SATware: Middleware for Sentient Spaces

WMSC 2007

Bijit Hore, Hojjat Jafarpour, Ramesh Jain, Shengyue Ji, Daniel Massaguer, Sharad Mehrotra, Nalini Venkatasubramanian, and Utz Westermann

Traditional sensor network

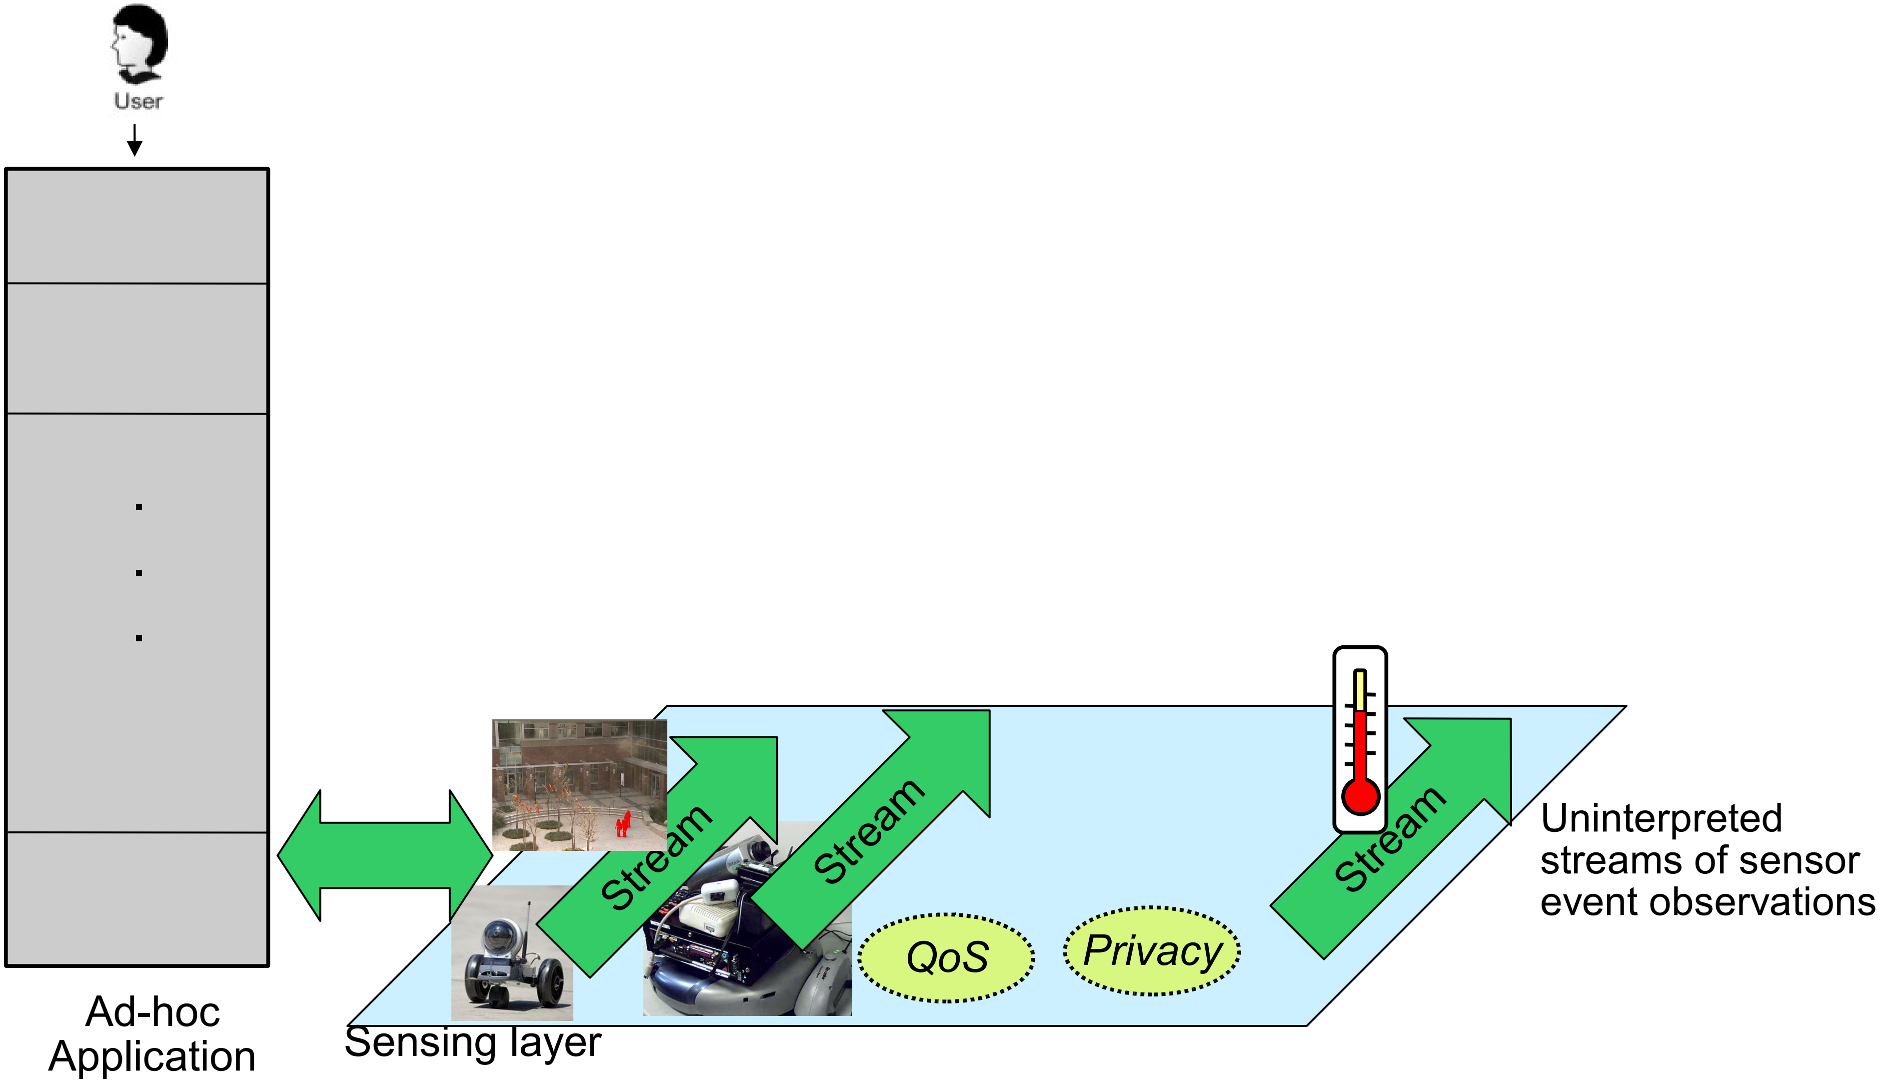

SATware: Semantic Layer

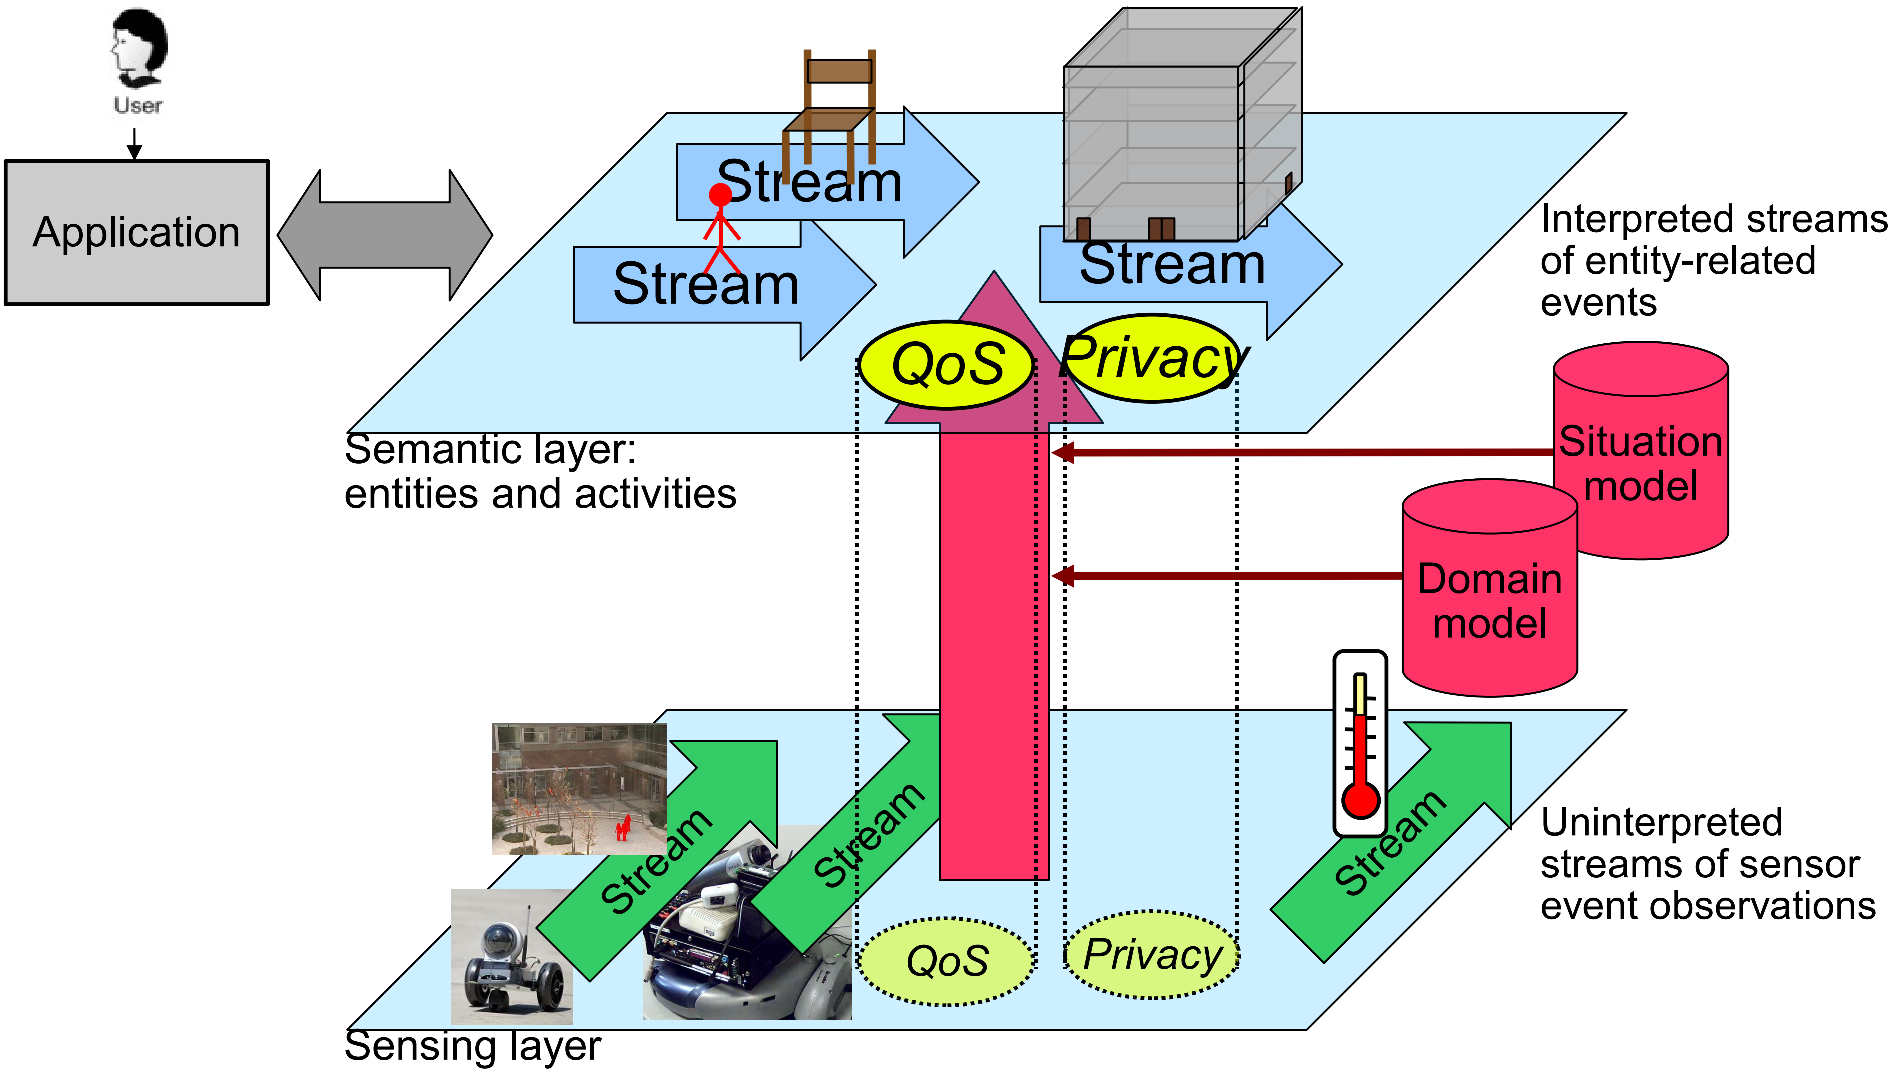

SATware: Challenges

In emergency response, being able to specify quality of the service and having the system guarantee it is crucial

Different levels of abstraction from user's high abstraction level to sensors' level
Users care about entities and events

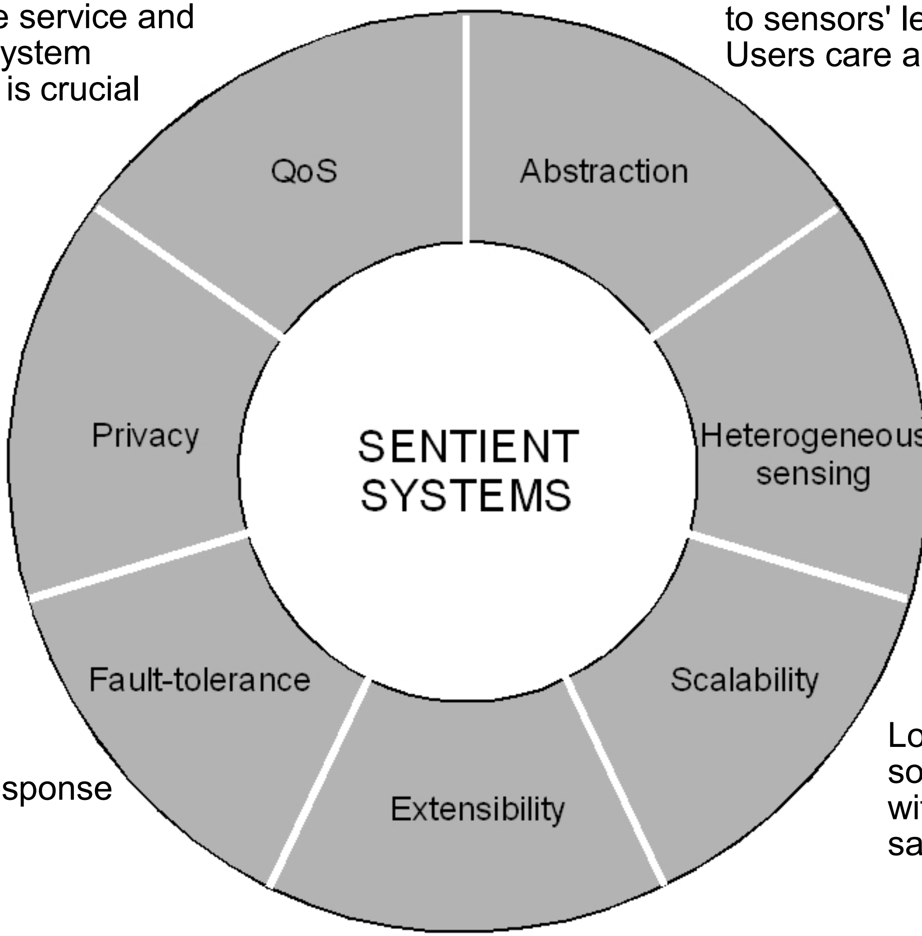

Privacy of users (both direct and indirect users) has to be protected

Different kinds of sensors: different resources, mobility capabilities, data streams produced (scalar, video), phenomena being sensed

In emergency response failures happen

Lots of sensors
some sensors produce streams with very high sample rates and sample sizes (e.g. video).

