Universität Stuttgart
Institute of Parallel and
Distributed Systems (IPVS)
Universitätsstraße 38
D-70569 Stuttgart

Identifying the Challenges in Reducing Latency in GSN using Predictors

Andreas Benzing, Klaus Herrmann, Boris Koldehofe and Kurt Rothermel

WowKiVS - March 6th 2009

Outline

- Motivation
- System Description
- Predictors
- Challenges
- Summary

Motivation

- Simulations depend on accurate and timely sensor information
  - Weather Monitoring
    - Early detection of hurricanes
  - Seismic Monitoring
    - Volcanic Eruptions
    - Earthquakes
  - Traffic Monitoring
    - Supply Chain Management

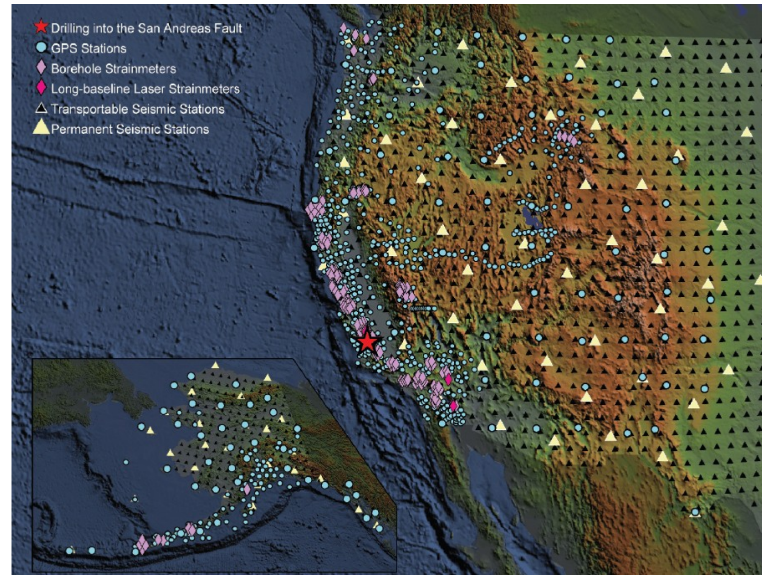

www.earthscope.org Source: Matt Welsh

- Global Sensor Networks produce a huge amount of data
  - In-network processing and filtering reduce data, but also incur additional latency overhead

Distributed Stream Processing

- Sensors produce data streams
- Applications pose queries as extended SQL statements
- Approach
  - Decompose each query to allow partial results
  - Distribute query in network
  - Results in overlay network formed by logical point-to-point connections
- Challenges
  - Highly bursty traffic
  - Placement of operators to achieve low latency

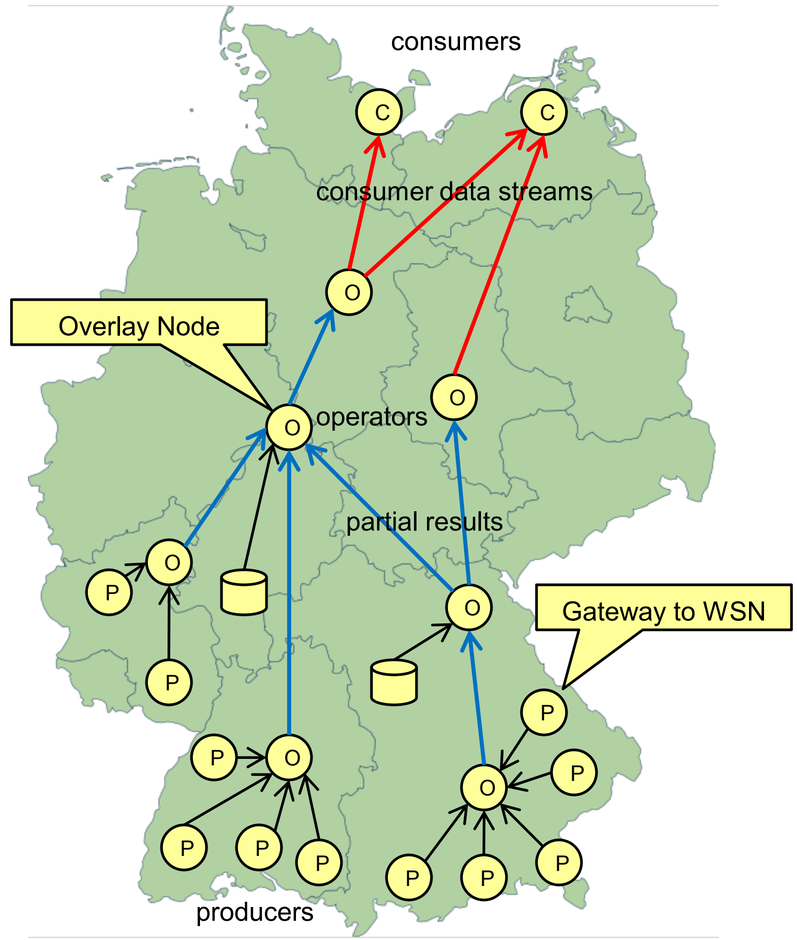

Outline

- Motivation
- System Description
- Predictors
  - Working Principle
  - Classification of Predictors
- Challenges
- Summary

