SensorBench: Benchmarking Approaches to Processing Wireless Sensor Network Data

Ixent Galpin[1], Alan B. Stokes[2], George Valkanas[3], Alasdair J. G. Gray[4], Norman W. Paton[2], Alvaro A. A. Fernandes[2], Kai-Uwe Sattler[5], Dimitrios Gunopulos[3]

SSDBM 2014
Ålborg, Denmark

[1]Universidad Jorge Tadeo Lozano, Colombia
[2]University of Manchester, UK
[3]University of Athens, Greece
[4]Heriot-Watt University, UK
[5]Ilmenau Institute of Technology, Germany

Wireless Sensor Networks (WSNs)

- Over last decade, used to monitor broad range of phenomena
  - Bird habitat monitoring
  - Volcanic activity
  - Glacier movement
  - Sniper localization
  - ...
- Tool to obtain data **cost-effectively** at higher spatial and temporal resolutions
- Scarce resources
  - Limited energy, memory and computational power
  - Trade-offs due to conflicting QoS requirements
- Intelligent
  - Nodes able to carry out data processing
  - **In-network processing** may yield tangible benefits

Data Processing in WSNs

Three broad categories, with different degrees of in-network processing and repurposability:

- **Warehousing** approach
  - Ship all raw sensor readings out of the WSN
  - Example: MultihopOscilloscope [6]
- **Bespoke, hand-crafted** approach
  - WSN carries out a fixed task
  - Examples: D3 outlier detection [10], LR linear regression
- **Sensor network query processing** (SNQP) approach
  - WSN evaluates *ad hoc* user-specified queries
  - Examples: TinyDB [8], AnduIN [4] and SNEE [2]

SensorBench: Why do we need it?

- Many different proposals for data processing techniques → complex design space
- Individual publications evaluate different
  - Tasks
  - Network topologies
  - Performance metrics
  - ...for a particular platform
- How to compare results?

SensorBench: What is it?

- Benchmark to enable comparison of data processing techniques that operate over **wireless sensor networks** (WSNs)
- Consists of workloads designed to:
  - Explore the **variables** (and associated trade-offs) within the complex design space of WSN deployments
  - Provide diverse **performance metrics** pertinent to a broad range of WSN application scenarios
- Scripts and instructions available at `http://code.google.com/p/sensorbench`

Paper Contributions

- Identification of **variables**, **tasks** and **performance metrics** that represent functional and non-functional requirements of WSN applications
- Specification of **workloads** that capture trade-offs inherent in WSN deployments
- Application of benchmark to analyse several different data processing techniques

Desiderata

- Aimed at **environmental monitoring applications**
  - Nodes at fixed locations, data sensed at regular intervals, energy is scarce, single gateway node
- Platform-agnostic
- Use of simulation
  - Allows systematic experimentation that covers broader region of WSN design space in efficient manner
- Agnostic about adaptivity
- Important benchmark properties include **relevance**, **portability**, **scalability** and **simplicity**

Variables

Acquisition interval	Amount of time between sensor readings	▶ Almost continuous ▶ Moderate (5-60 min) ▶ Very infrequent (4 hours)
Network size	Number of nodes in the WSN deployment	▶ Small (2-10) ▶ Medium (11-30) ▶ Large (30+)
Node layout	Spatial distribution of nodes throughout WSN	▶ Linear ▶ Grid ▶ Arbitrary
Node density	Measure of how close nodes are to one another	▶ Sparse topology ▶ Dense topology
Proportion of sources	Percentage of WSN nodes that have sensors	▶ Likely to be high to minimize costs
Radio packet loss rate	Percentage of radio packets not received successfully	▶ Average 30% reported in GDI deployment

Performance Metrics

Lifetime (days)	Amount of time taken for WSN to be unable to carry out data processing task due to energy depletion
Total energy consumption (Joules)	Sum of energy consumed by all nodes in the WSN
Delivery fraction (%)	Percentage of tuples delivered to the gateway of the total that could be delivered
Delivery delay (s)	Time elapsed between event occurring in environment and event being reported
Output rate (bytes/s)	Amount of data produced by the system per unit time.

Example Application Scenario

- Based on Great Duck Island deployment, a classical WSN application [11]
- Aim to monitor nesting patterns of Leach's Storm Petrel and micro-climatic conditions

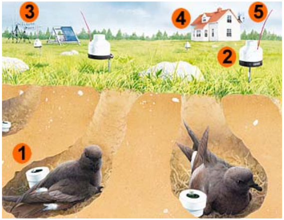

Source: wired.com

The following schema is assumed:

```
surface(node_id, time, light, temp, humidity)
burrow(node_id, time, light, temp, humidity)
```

Tasks

Select	Report raw data readings from the nodes in the WSN
Aggr	Report the average temperature readings for the current time
Join	Correlate data from different regions of the WSN
Join2	Correlate data from different regions of the WSN collected at different times
LR	Linear regression
OD	Outlier detection

