



Mythbusters:

Event Stream Processing v. Complex Event Processing

DEBS 2007, Toronto

June 20, 2007

Tim Bass



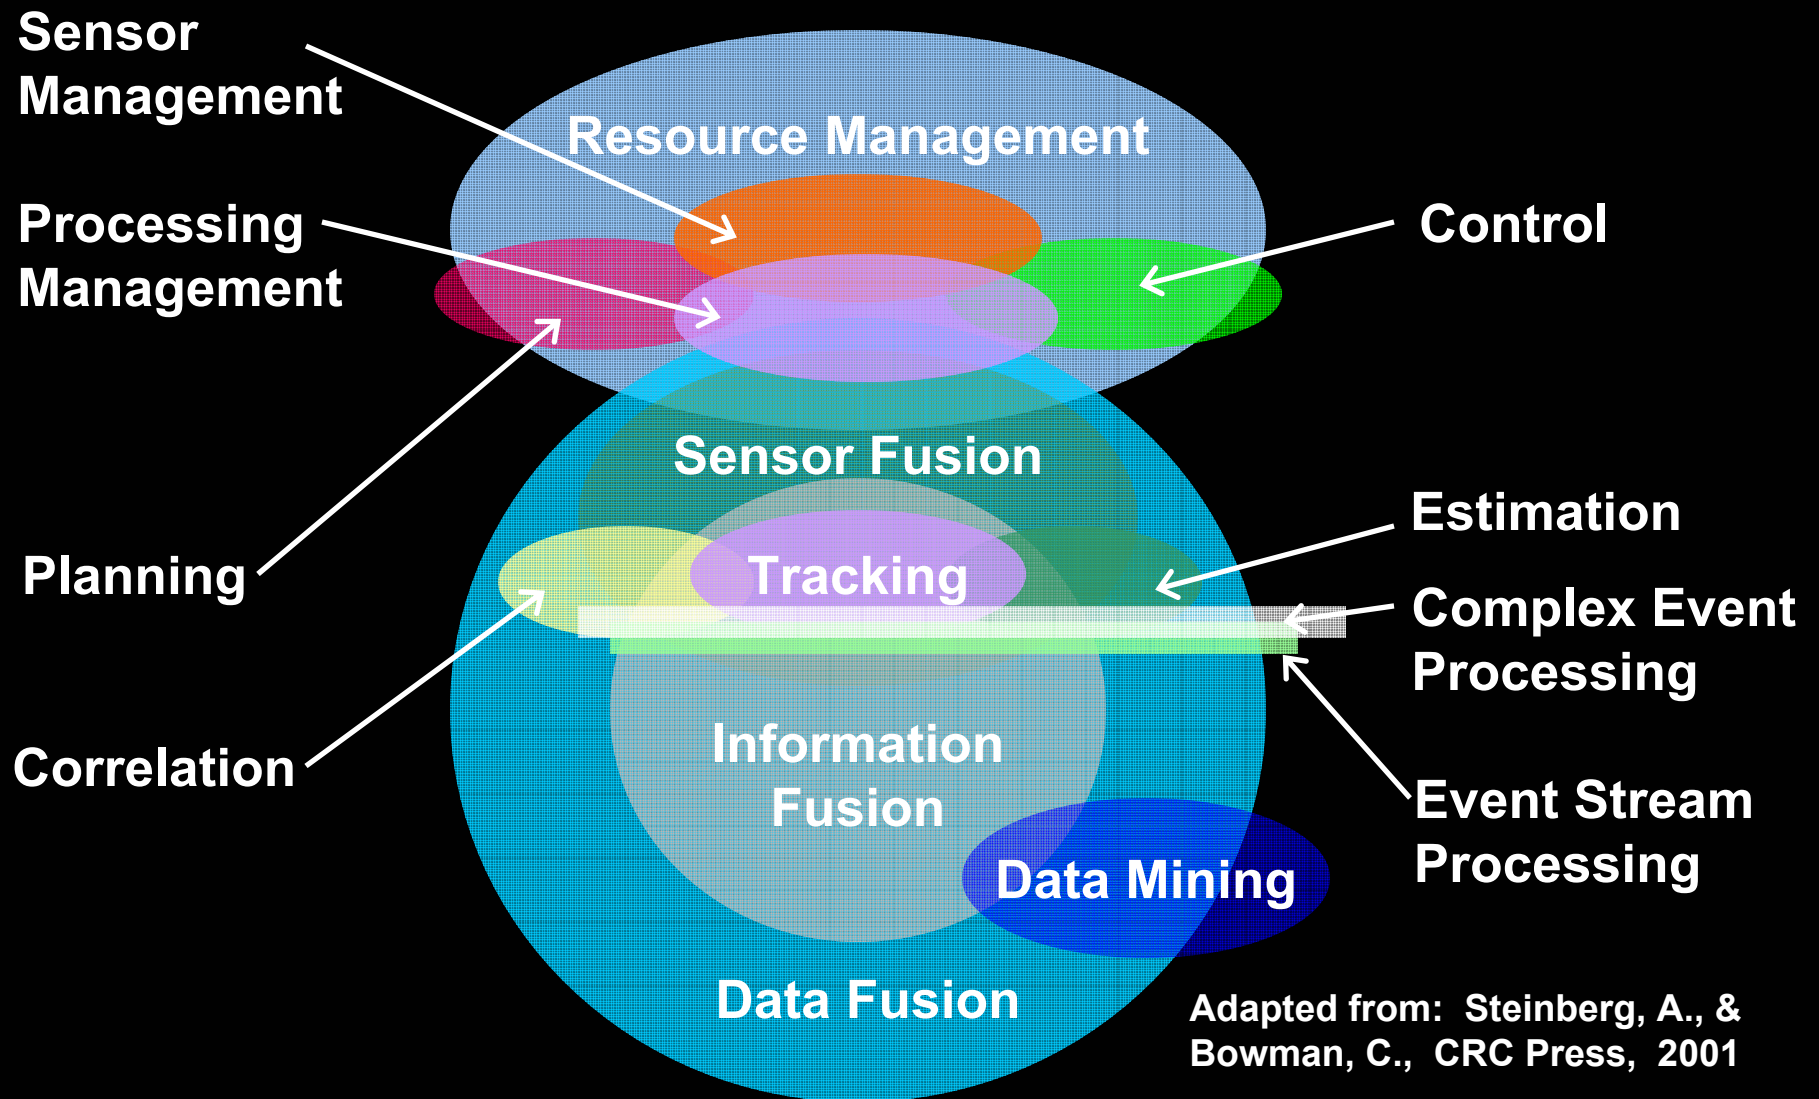
Our Agenda

An Illustrative Survey of Steams and Clouds


Event Processing Reference Architecture

Wrap-Up

A Vocabulary of Confusion (Work in Progress)