Working Principle

- Future measurements can be estimated from preceding ones
- Nodes in GSN negotiate on value predictors
  - Nodes exploit local knowledge to create predictors or
  - Application pushes predictors into GSN
- If a predictor is running, updates are only sent on deviation
  - Reduction of bandwidth usage
- Current value can be instantly estimated on sink
  - Reduction of latency

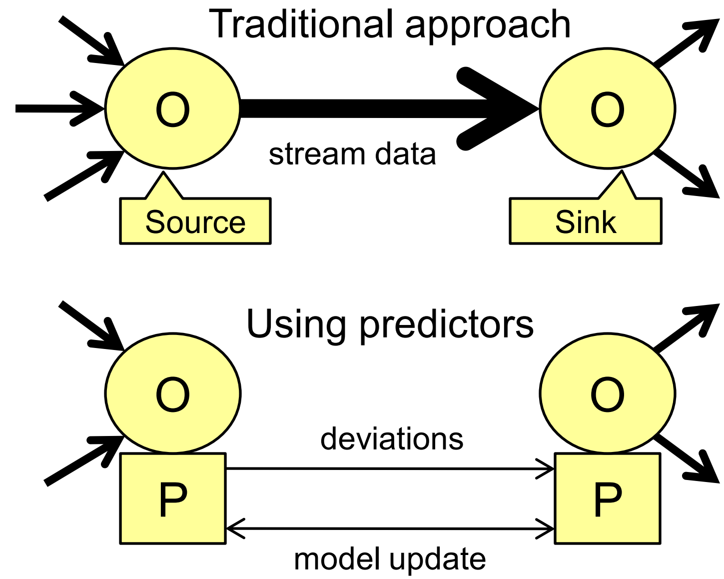

Temporal Predictors

- Synchronous
  - Next sensor reading is predicted using the same model on source and sink

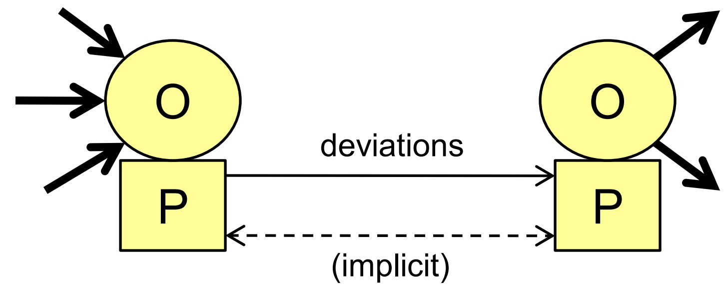

- Asynchronous
  - Prediction model is generated on the sink and sent to source where it is validated
  - Prediction model is derived on source and only model update is sent to sink

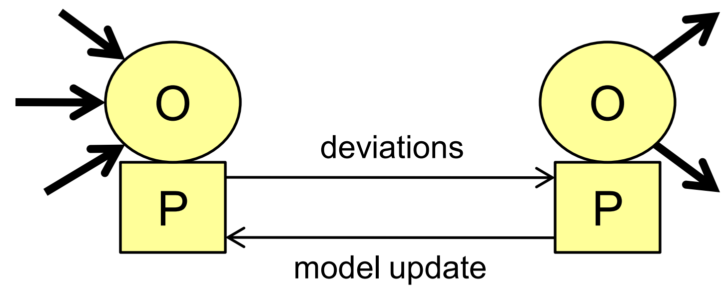

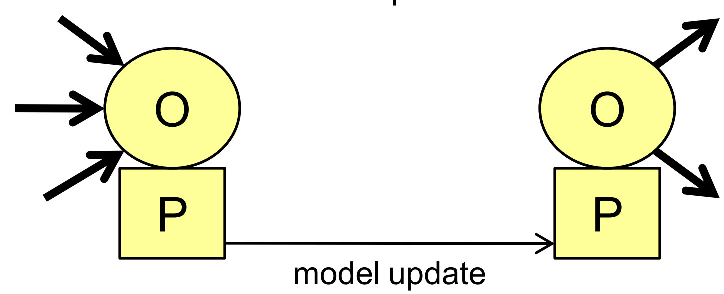

Spatial Predictors

- Basic approach in WSN: Query only sensors that most increase confidence in currently modeled world
  - Application specifies confidence threshold
- Exploit correlation between sensors
  - Different sensors (on the same node)
    - E.g. temperature readings and battery voltage
  - Proximate sensors of the same type
- Reduced number of actual sensor readings required
  - No need to wait for all nodes' results
  - Improved latency