Example of Join2 task expressed using a SNEEql query:

```
RSTREAM SELECT b.node_id, b.temp
FROM burrow[NOW] b, surface[NOW-1 MINUTE] s,
WHERE b.temp > s.temp;
```

SensorBench Workloads

Varying

1. network size
2. network layout
3. node density
4. acquisition interval
5. proportion of sources
6. radio loss rate
7. task

Running the Benchmark

- Sensor datafiles and topologies can be downloaded from `http://dx.doi.org/10.6084/m9.figshare.934307`
- Scripts to run jobs on Avrora emulator [13]
  - Optionally using HTCondor parallel computing platform [12]
- Scripts to parse total energy consumption, lifetime, output rate, delivery fraction and delivery delay from Avrora log files
- We ran it against MultihopOscilloscope, LR, OD, SNEE
- 10 topologies generated for each combination of ⟨*Network Size, Node Layout, Node Density, Proportion of Sources*⟩

Varying Network Size

Variable	Values
Tasks	{Select, Aggr, LR, OD}
Acquisition interval	32
Network size	{9, 25, 100}
Node layout	arbitrary
Node density	3
Proportion of sources	80
Radio loss rate	0

3 topology sizes $\times$ *4 tasks* $\times$ *10 topologies per topology size* $=$ **120 simulations***!*

Network Size vs. Delivery Fraction

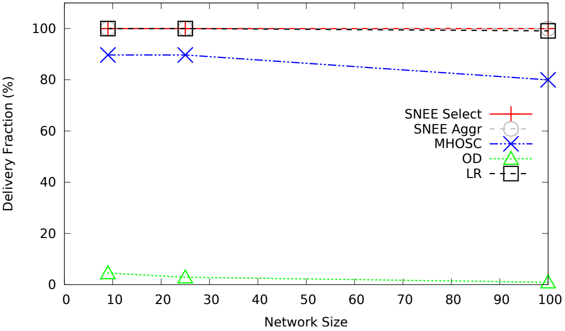

Network size vs. Delivery Delay

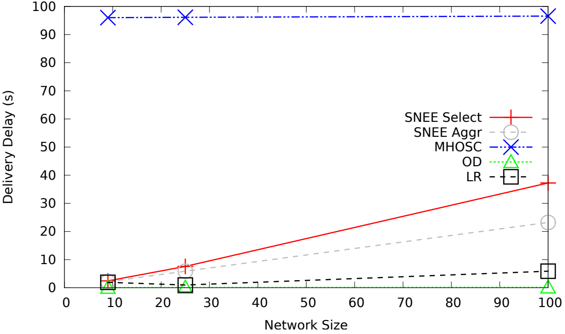

Varying Network Layout

Variable	Values
Tasks	{Select, Aggr, LR, OD}
Acquisition interval	32
Network size	25
Node layout	{linear, grid, arbitrary}
Node density	3
Proportion of sources	80
Radio loss rate	0

Node Layout vs. Lifetime

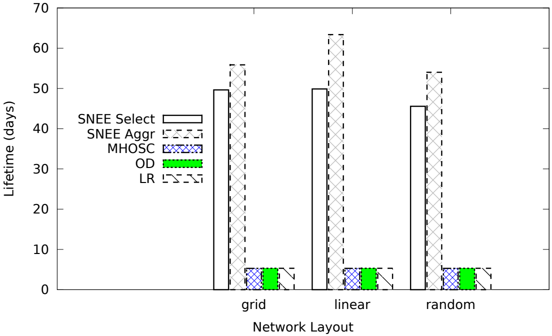

Varying Acquisition Interval

Variable	Values
Tasks	{Select, Aggr, LR, OD}
Acquisition interval	{1, 2, 4, 8, 16, 32, 64, 128}
Network size	25
Node layout	arbitrary
Node density	3
Proportion of sources	80
Radio loss rate	0

Acquisition Interval vs. Delivery Delay

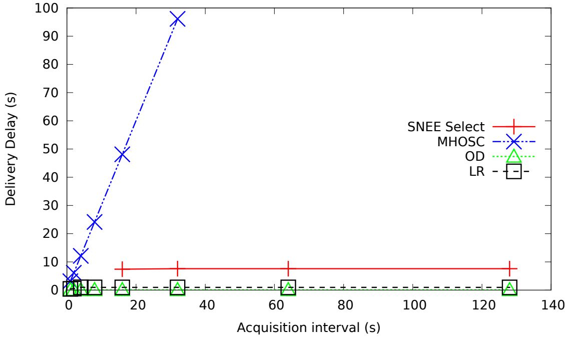

Related Benchmarks

- Stream Data management
  - Linear Road benchmark [1]
- Wireless Sensor Networks
  - Devices (TinyBench [3])
  - Processors (SenseBench [9])
  - Cryptographic algorithms [5]
  - Communications (LinkBench [14])
- Bisque [7] is a proposals for a WSN query processing benchmark
  - We cover more varied variables, tasks and metrics