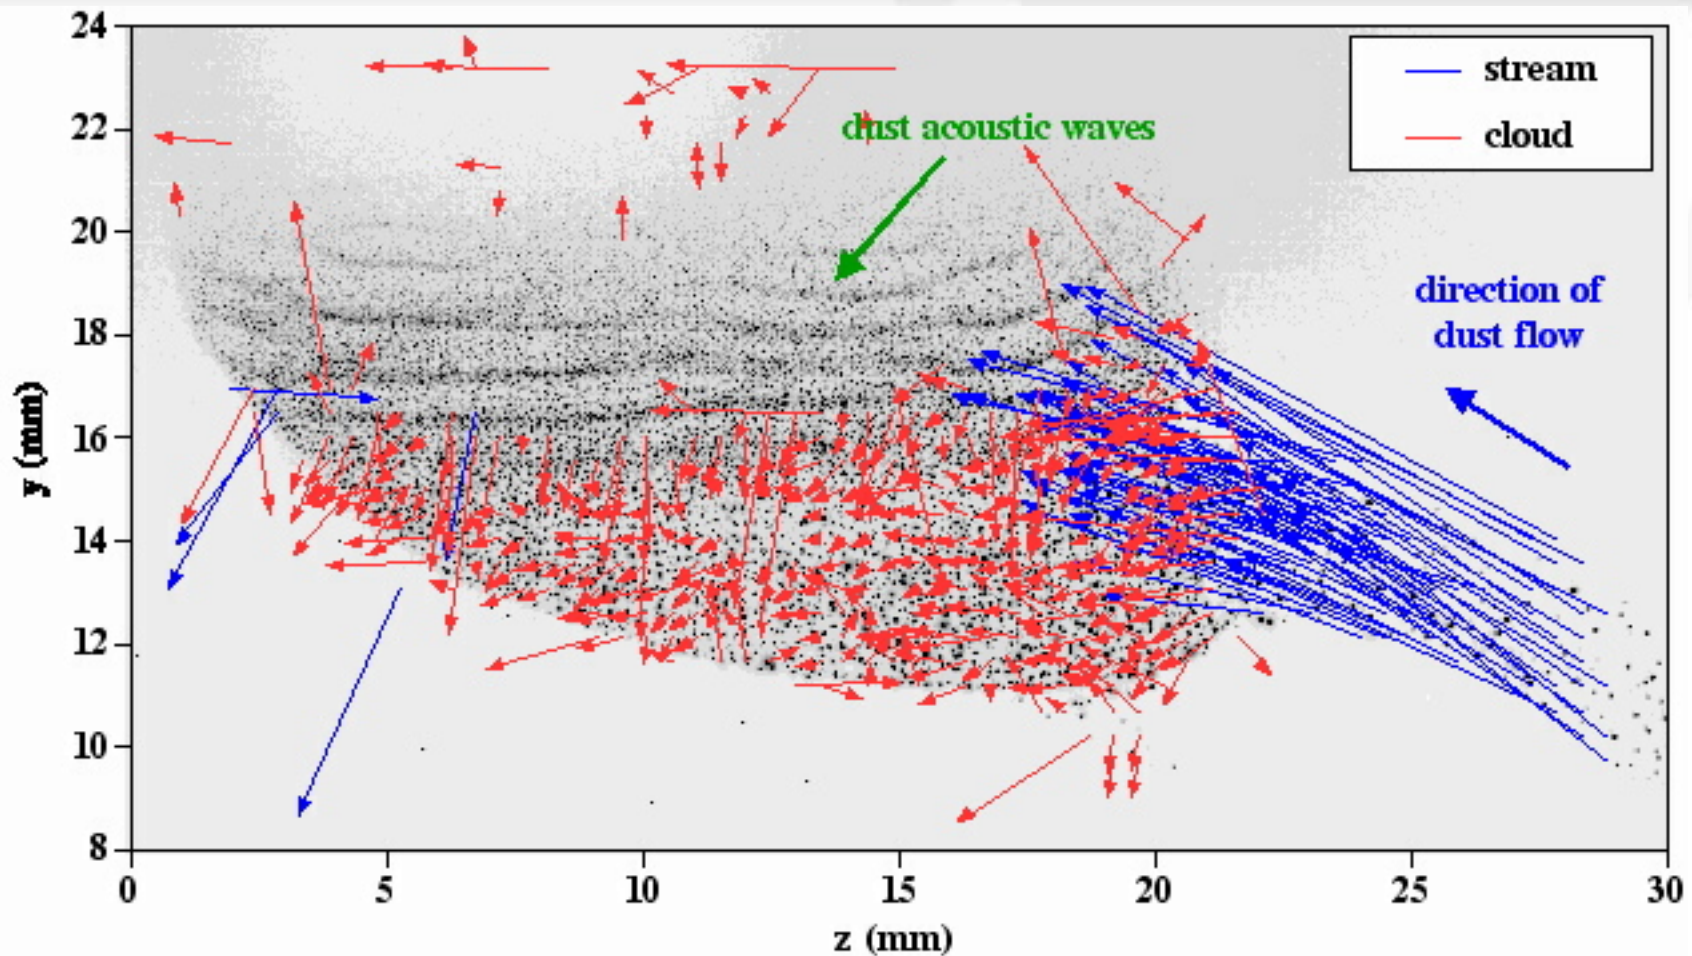


Clouds from Thermal Streams




Ref: www.paragliding.gr/cd-rom/aerology.htm

Dust Clouds and Streams

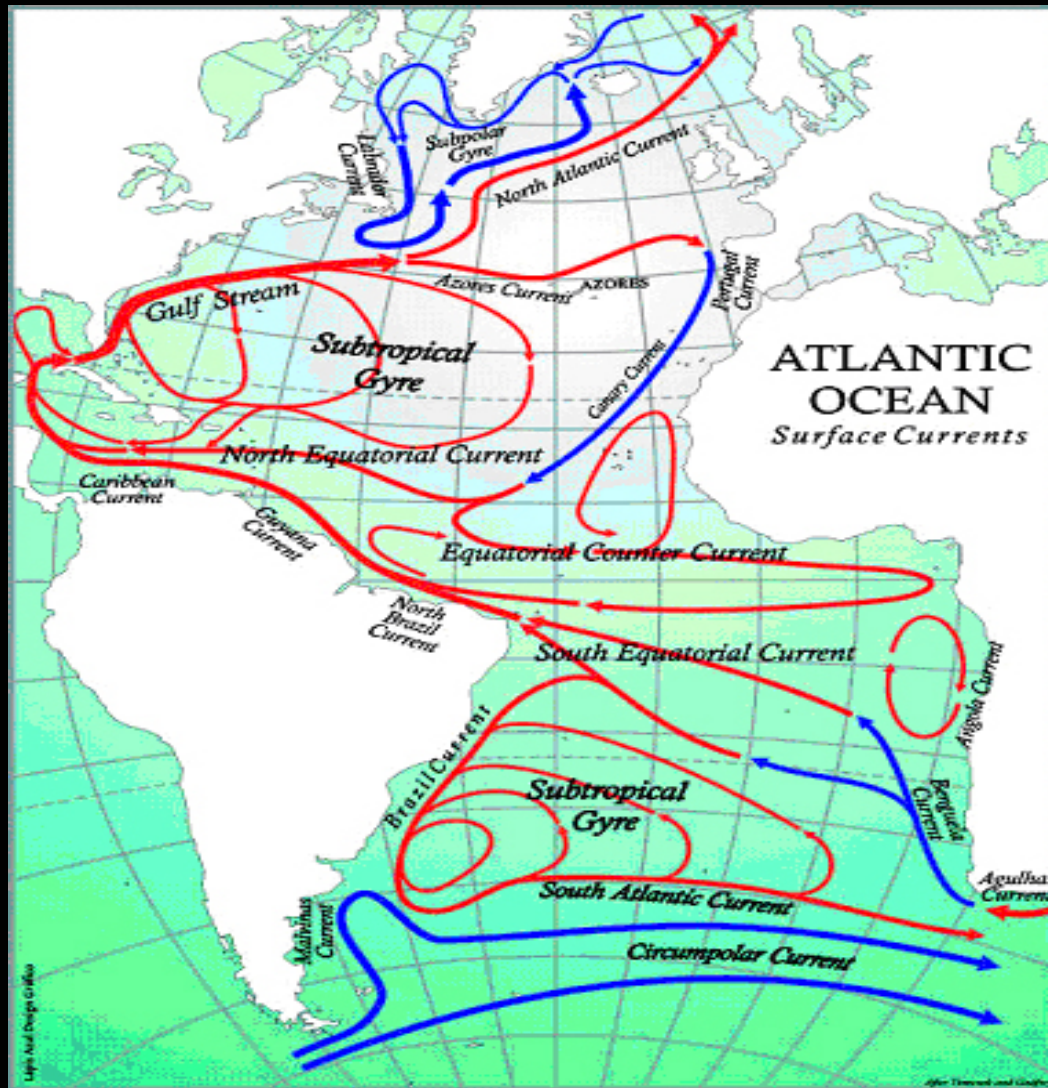


Ref: <http://narn.physics.auburn.edu/research/dusty/images/stream.jpg>

Gulf Stream

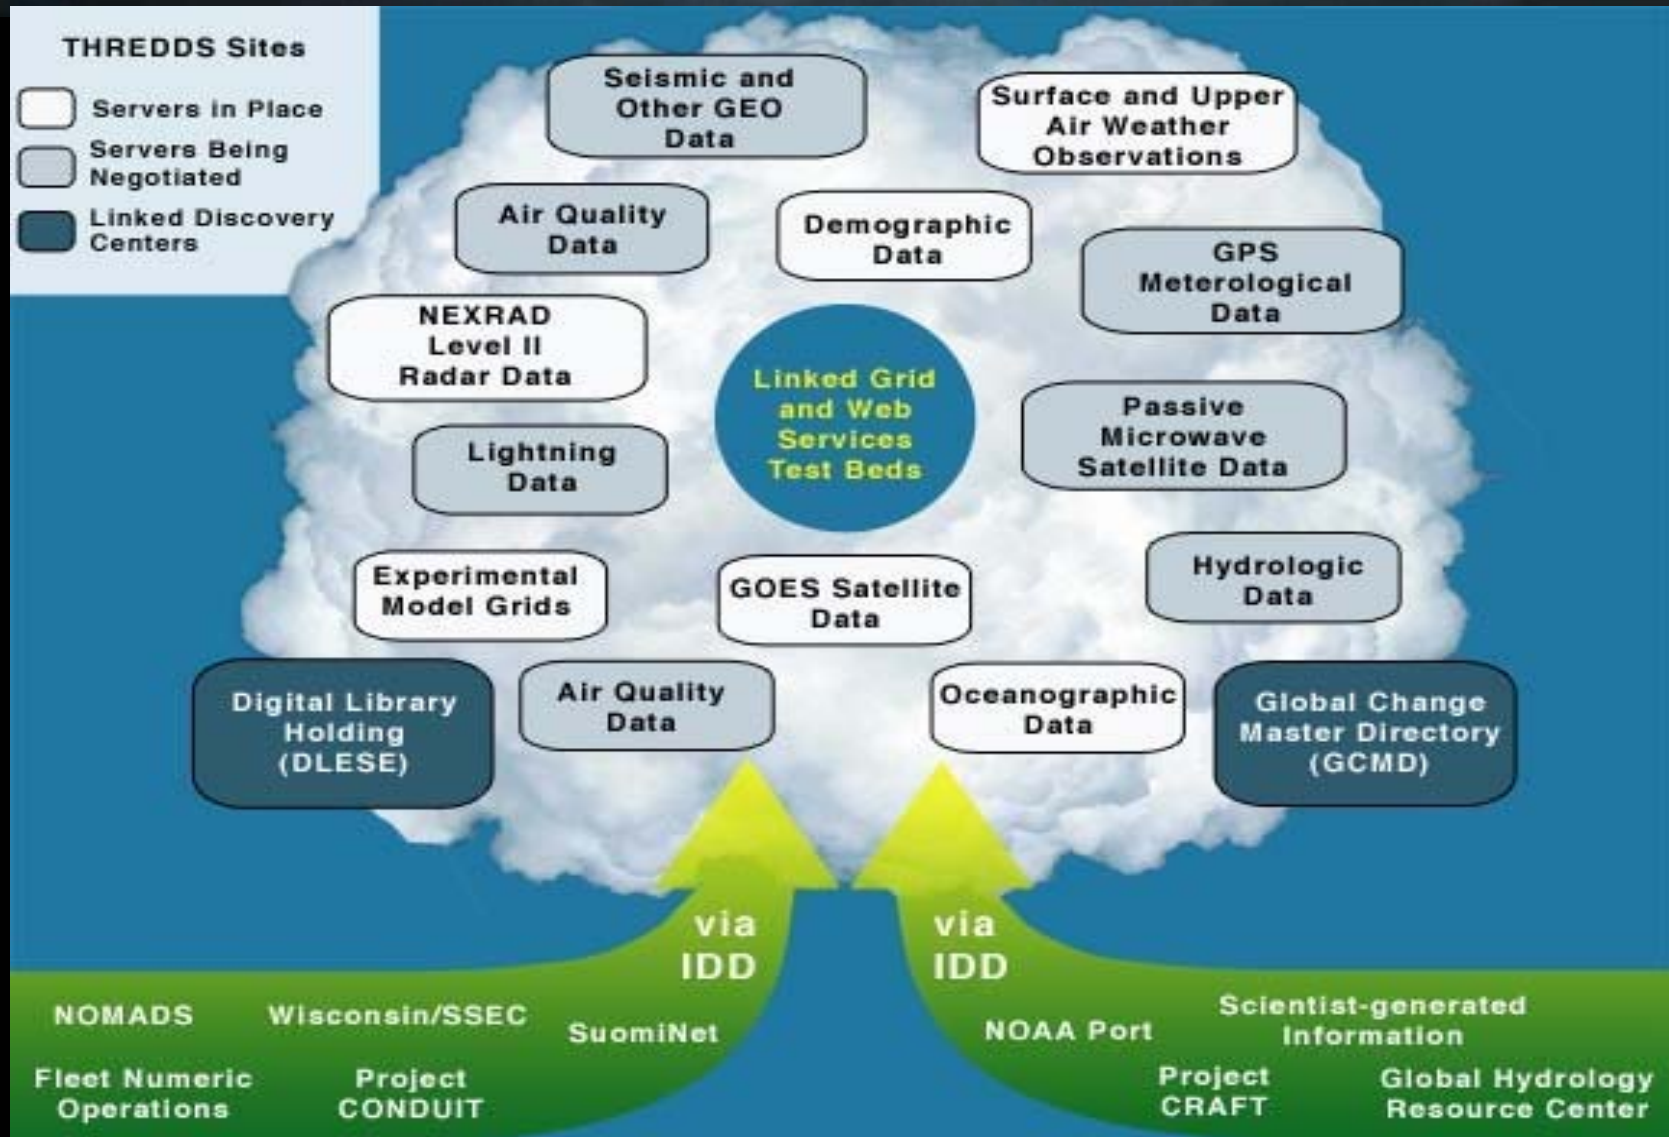


Ocean Currents as Streams