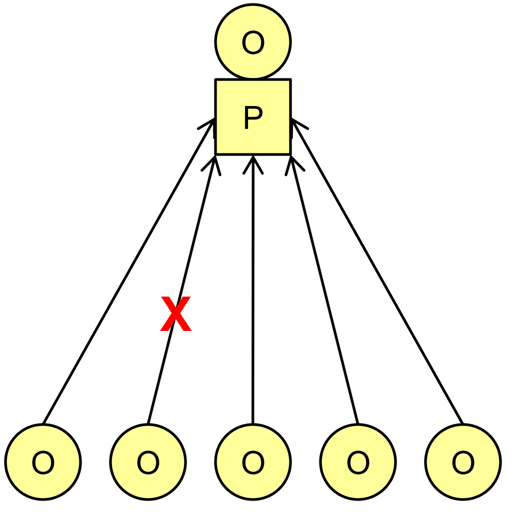

Spatio-Temporal Predictors

- Exploit knowledge about stream phenomena of air or water
  - E.g. stream of warm air leads to rising temperatures

- Employ MPEG motion pattern predictors
  - Treat sensors as pixels of an image
  - Value patterns represent video
  - Derive value predictor for single pixel and push it to the corresponding sensor node

Location Predictors

- Main use in Object Tracking Sensor Networks
- Dual predictor moves with the object
- Usage of behavior profiles / movement patterns to improve predictions

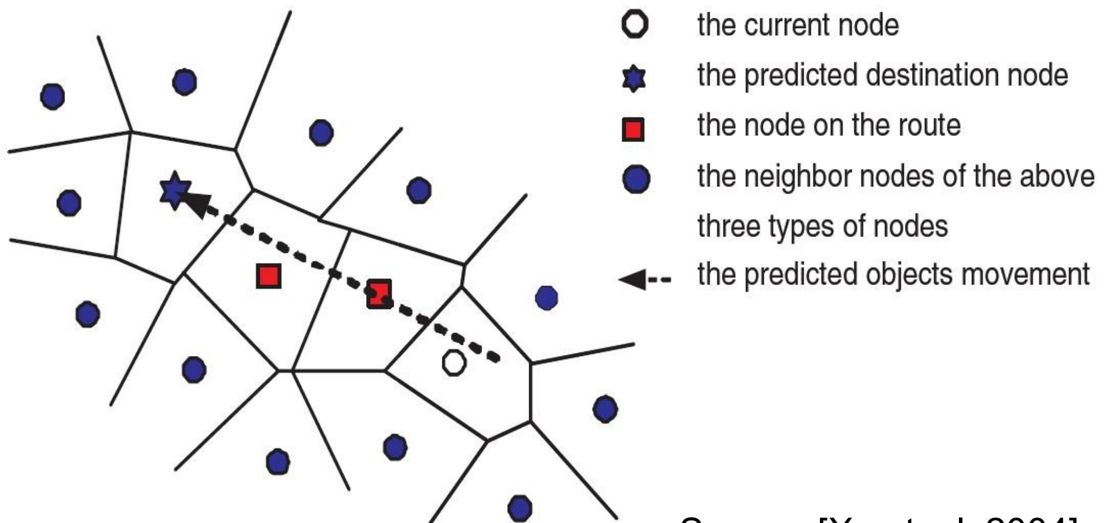

Source: [Xu et. al. 2004]

Outline

- Motivation
- System Description
- Predictors
- Challenges
  - Placement
  - Distributed Stream Processing
  - Application Interaction
- Summary

Placement

- Predictors close to the application
  - Improved latency due to direct connection
  - Data of all relevant sensors is available

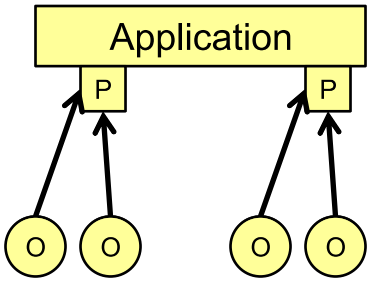

- Predictors close to the sensors
  - Bandwidth reduction with temporal predictors
  - Data required by aggregating operators multi-hop predictor