Sensors and services need to be added, removed and discovered

Queries in SATware

- Queries written at the **entity** level
  - E.g., *Identify who left the coffee pot burning?*
  - E.g., *Notify when someone leaves a mess in the shared kitchen.*
  - E.g., *Notify when Sharad is in the shared area and his status is "Not Busy".*
  - E.g., *Notify when Iosif and Sharad are both in the shared area*

Queries are translated into Operators

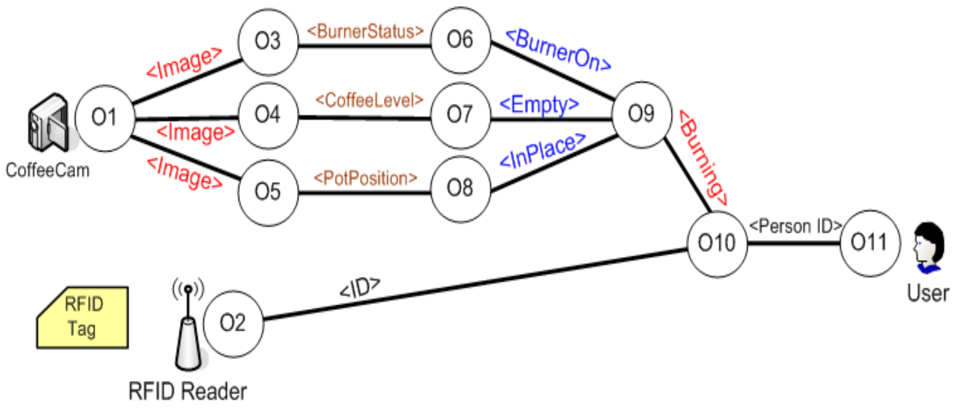

Stream processing model -Data model-

A stream is an infinite discrete flow of packets:

stream = list(packet)

packet = <*t* , *c*>

t : time stamp

c: content

- C ϵ domain(DT)
- DT := int | float | char | byte | void | bool | event
- DT := list(DT) | tuple(DT)
- Event := <event_id, confidence, DT>

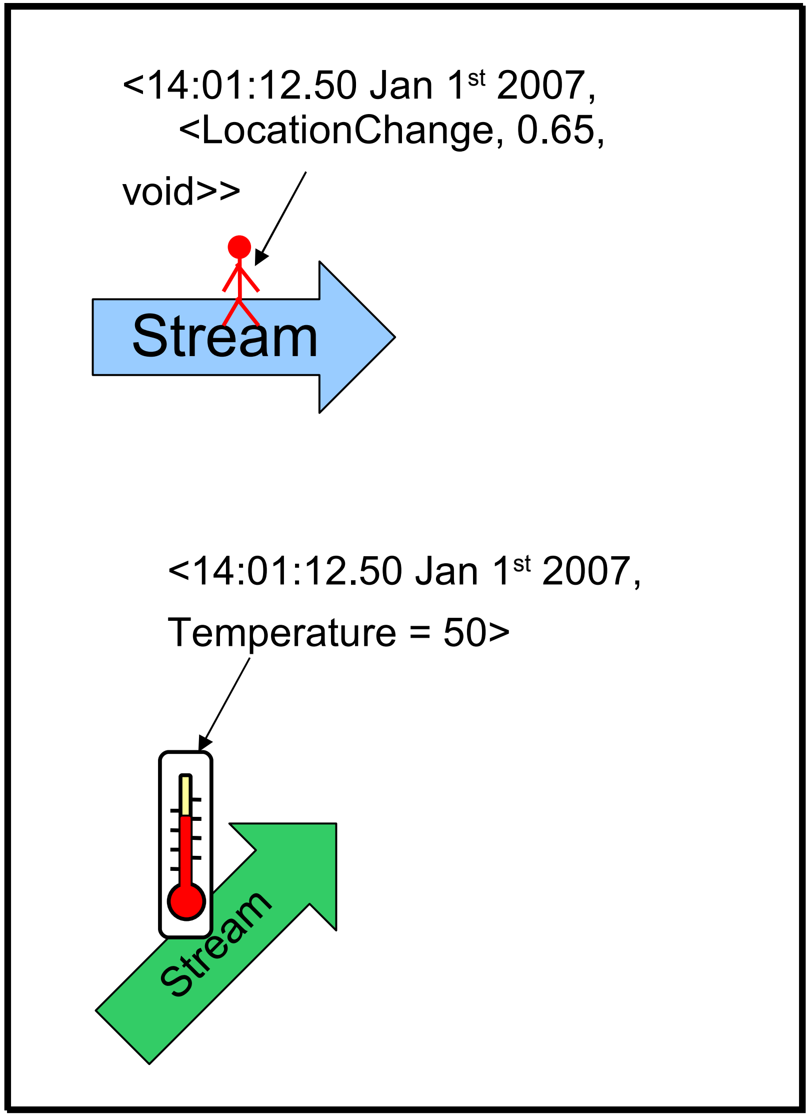

Stream processing model -Processing model-

Processing based on operators

$OP = \{ f_{DTx \ldots xDT} \mid f_{DTx \ldots xDT}: \text{streamx} \ldots \text{xstream} \rightarrow \text{stream} \}$