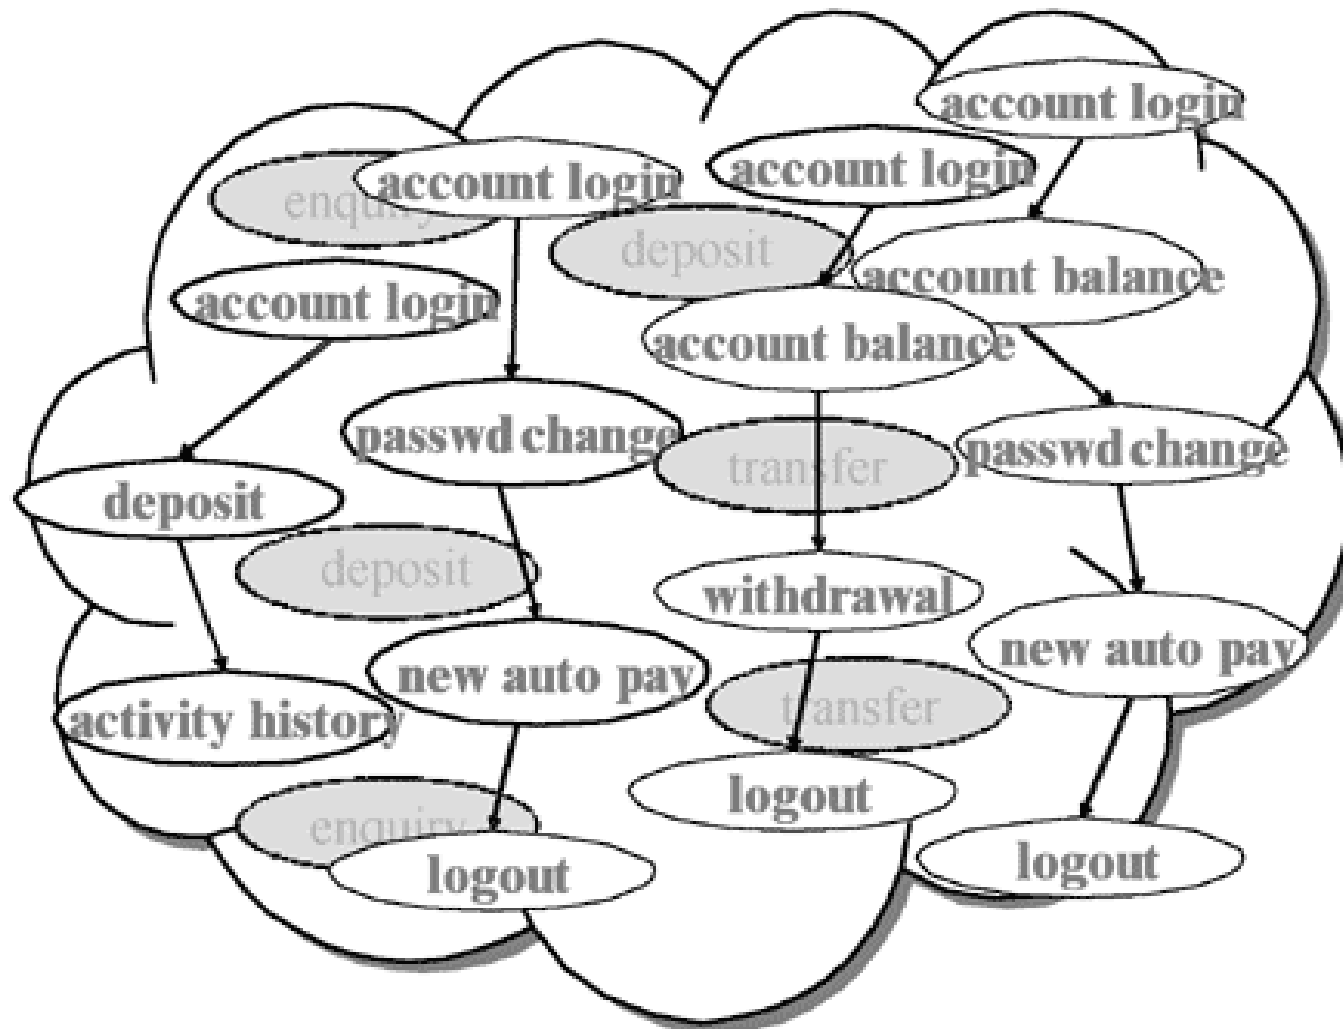
Weather Data as a Cloud

<http://meiyu.atmphys.howard.edu/hu-lead/picture/cloud.png>




Business Events as Clouds

David Luckham



David Luckham on Streams and Clouds



- An **event stream** is a special case of an **event cloud**.
- An **event stream** is a sequence of events ordered by time, such as a stock market feed.
- An **event cloud** is the result of many event generating activities going on at different places in an IT system.
- A **cloud** might contain many **streams**.

Luckham, D., What's the Difference Between ESP and CEP?,
<http://complexevents.com/?p=103> , 1st August 2006

Formally: Streams and Clouds

David Luckham, CEP-Interest Yahoo! Feb 16, 2007

STREAM a linearly ordered sequence of events.

Examples: 1, 2, 3, 4, 5, ... i.e., the integers under $<$ order.
- also, a stock market feed.

CLOUD: a partially ordered set of events.

Examples: planar points (x,y) under the order, $(x,y) < (u,v)$ if and only if $x < u$ and $y < v$. - also, all email messages on the Internet at any instant.



Clouds



Streams


An Event Stream as a type of Data Stream

- **Formally: A data stream is an ordered pair (s, Δ) where:**

- 1. s is a sequence of tuples,
- 2. Δ is the sequence of time intervals (i.e. rational or real numbers) and each $\Delta n > 0$.

- **Examples:**

- stock quotes
- click streams
- network traffic
- GPS signals
- sensor network applications.



