involved, so that they persist for other users to discover and access, even if the original creator is no longer present.

Mobile Streams

An ICPS should foster human-to-human collaboration and data sharing. In SenQ, a user can share a data stream by dynamically adding another user's device to the stream recipients. The original user can drop out of the stream while letting it persist, effectively pushing the stream to other devices. For example, a doctor creates a virtual Fall sensor as before, and pushes it to an assisted-living nurse's station for continuous monitoring. He also copies the data to a database for archival storage. A virtual Fall Risk sensor is pushed to the resident's wrist-worn display, providing real-time feedback to prevent falls before they occur.

A prototype of SenQ has been developed for MicaZ motes that supports diverse sensor hardware (see Figure 1) for medical and environmental data. An embedded SeeQuery application provides in-situ display for reminders and sensor data on the SeeMote, a hardware module with a color LCD screen, buttons, and removable storage.

Immersive CPSs raise other questions for researchers to address. Large-scale networks with overlapping administrative domains bring trust issues to the fore. How can users gain confidence in raw and virtual sensors created by other users? Are authorization policies to control the creation and publishing of data streams needed, and how are they managed in a distributed network? Streams of variable-quality and availability will have adverse effects on control loops in autonomous actuators, requiring robust control approaches.

The nascent field of Immersive Cyber-Physical Systems requires the integration of computation, communication, sensing, control, and human interaction. We presented a design for people-centric sensor stream management that accommodates the diverse requirements of ICPSs, and outlined some open research questions. Support for persistent virtual sensors and mobile streams in SenQ is under-way, and a demo of the prototype is available.

References

- [1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. Wireless sensor networks: A survey. *Computer Networks*, 38(4):393–422, Mar. 2002.
- [2] S. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. TinyDB: An acquisitional query processing system for sensor networks. *ACM TODS*, 30(1):122–173, 2005.
- [3] A. Wood, L. Selavo, and J. A. Stankovic. Senq: An extensible query system for streaming data in heterogeneous interactive wireless sensor networks. In submission.