- Code reuse
- Sharing operator instances
- Simplify design
- Load distribution

- Virtual sensors

$WhoBurnedCoffee_{x,y} = O10(CoffeeBurn_x, PersonID_y)$

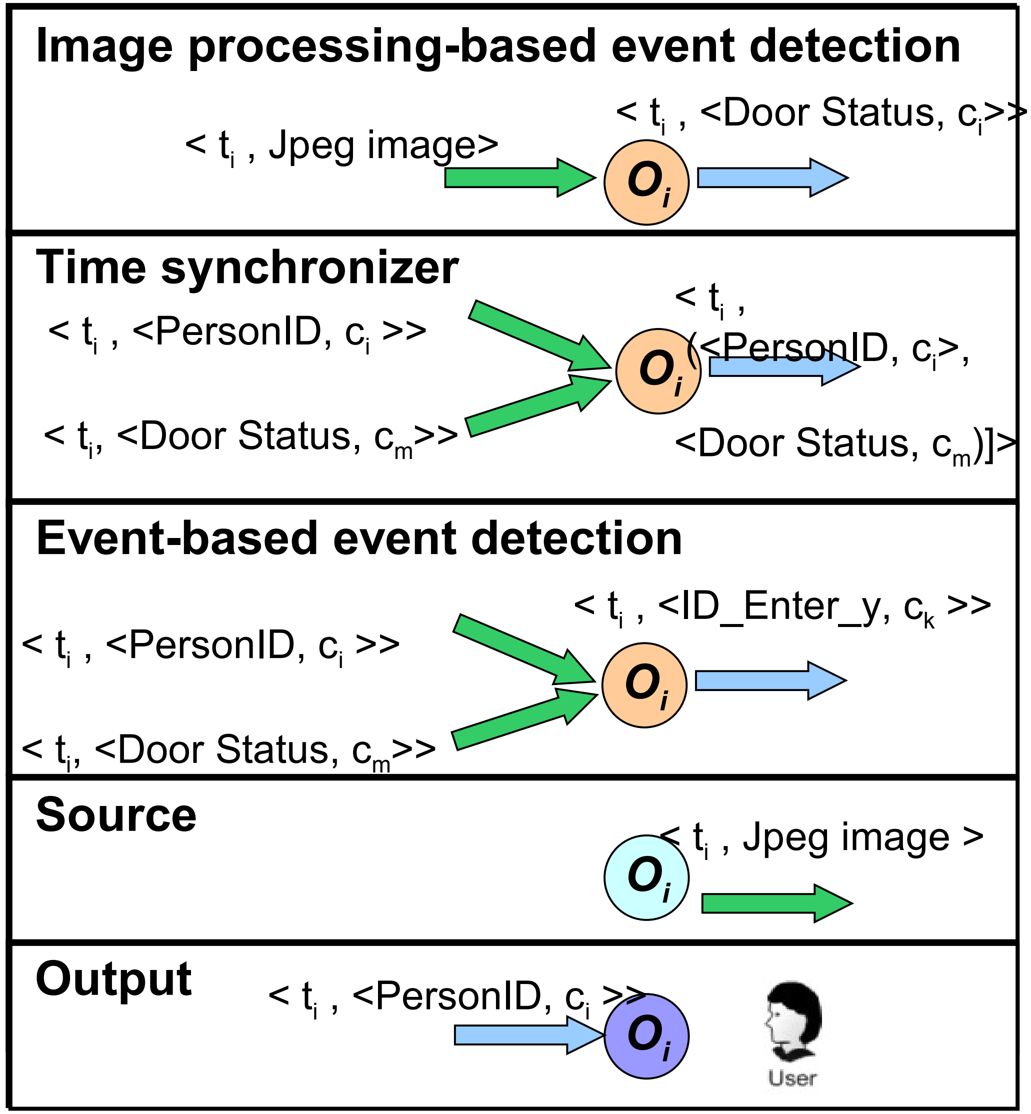

Sample application

Remote Windshield Assessment

UCI has been instrumented with a large set of sensors throughout all campus.

The instrumentation would allow conducting a **quick assessment** of the status of buildings from any device that has access to Internet by accessing an online *Remote Windshield Assessment* application.

The system would allow any user (for example, UCIPD) to **register for events**.

Sample query

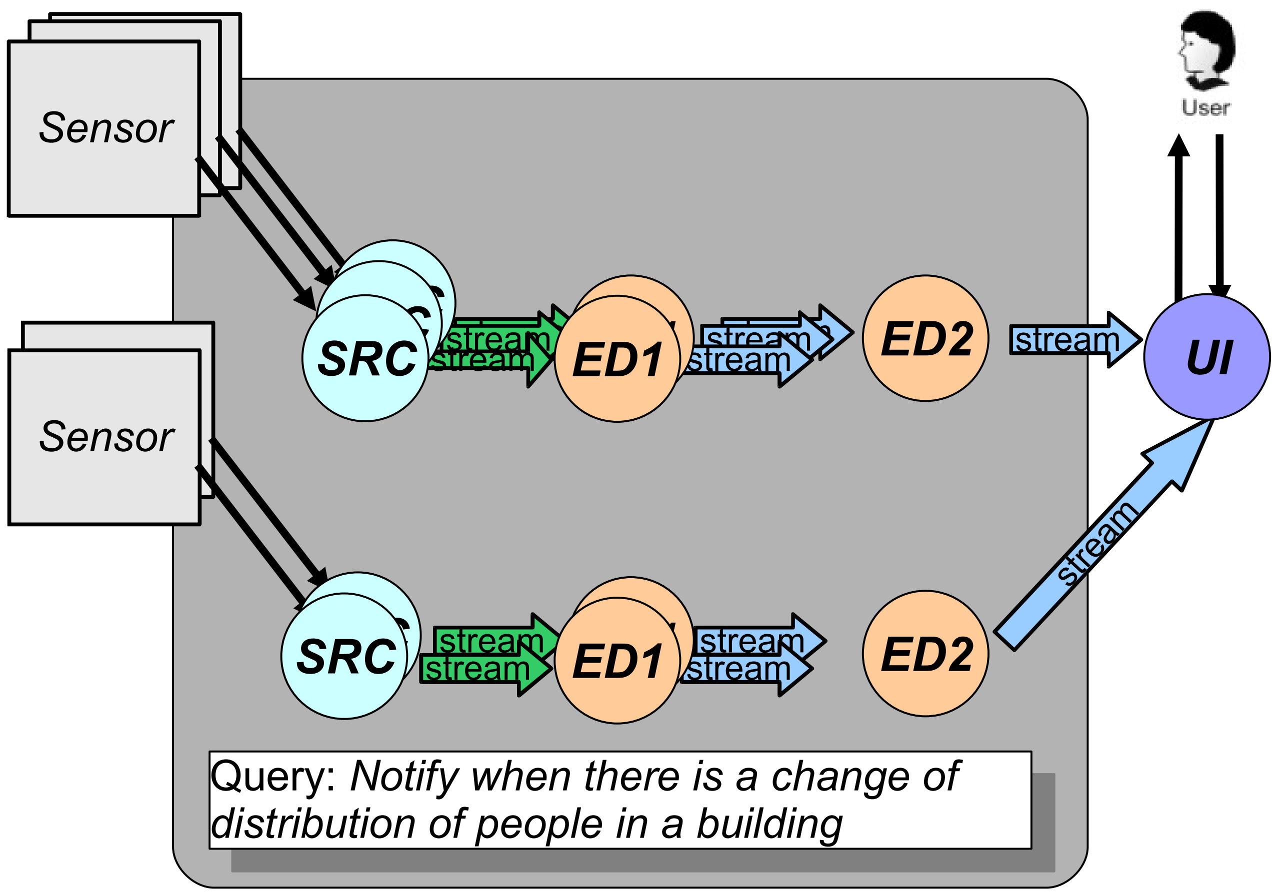

- **ED1** — **Event detection1:** Estimation of number of people in and out a certain door.
- **ED2** — **Event detection 2:** Estimation of number of people in a room.
- **SRC** — **Source agent:** Gets JPEG frames from a network camera.
- **UI** — **Source agent:** User interface.