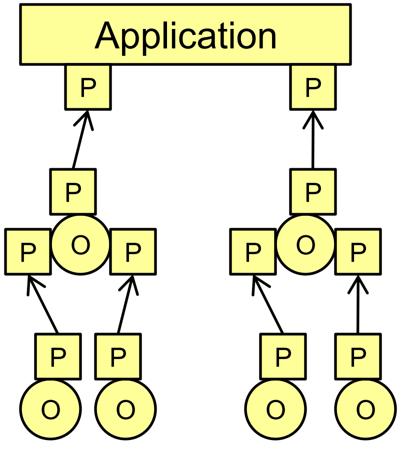

- Transition from / to WSN via gateways
  - Heterogeneous architecture and optimizations inside WSN

Distributed Stream Processing

- Window operators
  - Materialize stream for join and aggregate operations
    - Join temperature and humidity measurements to single stream
    - Average temperature over last 10 measurements or last 5 minutes
  - Semantics of tuple-based window operators changes with predictors since number of tuples is not preserved
  - Conversion to time-based windows required

- Discrete events
  - Continuous occurrence rate can be predicted but no common models for the time of the next event

Application Interaction

- Applications have different quality of data / service requirements
  - Transparent re-usage of data streams requires adjustment
  - Tradeoff precision vs. latency
- Predictors can be used for intelligent load shedding
  - Link overload: Coarser update threshold can reduce data
  - Node overload: Switch to simpler predictor model
- Supplying application knowledge also requires unified interface
  - Movement patterns for objects
  - Correlations between sensors
  - Expected and historical measurement patterns

Conclusion & Future Work

- Conclusions
  - Predictors can also be used in DSPS
    - In-time approximate query results
    - Reduction of data rates for load shedding
  - Predictors need a transparent unified interface
    - Application on arbitrary data links
    - Combination with arbitrary operators
- Future work
  - Additional application knowledge required to improve predictors
    - QoD/QoS restrictions, movement patterns, correlations, …
  - Find and adapt existing data streams for re-usage

Questions?

Thank you for your attention!

Contact:
Andreas Benzing
Institute of Parallel and Distributed Systems (IPVS)
Universität Stuttgart
Universitätsstraße 38
70569 Stuttgart, Germany

Email: andreas.benzing@ipvs.uni-stuttgart.de
Web: http://www.ipvs.uni-stuttgart.de
Phone: (+49) 711 – 7816 280

References

- [BSB07] Borgne, Y. L.; Santini, S. & Bontempi, G.
  Adaptive model selection for time series prediction in wireless sensor networks
  Signal Processing, **2007**, *87*, 3010 - 3020
- [DGL+05] Desnoyers, P.; Ganesan, D.; Li, H.; Li, M. & Shenoy, P.
  PRESTO: A Predictive Storage Architecture for Sensor Networks
  HOTOS'05: Proceedings of the 10th conference on Hot Topics in Operating Systems, USENIX Association, **2005**, 23-23
- [DGM+04] Deshpande, A.; Guestrin, C.; Madden, S. R.; Hellerstein, J. M. & Hong, W.
  Model-Driven Data Acquisition in Sensor Networks
  VLDB '04: Proceedings of the Thirtieth international conference on Very large data bases, VLDB Endowment, **2004**, 588-599
- [GI01] Goel, S. & Imielinski, T.
  Prediction-based Monitoring in Sensor Networks: Taking Lessons from MPEG
  SIGCOMM Comput. Commun. Rev., ACM, **2001**, *31*, 82-98
- [GQP06] Garcia, O.; Quintero, A. & Pierre, S.
  Profile-based energy minimisation strategy for Object Tracking Wireless Sensor Networks
  WiMob'2006: IEEE International Conference on Wireless and Mobile Computing, Networking and Communications, **2006**, 372-379

References

- [LCLC04] Lam, K.; Cheng, R.; Liang, B. & Chau, J.
  Sensor Node Selection for Execution of Continuous Probabilistic Queries in Wireless Sensor Networks
  VSSN '04: Proceedings of the ACM 2nd international workshop on Video surveillance & sensor networks, ACM, **2004**, 63-71
- [SR06] Santini, S. & Römer, K.
  An Adaptive Strategy for Quality-Based Data Reduction in Wireless Sensor Networks
  Proceedings of the 3rd International Conference on Networked Sensing Systems (INSS'06), **2006**, 29-36
- [TLH08] Tseng, V.; Lin, K. & Hsieh, M.
  Energy Efficient Object Tracking in Sensor Networks by Mining Temporal Moving Patterns
  IEEE International Conference on Sensor Networks, Ubiquitous and Trustworthy Computing. SUTC '08, **2008**, 170-176
- [XWL04] Xu, Y.; Winter, J. & Lee, W.
  Dual Prediction-based Reporting for Object Tracking Sensor Networks
  MOBIQUITOUS 2004: The First Annual International Conference on Mobile and Ubiquitous Systems: Networking and Services, **2004**, 154-163