MiddleWare for Sensor Systems – keeping things Open

Ken Moody

Computer Laboratory, University of Cambridge

Overview of the talk

the Opera group and what we do - *past and present activities*

why we're at this Workshop – *TIME-EACM and* ***CareGrid***

fundamental problems of (sensor) data interpretation

volume of data, how and where to cope with it, audit and reliability

timeliness, feedback and stability – ***too hard*** for *TIME-EACM*

presenting relevant information as the context changes

OPERA Group

part of the *Systems Research Group* (SRG) in the Computer Laboratory

- ***Jean Bacon***, ***Ken Moody*** (*Faculty*)
- *about 6 RAs* , *and 12 PhD students* at present

MiddleWare for wide-area application support

- ***event-based systems*** (CEA - Cambridge Event Architecture, 1995)
- ***Role-Based Access Control*** (*service-based, parametrised roles*)
- ***scalable publish/subscribe system*** (*HERMES*)
- ***Trust-based systems*** (*SECURE* - *EU Framework 5* grant)
- ***public-sector*** applications (*police, health, now transport*)

Why we're at this Workshop ...

Current research grants

TIME-EACM **(EPSRC *WINES* program)**

Transport Information Monitoring Environment

umbrella framework supported by Boeing

– ***Event Architecture and Context Management***

5-group consortium, Cambridge and Birkbeck, + ***BT***

Opera group (CL) *generic event-architecture*

Digital Technology Group (CL) *sensor deployment*

Richard Gibbens (CL) *transport data analysis*

Bill Fitzgerald (CUED) *audio-sensor interpretation*

Niki Trigoni (Birkbeck) *sensor data management*

main thrust of ***TIME-EACM*** to realise sensor-data as information

CareGrid – aspects of Trust in health-care monitoring (with ***Imperial C***)

ASIDE – *how I got into Database ...*

with Alan Macfarlane – starting in 1972 – 25-year project

parish records of Earls Colne 1400-1750 – 13 million English words

mark-up language + editor for *syntax* and *simple semantic categories*

intelligent algorithms for data interpretation (i.e. medieval Latin readers)

realise marked text as a semantic network – store in a relational database

high-level query language via relational algebra – *original text available*

demographic model – *people* and relationships, property, court records

not actually *people* – just the *names of people* – *grounding problem*

raw data – *intelligent algorithms* – information base – *grounding*

Converting data to information in the context of a sensor network

thrust of *TIME-EACM* is to deploy traffic sensors incrementally

fixed infrastructure and wireless networks, including mobile sensors

information bases available to applications, including mobile clients

intelligent algorithms to convert ***data*** to ***information***

statistical and information theoretic approaches

- *hidden Markov Models*, *e.g. for speech recognition*
- *Bayesian networks*, *e.g. for classification*

other approaches based on models at the information level

potential bulk of sensor data – where to carry out information fusion

requirement for event data archive – *audit* and *replay*

Relevant recent PhD dissertations

Middleware support for context-awareness in distributed sensor-driven systems

Eli Katsiri (DTG - then LCE in CUED) *Computer Lab TR 620*

Examples of the approaches mentioned already

- *hidden Markov Models*, *e.g. for identifying movements*
- *Bayesian networks*, *e.g. for predicting the location of users*

Separate approach based on *Temporal First-Order Logic*

- *Horn clause productions*, *inferring high-level predicates* (*CLIPS*)
- *maintaining state-based high-level predicates* (*SPASS*)

SCAFOS – application development for the *active BAT* environment

- *infers state-level predicates* from *sensor-level predicates*

PhD dissertation (Runner-up in the BCS DD series)

Ontology based Visual Information Processing

Chris Town (AI group in CL, with John Daugman of *iris* fame)

available at http://www.bcs.org/awards/dissertations/ctown.pdf

Frame and sequence analysis via specific components

- *low-level blob trackers*, *e.g. for identifying movements*
- *expert witnesses*, *e.g. for identifying texture types* (*skin*, *metal*)

Higher-level description via an Ontological Query Language *OQUEL*

- classification of static images (e.g. photograph collections)
- scene and sequence analysis by tracking multiple blobs
- *performance enhancement by data fusion of video and BAT sensors*

Fusion experiments allow a combination of sensor sources to support the same high-level scene descriptions through *OQUEL*

- the results are encouraging rather than overwhelming

Thought for the day

"You might want to say a little about how machine learning is used to tie together the different feature detectors and classifiers. The main value of using an ontology is twofold: as a human readable specification of a problem domain and as ***hierarchical prior*** that allows one to train a Bayesian Network or similar learning framework to integrate different sources of evidence (sensors, detectors, classifiers etc.)."