SATware architecture

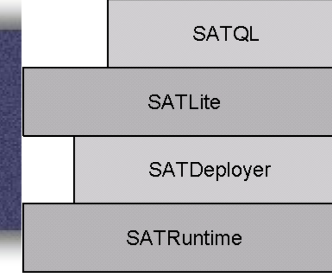

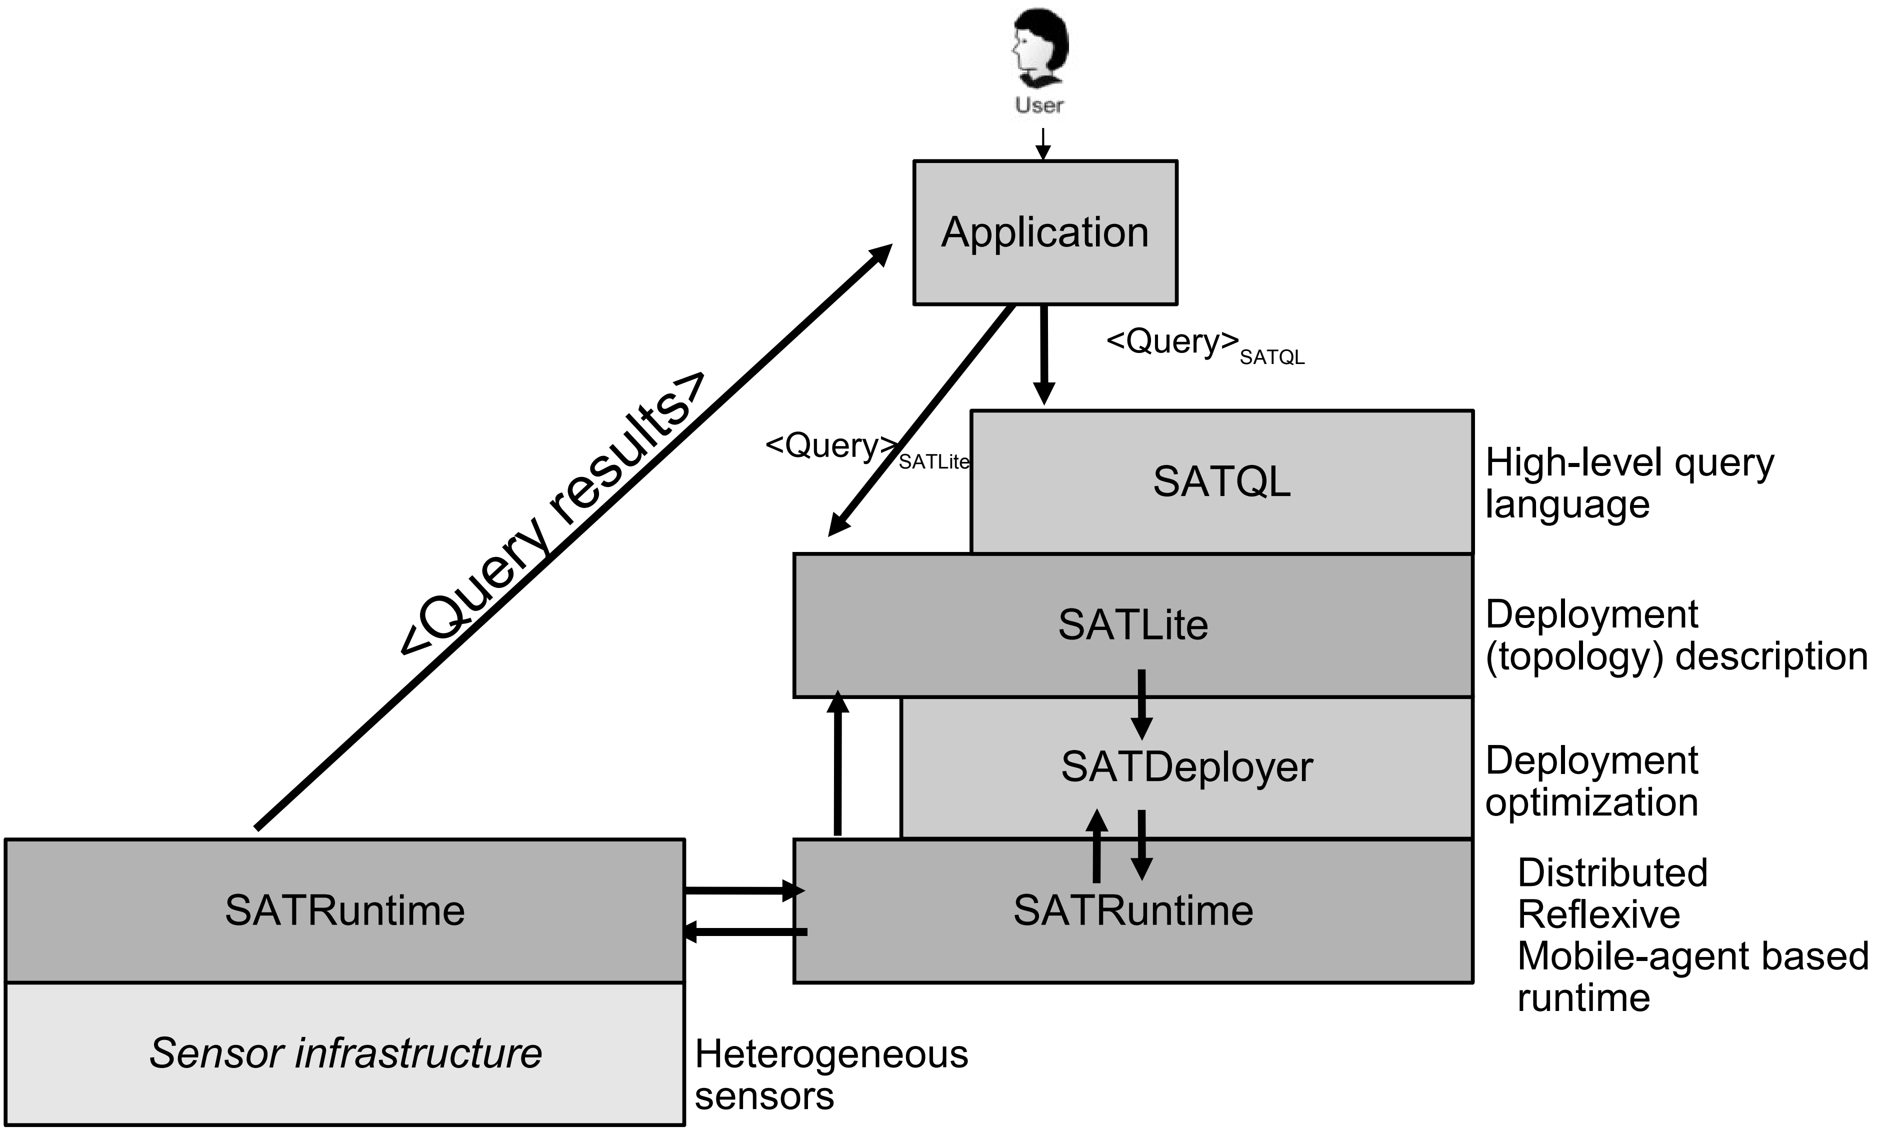

SATRuntime

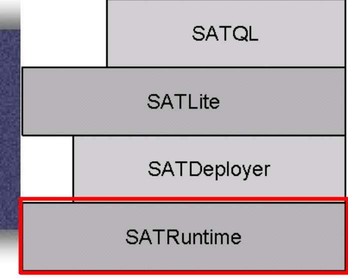

SATQL
SATLite
SATDeployer
SATRuntime

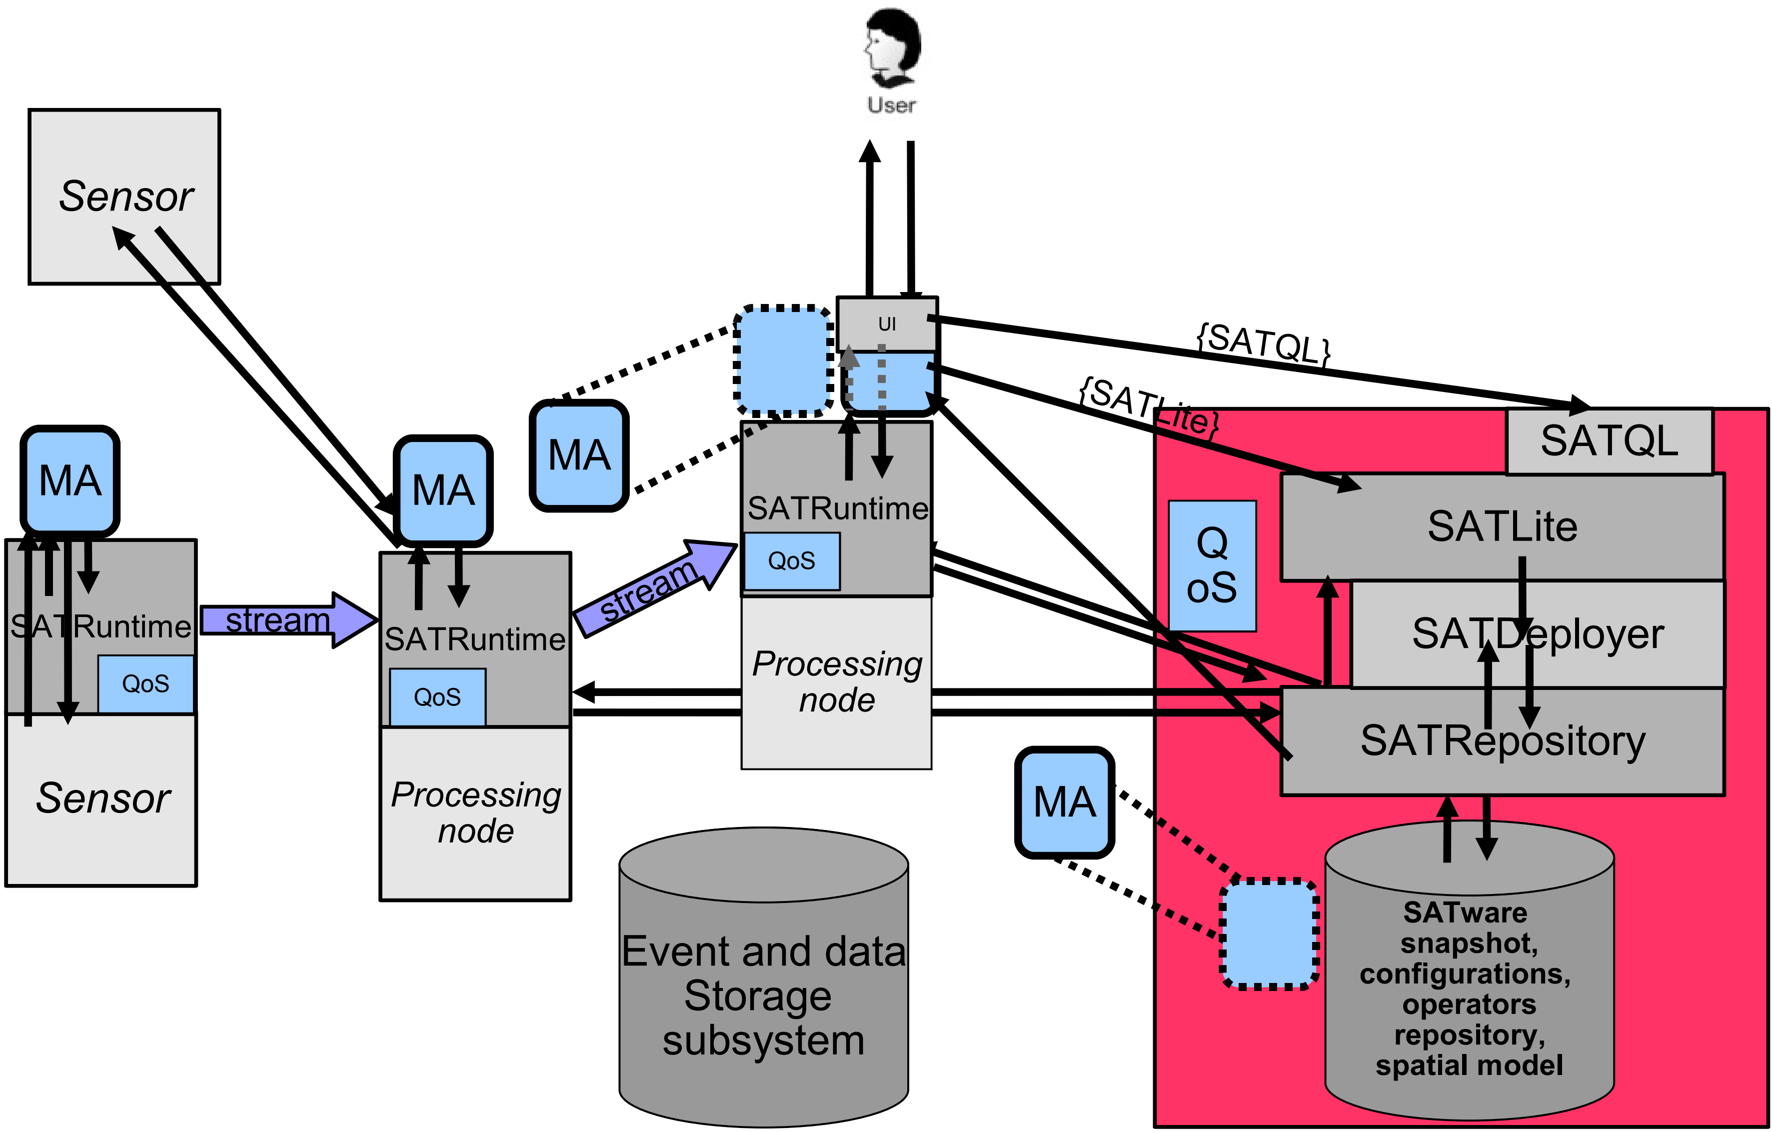

User
Sensor
UI
{SATQL}
{SATLite}
MA
SATRuntime
QoS
Sensor
stream
MA
SATRuntime
QoS
Processing node
stream
MA
SATRuntime
QoS
Processing node
SATQL
SATLite
Q oS
SATDeployer
SATRepository
MA
Event and data Storage subsystem
SATware snapshot, configurations, operators repository, spatial model

SATRuntime

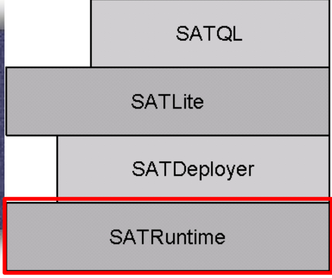

- Runtimes are connected instead of the agents.
- Agents are programmed in a topology-agnostic manner.