Evaluations of Sensor Data Management Systems: Variables

Proposal	Acquisition interval	Node layout	Node density	Network size	Proportion of Sources	Packet loss rate	Other
SensorBench	•	•	•	•	•	•	
TinyDB	•						Selectivity, Time
AnduIN							Time, Window size
MicroPulse			•				Time
SNEE	•	•					Delivery Time
Aspen		•					Selectivity, Window size, Time
Bisque				•			Selectivity

Evaluations of Sensor Data Management Systems: Metrics

Proposal	Network energy	Lifetime	Delivery fraction	Delivery delay	Output rate	Other
SensorBench	•	•	•	•	•	
TinyDB		•	•		•	Maintenance overhead
AnduIN		•				Computation time
MicroPulse		•				
SNEE	•	•				Memory Usage
Aspen					Network traffic, Node load	
Bisque			•	•		Node Energy Consumption

Evaluations of Sensor Data Management Systems: Tasks

Proposal	Select	Aggr	Join	Regression	Outlier Detection
SensorBench	●	●	●	●	●
TinyDB	●	●			
AnduIN	●	●	●		●
MicroPulse	●				
SNEE	●	●	●		
Aspen			●		
Bisque	●	●			

Evaluations of Sensor Data Management Systems: Tasks

- SensorBench provides means to perform descriptive and comparative analysis of broad range of WSN data processing proposals
  - relevance, portability, scalability and simplicity
- Subsumes most relevant empirical analysis in terms of scope while remaining simple to run
- Scripts provided to facilitate implementation of the benchmark using popular simulator

References

- [1] A. Arasu, M. Cherniack, E. F. Galvez, D. Maier, A. Maskey, E. Ryvkina, M. Stonebraker, and R. Tibbetts. Linear road: A stream data management benchmark. In *VLDB*, pages 480–491, 2004.
- [2] I. Galpin, C. Y. A. Brenninkmeijer, A. J. G. Gray, F. Jabeen, A. A. A. Fernandes, and N. W. Paton. SNEE: a query processor for wireless sensor networks. *Distrib. Parallel Dat.*, 29(1-2):31–85, Nov. 2011.
- [3] M. Hempstead, M. Welsh, and D. Brooks. TinyBench: The case for a standardized benchmark suite for TinyOS based wireless sensor network devices. In *LCN*, pages 585–586, 2004.
- [4] D. Klan, M. Karnstedt, K. Hose, L. Ribe-Baumann, and K.-U. Sattler. Stream engines meet wireless sensor networks: cost-based planning and processing of complex queries in AnduIN. *Distrib. Parallel Data.*, 29(1-2), 2011.
- [5] Y. W. Law, J. Doumen, and P. Hartel. Survey and benchmark of block ciphers for wireless sensor networks. *TOSN*, 2(1):65–93, 2006.
- [6] P. Levis, S. Madden, J. Polastre, R. Szewczyk, K. Whitehouse, A. Woo, D. Gay, J. Hill, M. Welsh, E. Brewer, et al. TinyOS: An operating system for sensor networks. In *Ambient intelligence*, pages 115–148. Springer, 2005.
- [7] Q. Luo, H. Wu, W. Xue, and B. He. Benchmarking in-network sensor query processing. Technical Report HKUST-CS05-09, Department of Computer Science, HKUST, 2005.
- [8] S. Madden, M. J. Franklin, J. M. Hellerstein, and W. Hong. TinyDB: an acquisitional query processing system for sensor networks. *ACM Trans. Database Syst.*, 30(1):122–173, 2005.

References

- [9] L. Nazhandali, M. Minuth, and T. Austin. SenseBench: toward an accurate evaluation of sensor network processors. In *Proc. Int. Workload Characterization Symp.*, pages 197–203. IEEE, 2005.
- [10] S. Subramaniam, T. Palpanas, D. Papadopoulos, V. Kalogeraki, and D. Gunopulos. Online outlier detection in sensor data using non-parametric models. In *VLDB*, pages 187–198, 2006.
- [11] R. Szewczyk, A. M. Mainwaring, J. Polastre, J. Anderson, and D. E. Culler. An analysis of a large scale habitat monitoring application. In *SenSys*, pages 214–226, 2004.
- [12] D. Thain, T. Tannenbaum, and M. Livny. Distributed computing in practice: the Condor experience. *Concurrency - Practice and Experience*, 17(2-4):323–356, 2005.
- [13] B. Titzer, D. K. Lee, and J. Palsberg. Avrora: scalable sensor network simulation with precise timing. In *IPSN*, pages 477–482, 2005.
- [14] K. Veress and M. Maroti. Linkbench: Benchmark and metric framework for wireless sensor networks. In *IPSN*, pages 171–172, 2011.