Chris Town

Integrating sensor networks

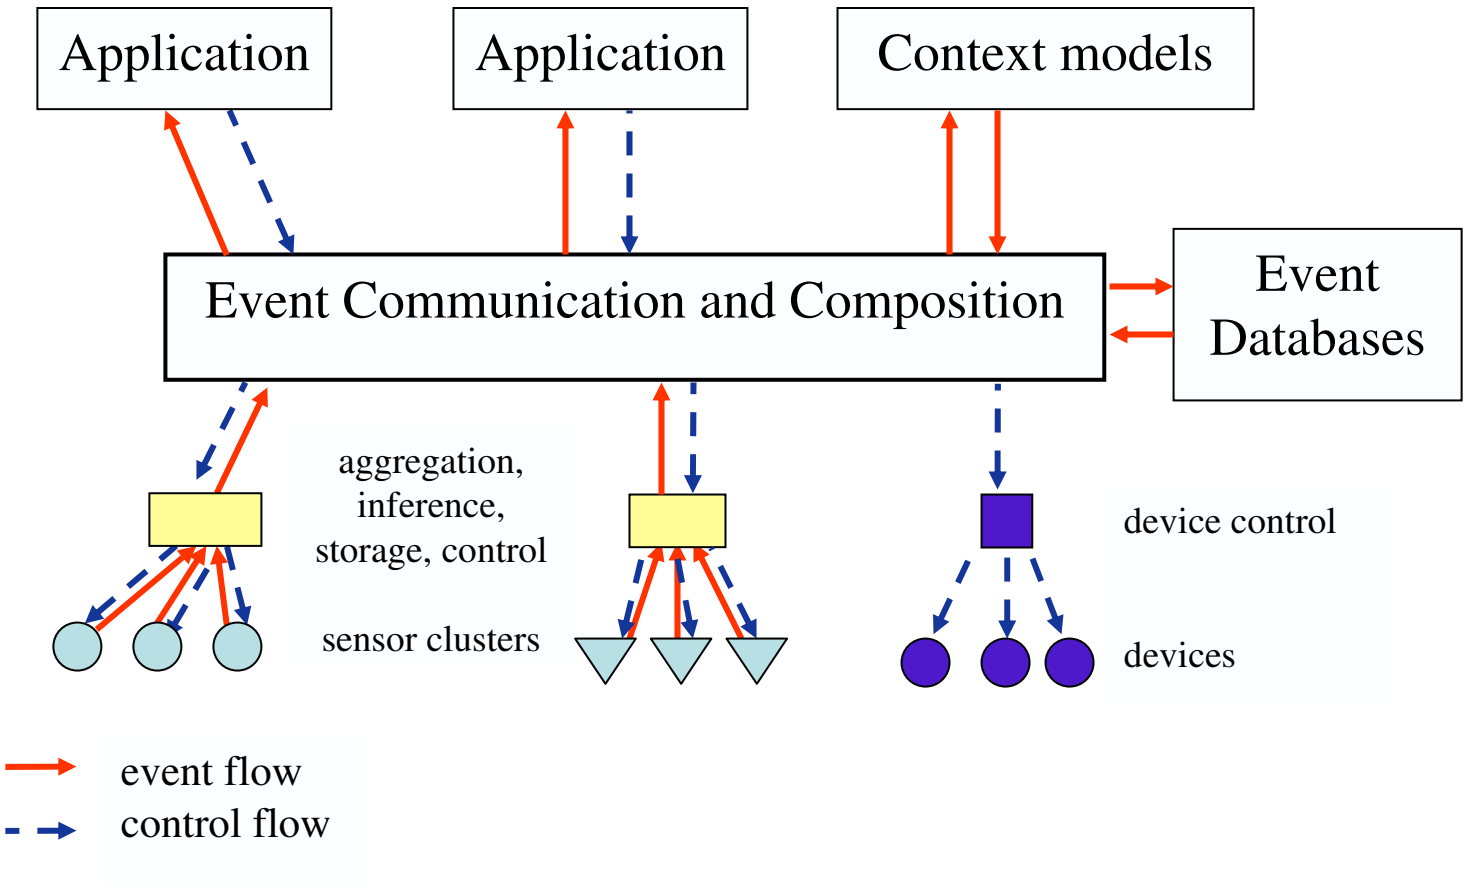

Timeliness, feedback and stability

outside the scope of *TIME-EACM* – *walk before you try to run ?*

Nyquist Theorem, need ***hard real-time*** guarantees

intention is to support real-time sensor data feeds – *latency* matters

distributed systems require sophisticated handling of ***time***

control applications need *guarantees* and *proofs* of stability properties

target applications are such ***soft real-time*** tasks as route planning

future control support implies an ***event data archive*** – *audit* and *replay*

Boeing are particularly interested in support for control applications

Supporting context-sensitive applications

clients above *TIME-EACM* may be mobile – *motion-sensitive needs*

clients may have limited bandwidth/energy – *resource-aware views*

whose responsibility is it to support context-sensitive views ?

important
- but outside the initial scope of *TIME-EACM*
- must be aware of requirements as the architecture evolves

target applications are such ***soft real-time*** tasks as route planning
- these will often be ***context-aware***
- important to consider ***use-case*** higher-level applications