- SATRuntime provides mobility support and message-passing.

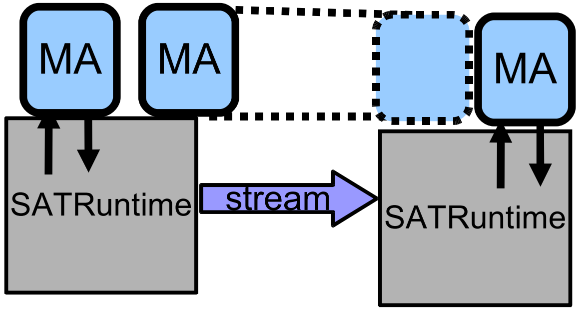

- SATRepository contains
  - repository of operators and
  - state of SATware (sensors, runtime nodes, network topology, and current operator deployment).

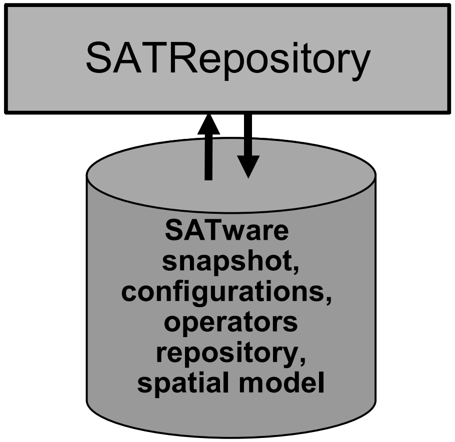

Sample query

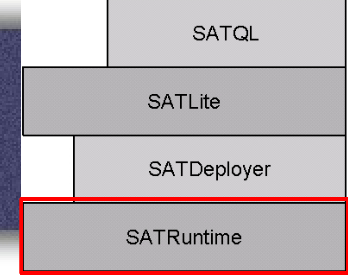

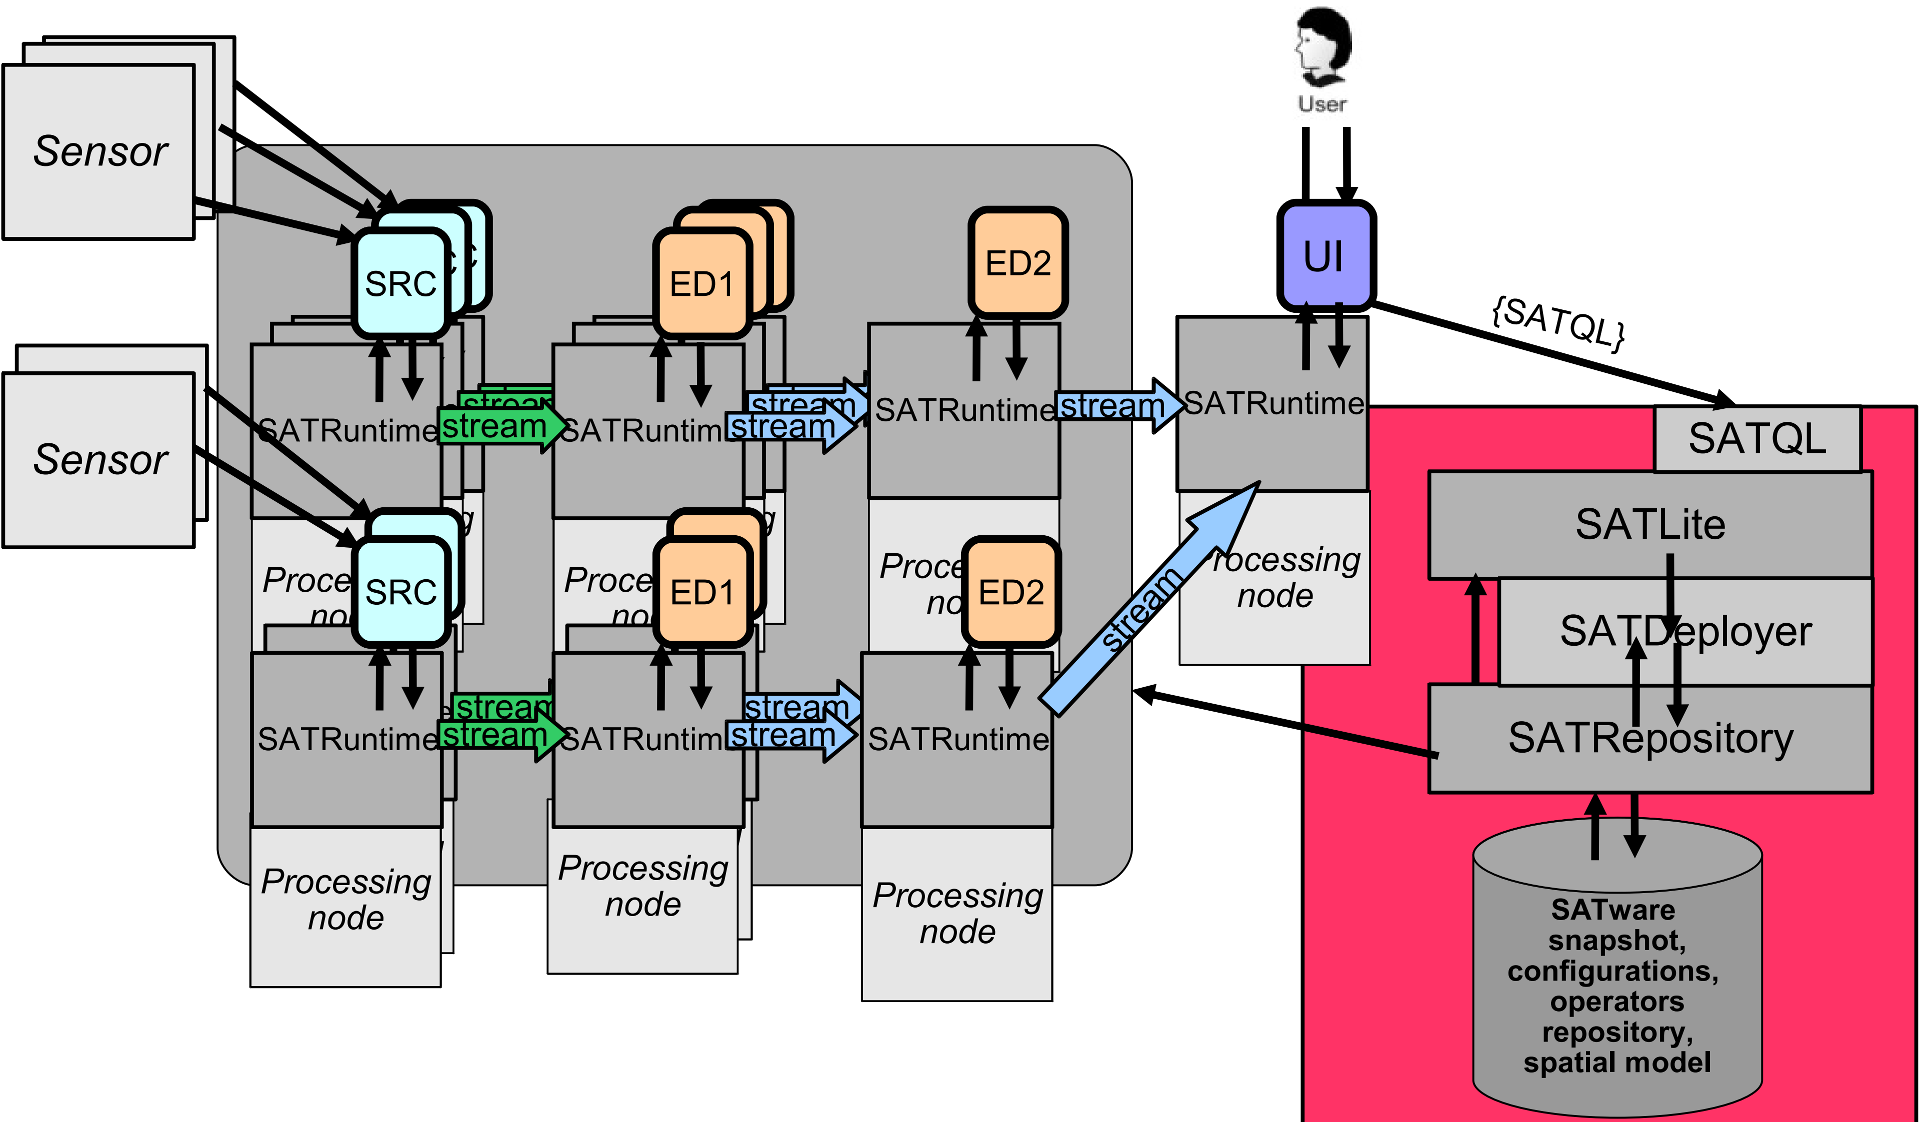

SATLite

SATLite is a language for describing queries

```
// operators
O1 = CreateOperator(ReadCam, camera_url);
O3 = CreateOperator(DetectBurner);

// data flows
Image = O1();
BurnerStatus = O3(Image);
BurnerOn = O6(BurnerStatus);
O11(PersonId);
```

