



Software Defined Buildings

Efficient Analytics with BTrDB: From Grids to Buildings



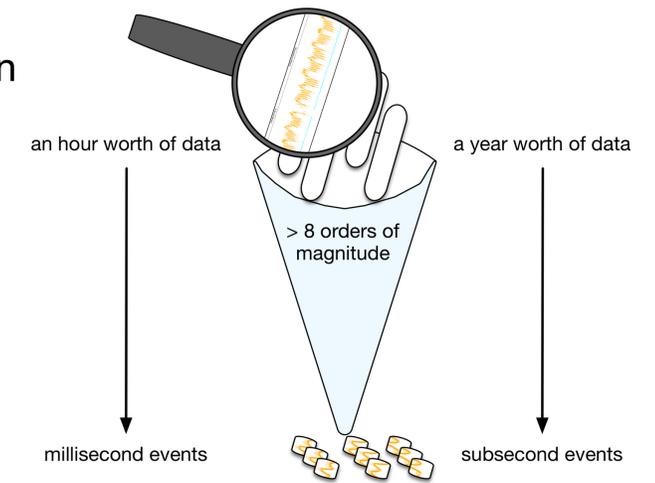
Omid Ardakanian, Michael Andersen, David Culler

Problem Definition

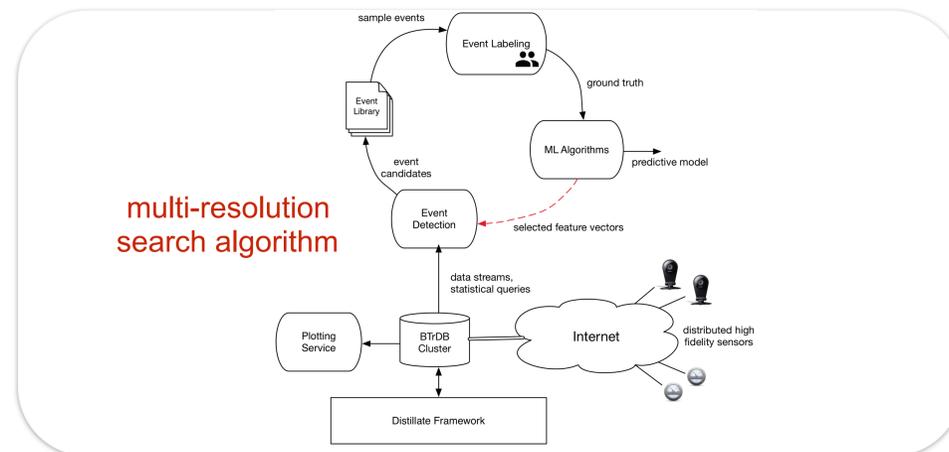
- High-precision high-sample-rate data collected over extensive time periods by an enormous number of high fidelity sensors installed in various cyber physical systems
 - electrical grids, commercial buildings, transportation systems, ...
- It is crucial to identify, analyze, and react to critical events in a timely manner
 - anomalies and critical events are rare
- Exploiting statistical queries of data offered by BTrDB at various temporal resolutions, we design an efficient search algorithm that is **logarithmic in the size of data** and **linear in the number of events**

Multi-Resolution Search Algorithm

- Query statistical summaries (max, min, mean, count, ...) at a temporal resolution
- Compare a function of these statistical summaries with a threshold
- Dive down if some condition is satisfied
- Query raw data only if a desired temporal resolution is reached
- Run compute intensive tasks on smaller chunks of data



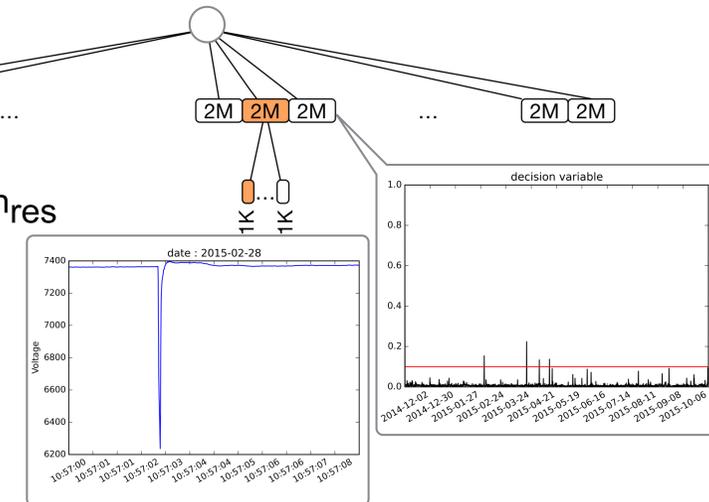
System Architecture



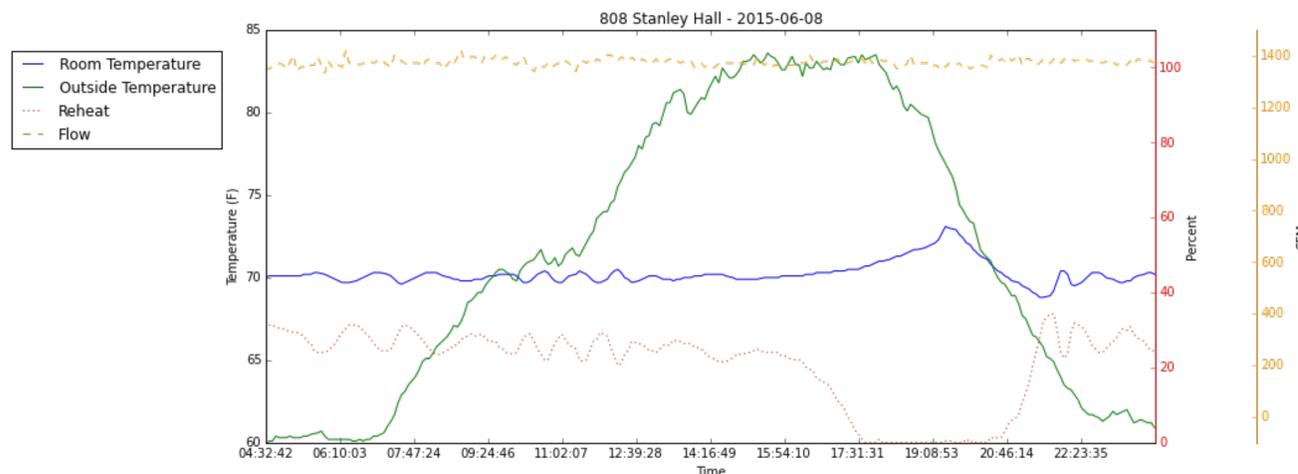
storage and processing infrastructure

Detecting Critical Events in Distribution Grids

voltage sag detector
kernel: $(\text{mean}_{\text{res}} - \text{min}_{\text{res}}) / \text{mean}_{\text{res}}$
statistical records: min, mean



Detecting Anomalies in Buildings



Conclusion and Future Work

- Complexity of the multi-resolution search algorithm is $O(n \log(L))$
 - quite efficient in finding rare events
- Building a detector can be challenging
 - how to define the threshold and the kernel function?
- Using a data-driven approach to develop sophisticated detectors