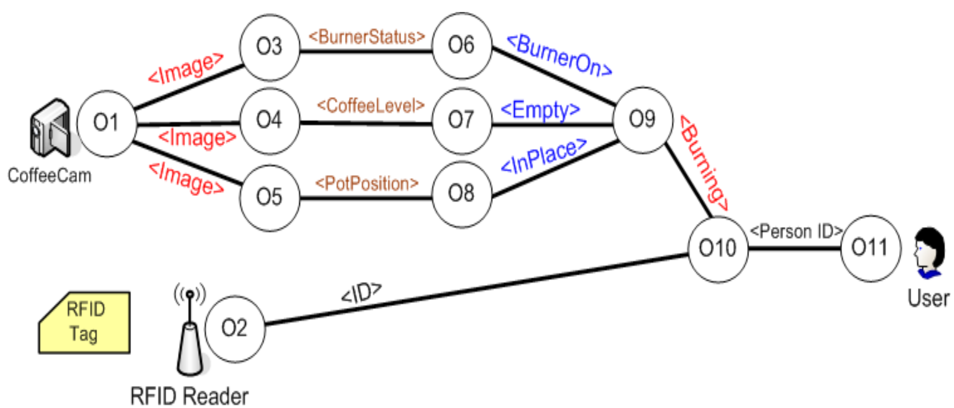

Predicate optimization

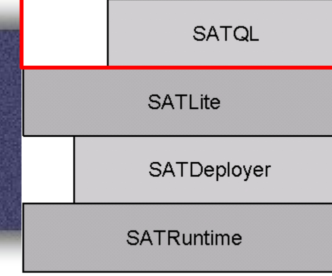

"Which image contains Iosif?"

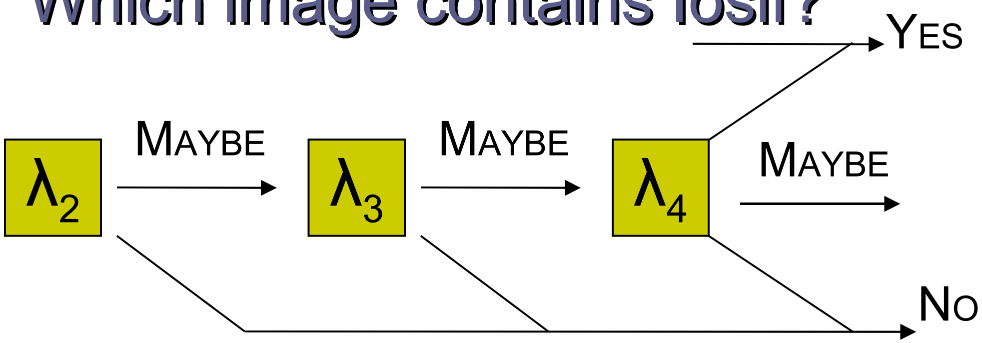

Assume RFID also available

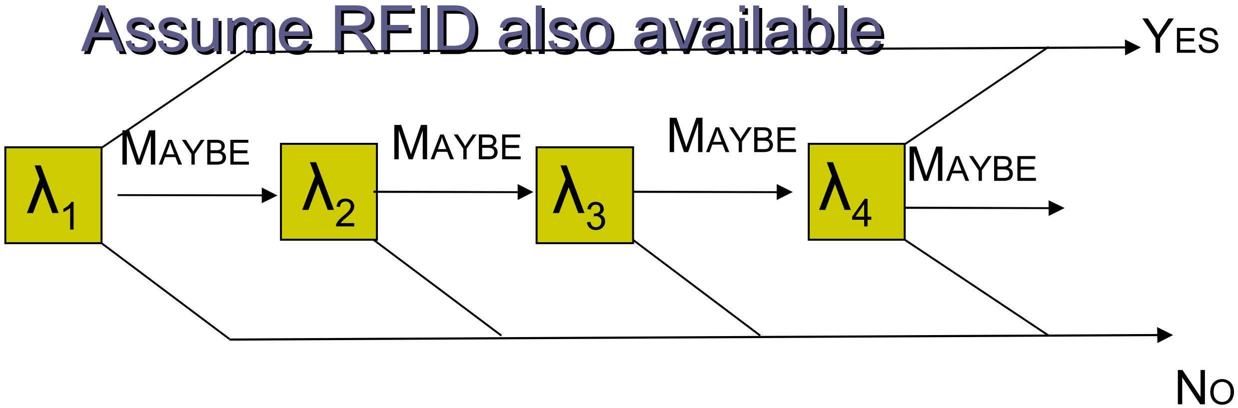

- λ_1 RFID tag matches Iosif's tag?"
- λ_2 Which pictures contain an object?"
- λ_3 Which pictures contain a person?"
- λ_4 Which pictures contain Iosif?"

[1]Iosif Lazaridis and Sharad Mehrotra, "Optimazation of Multi-Version Expensive Predicates", SIGMOD 2007.

Problem formulation

- Given
  - Sequence of selection predicates
    - $\lambda_1, \lambda_2, \ldots, \lambda_n$
  - With Per-tuple cost:
    - $c_1, c_2, \ldots, c_n$
  - And MAYBE selectivity: #Maybe / #All
    - $m_1, m_2, \ldots, m_n$
  - Such that:
    - $c_1 < c_2 < \ldots < c_n$
    - $m_1 > m_2 > \ldots > m_n$
- Determine the least cost execution plan for λ_n

$\lambda_k(t) =$ Yes / No / Maybe

[1]Iosif Lazaridis and Sharad Mehrotra, "Optimazation of Multi-Version Expensive Predicates", SIGMOD 2007.

Solutions

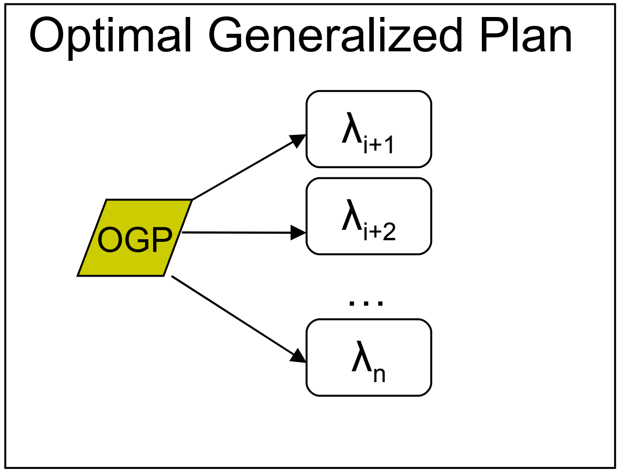

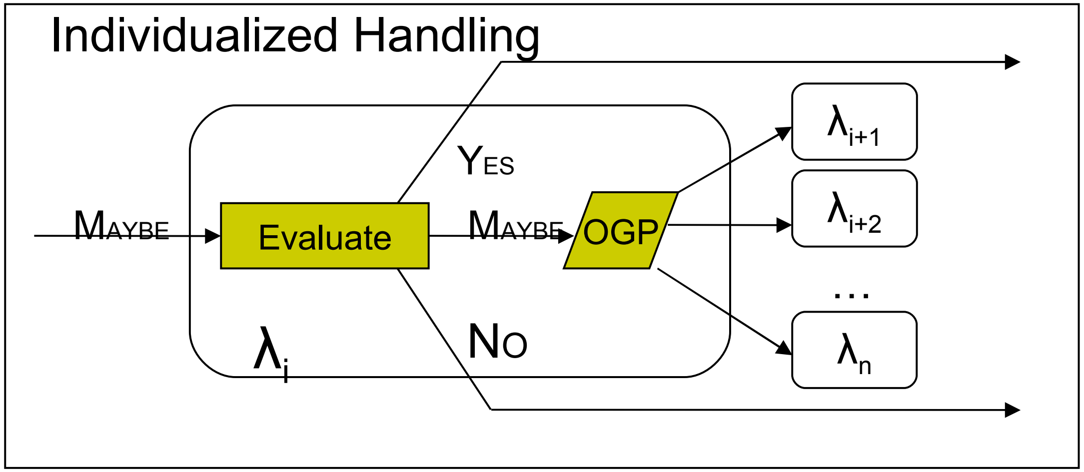

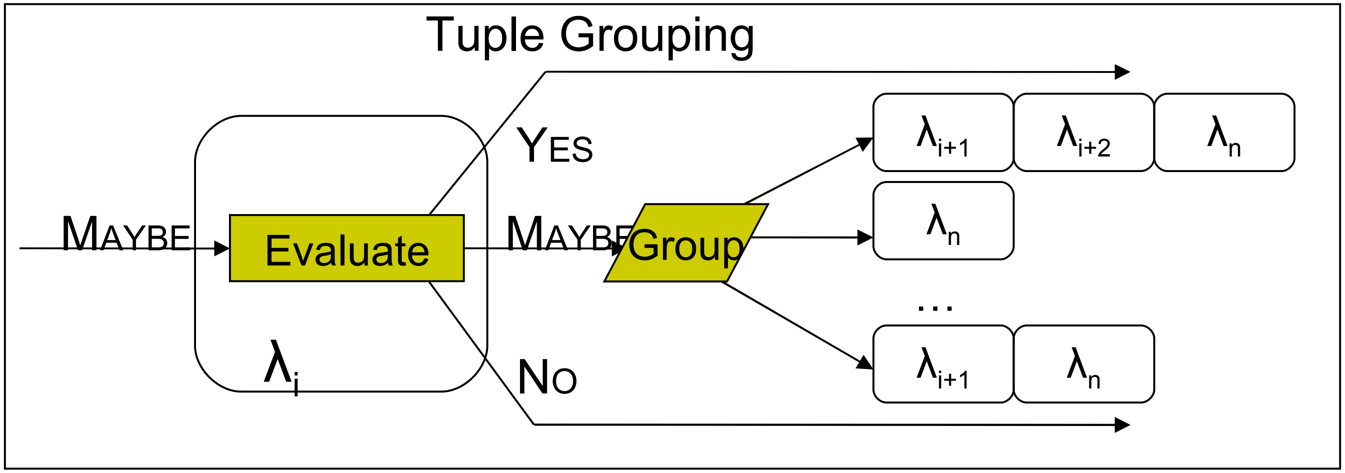

[1] **Iosif Lazaridis and Sharad Mehrotra, "Optimazation of Multi-Version Expensive Predicates", SIGMOD 2007.**

SATDeployer

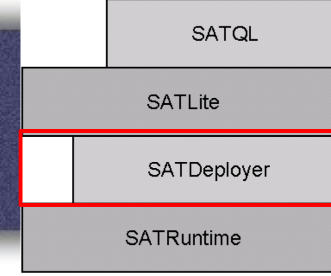

- Translating a query plan into a deployment plan
  - Mapping query plan nodes into processing nodes
- Minimizing query processing cost
  - Network cost
  - Computing cost
  - Query evaluation time
- Optimizations
  - Operator reuse
  - Load balancing

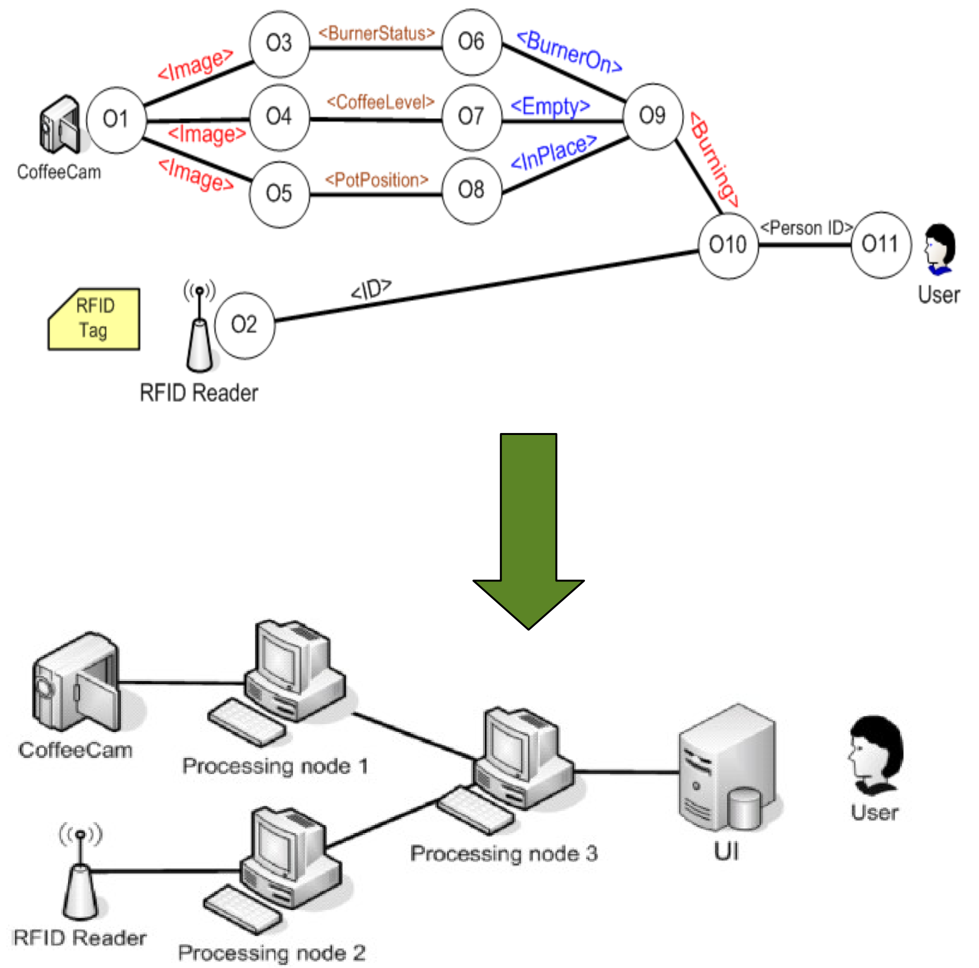

Privacy

Privacy is protected unless the user violates a specific rule

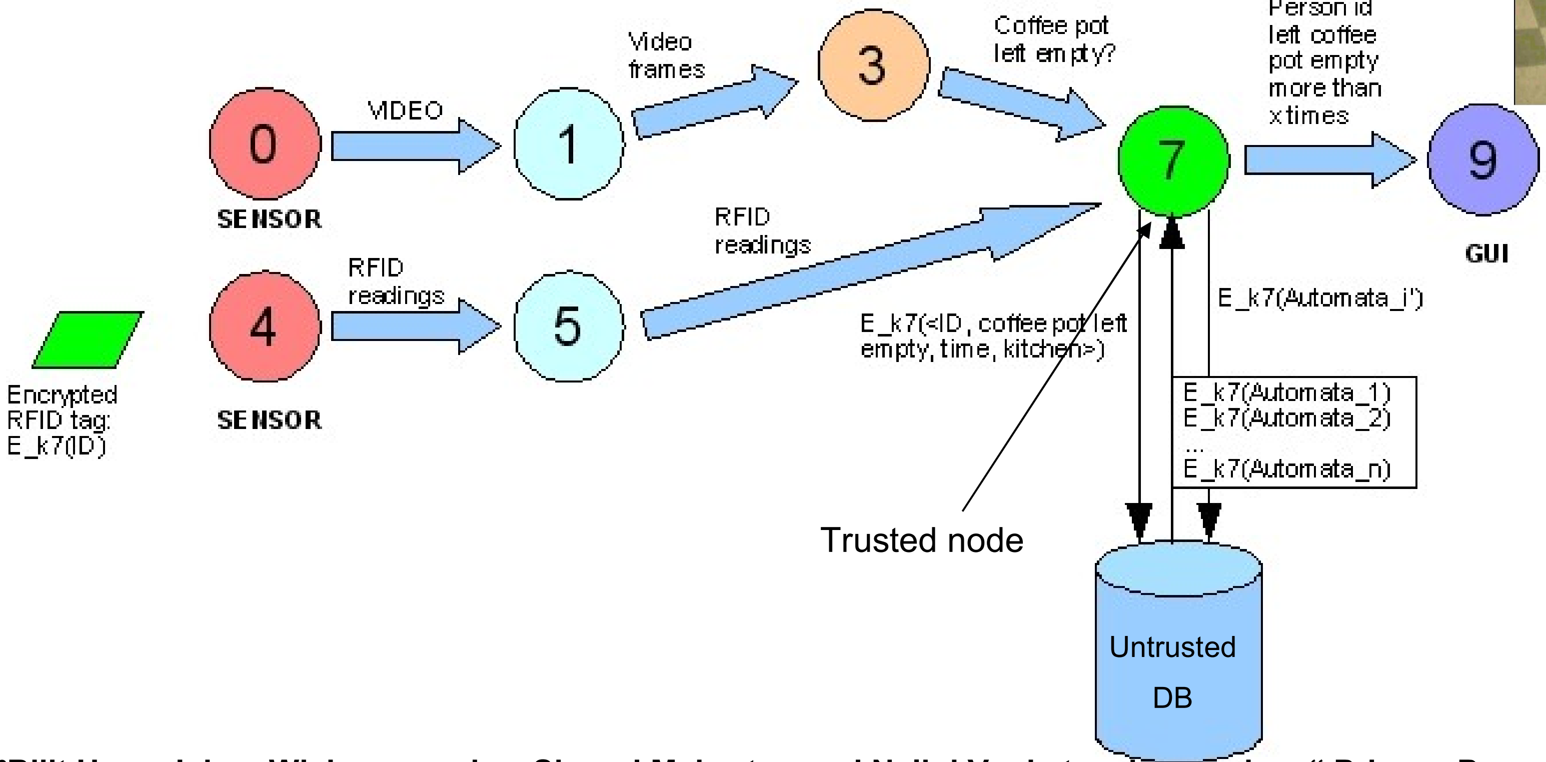

[2]Bijit Hore, Jehan Wickramasuriya, Sharad Mehrotra, and Nalini Venkatasubramanian, " Privacy-Preserving Event Detection for Pervasive Spaces" (2007, UCI-ICS tech-report)

Abnormal event detection

Aggregated behavior of human beings:

$N(t) = N_0(t) + N_E(t)$ ← *Markov Modulated Poisson*

Normal behavior: periodicity

$N_0(t) \sim P(N; \lambda(t))$ ← *Time-Varying Poisson*

$$\lambda(t) = \lambda_0 * \delta_{d(t)} * \eta_{d(t),h(t)}$$

$$1 \leq d(t) \leq 7 \text{ and } 1 \leq h(t) \leq 48/288/\ldots$$

Abnormal behavior: rare and short

$$N_E(t) \sim \begin{cases} 0 & z(t) = 0 \\ P(N; \gamma(t)) & z(t) = 1 \end{cases}$$

$z(t) \sim$ Markov Process

Unsupervised learning of model parameters via statistical methods (MCMC)

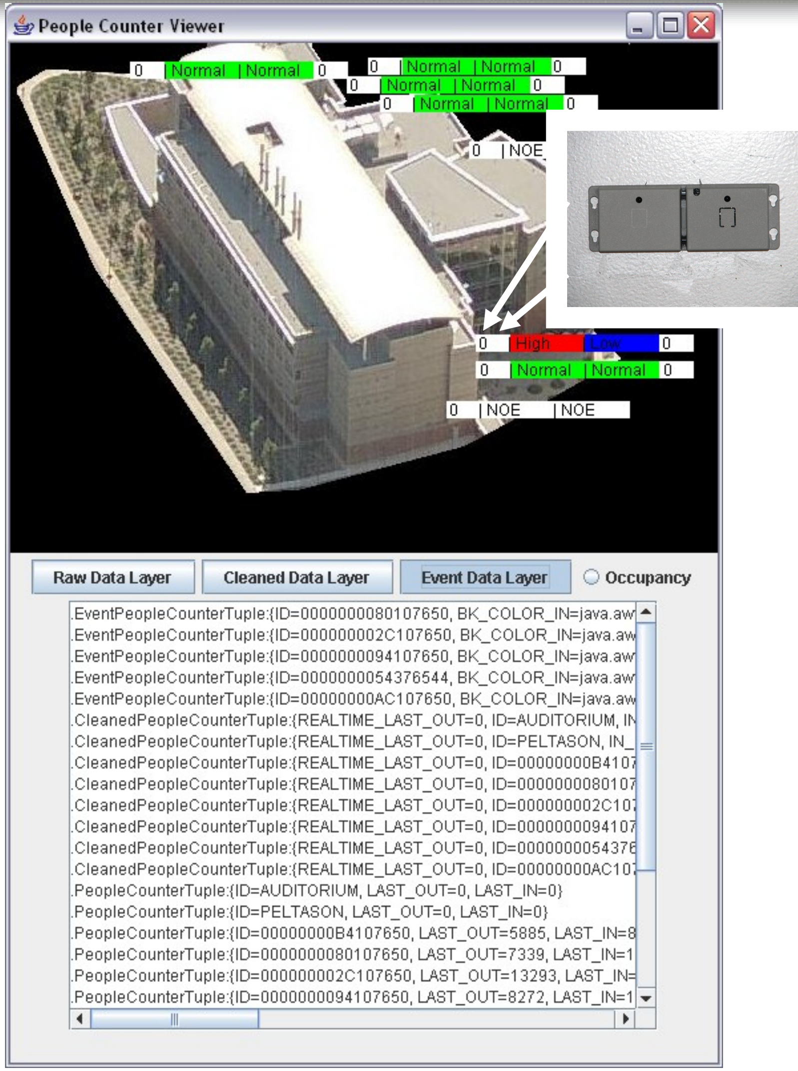

[3]Alexander Ihler, Jon Hutchins, and Padhraic Smyth, "Adaptive Event Detection with Time–Varying Poisson Processes", KDD 2006

Implementation

Responsphere:

1/3 UCI campus:
2 Indoor buildings
+
Outdoors

+200 sensors
+10 types

Summary

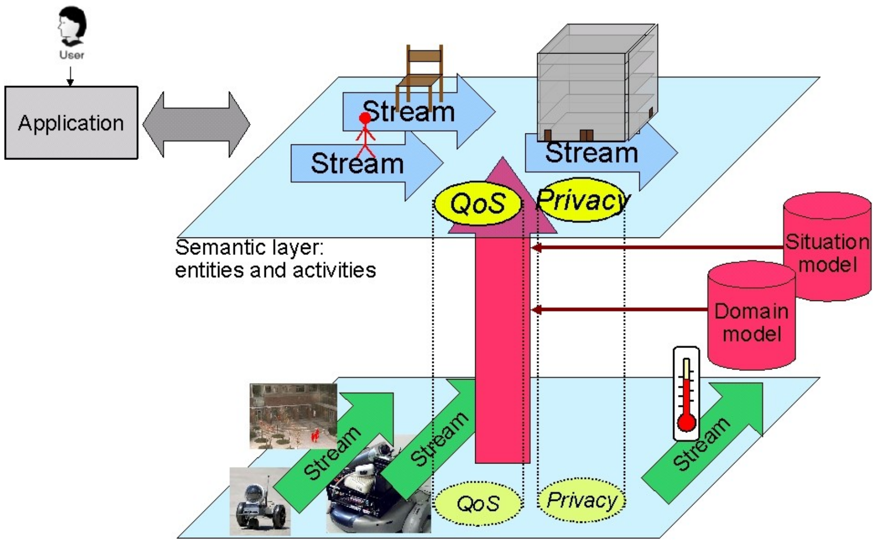

Sentient systems

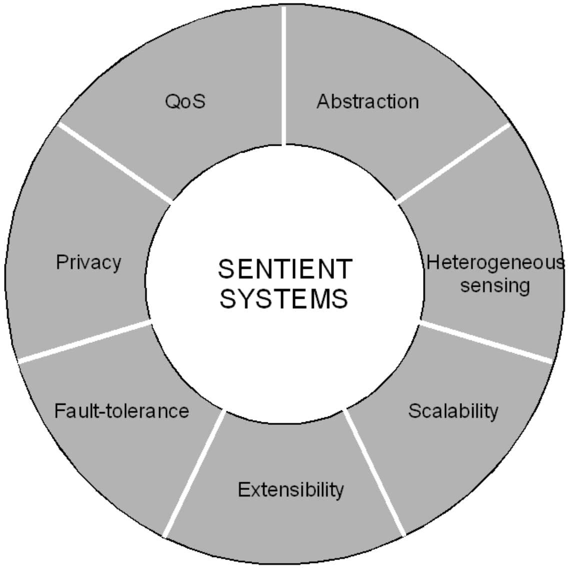

Challenges

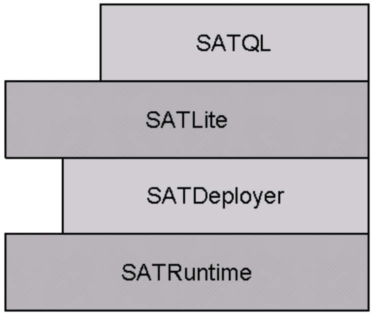

Layered architecture

Privacy applications

Thank you

http://satware.ics.uci.edu

Extra slides

Project RESCUE
June 20, 2007

http://www.itr-rescue.org

Funded through National Science Foundation Funding:
ITR award
NSF Research Infrastructure award

rescue Responding to Crises and Unexpected Events

Mission

The mission of RESCUE is to enhance the ability of emergency response organizations and the public to mitigate crises, save lives, and prevent secondary and indirect human and economic loss by radically transforming ways in which these organizations gather, process, manage, use and disseminate information during man-made and natural catastrophes.

Approach: Information Technologies for Improved Situational Awareness

Hypothesis: ***Right Information*** to the ***Right Person*** at the ***Right Time*** can result in dramatically better response

Research Team

Multidisciplinary Team
IT Experts, Scientists, Engineers

25 Investigators

50+ Students

6 Universities

1 Southern California Company

BYU Brigham Young University	• Privacy • Security • Trust
Colorado University of Colorado at Boulder — Natural Hazards Center	• Natural Hazards Center • Social Science
Illinois University of Illinois at Urbana-Champaign	• Data Management • Security and Trust
ImageCat, Inc.	• Disaster Analysis • Earthquake Engineering • GIS • Transporation Modeling • Urban Planning
UCIrvine University of California, Irvine	• Civil Engineering • Data Analysis & Mining • Data Management • Middleware & Distributed Systems • Privacy • Social Science • Transportation Science
University of Maryland	• Civil Engineering • Transportation Engineering
UCSD	• Computer Vision • Networking • Multimodal Speech • Wireless

Industrial Partners

5G Wireless *Broad-ranged IEEE 802.11 networking*	**AMD** *Compute Servers*
Apani Networks *Data security at layer 2*	**Asvaco** *1st responder (LAPD), and threat analysis software*
Boeing *Community Advisory Board Member*	**Canon** *Visualization equipment SDK*
Convera *Software partnership*	**Cox Communications** *Broadcast video delivery*
D-Link *Camera Equipment and SDK*	**Ether2** *Next-generation ethernet*
IBM *Smart Surveillance Software (S3) and 22 e330 xSeries servers*	**ImageCat, Inc.** GIS loss estimation in emergency response
Microsoft *Software*	**Printronix** *RFID Technology*
The School Broadcasting Company *School based dissemination*	**Vital Data Technology** *Software partnership*
Walker Wireless *People-counting technology*	

Government Partners

California Governor's Office of Emergency Services	California Governor's Office of Homeland Security
City of Champaign	City of Dana Point
City of Irvine	City of Los Angeles
City of Ontario Fire Department	City of San Diego
Department of Health and Human Services – Centers for Disease Control	Lawrence Livermore National Laboratory
Los Angeles County	National Science Foundation
Orange County	Orange County Fire Authority
U.S. Department of Homeland Security	

RESCUE Research

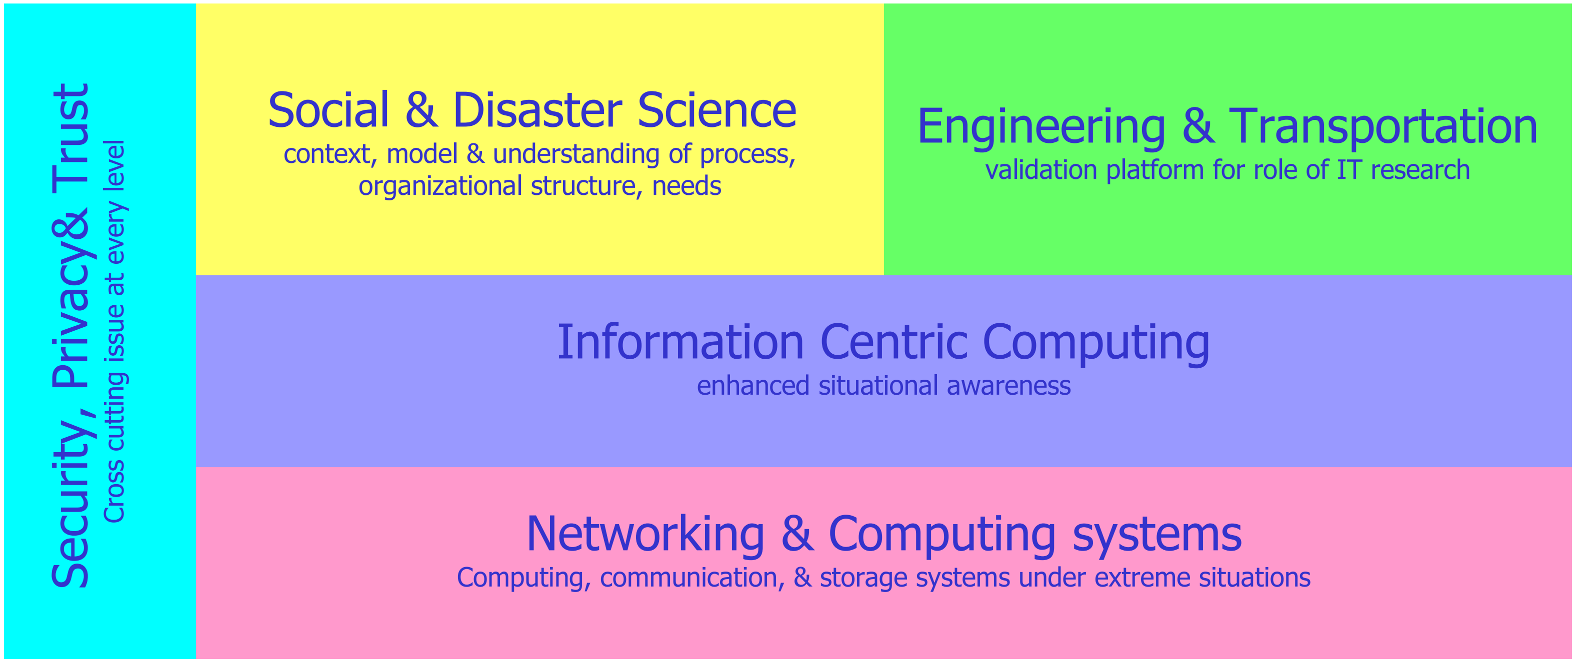

- Social Science
  - context and understanding of crisis domain
- Information Technology
  - infrastructure & tools to enhance flow of information & situational awareness
- Engineering
  - platform for realization, real-world physical constraints that help test, and validate

Project Structure