David Luckham on Streams and ESP

- **Processing a stream of events in their order of arrival has advantages.**
 - Algorithms for processing the data in the events that use very little memory because they don't have to remember many events.
 - The algorithms can be very fast.
 - They compute on events in the stream as they arrive, pass on the results to the next computation and forget those events.
- **Event stream processing is focused more on high-speed querying of data in streams of events and applying mathematical algorithms to the event data.**

David Luckham on Clouds and CEP

- In clouds, you can't assume that events arrive in a nice order.
- You may be looking for sets of events that have a complex relationship.
- CEP applies to a richer set of business problems, not only event data processing, but also business process management, for example.
- CEP is designed for extracting information from clouds of events created in enterprise IT and business systems.
- CEP includes event data analysis, but places emphasis on patterns of events, and abstracting and simplifying information in the patterns, to span the broadest possible area of enterprise management decision.
- CEP takes more memory and more time!

- “Event streams are generally confined to an individual organization [...] and correlation systems remain within the organization’s network [...]” pp .1
- “Producers publish events into the event “cloud”, and subscribers indicate interest in classes of events, either via subscriptions to channels or by declaring interest in certain classes of context.” pp. 6

Ref: <http://www.cs.columbia.edu/techreports/cucs-049-05.pdf>

- “[Cloud ...] results from a large amount of distributed activities around the world and it’s usually unstructured.” pp .17
- “Producers publish events into the event “cloud”, and subscribers indicate interest in classes of events, either via subscriptions to channels or by declaring interest in certain classes of context.” pp. 6

Ref: <http://wwwhome.cs.utwente.nl/~zhangyy/thesisv2.pdf>

Event Processing Technical Society WG

<http://complexevents.com/?p=124>

- ***Event stream:*** a linearly ordered sequence of events.
 - *Notes:* Usually, streams are ordered by time, e.g., arrival time. An event stream may be bounded by a certain time interval or other contextual dimension (content, space, source, certainty), or be open ended and unbounded.

- ***Event cloud:*** a partially ordered set of events (poset), either bounded or unbounded, where the partial orderings are imposed by the causal, timing and other relationships between the events.
 - *Notes:* Typically an event cloud is created by the events produced by one or more distributed systems. An event cloud may contain many event types, event streams and event channels. The difference between a cloud and a stream is that there may not be an event relationship that totally orders the events in a cloud. A stream is a cloud, but the converse is not necessarily true.

POSET (A,R) : SET (A) and Relationship (R)

Set A	Relationship R	Abstraction
Set of all stock trades for GOOG within a 5 minute time window	Linearly ordered by time. A chain of events.	Event stream
Set of all stock trades in NASDAQ for a single day.	Partially ordered Many incomparable chains of events.	Event cloud
Set of all log file entries in a single banking application	Linearly ordered by time. A chain of events.	Event stream
Set of all banking systems	Partially ordered.	Event cloud

CEP Clouds and ESP Streams

A Stream is a Special Case of a Cloud.

EVENT CLOUDS	EVENT STREAMS
Network Management	Web Traffic Session Extraction
Weather	Temp Data from Station Zulu
Enterprise Security / SEM	Alerts from Firewall or IDS
Insider Trading / Fraud	EOG Financial Transactions
Logistics / Supply Chain	Tracking RFID Information
Global Epidemiology	Tylenol OTC Sales
Homeland Security	San Ysidro Border Crossings
Investor Sentiment	Market Data from NASDAQ Sales

Philip Howard on CEP v. ESP

Bloor Research Analyst Opinion

- *“A typical ESP application is one such as algorithmic trading ...”*
- *“CEP is about what we might call über-events or, more specifically, patterns of events.”*
- *“CEP is a superset of ESP.”*

Howard, P., *ESP and CEP...what's the difference?*,
Reg Developer, 6 June 2005



Our Agenda

An Illustrative Survey of Steams and Clouds

Event Processing Reference Architecture

Wrap-Up

Data Clouds → Actionable Knowledge

Impact Assessment

Situational Assessment

Relationship of Events

Identify Events

Location, Times and Rates
of Events of Interest

Existence of Possible
Event of Interest

Data/Event Cloud

HIGH

MED

LOW

Analysis of Situation & Plans

Contextual and Causal
Analysis

Causal Analysis, Bayesian
Belief Networks, NNs,

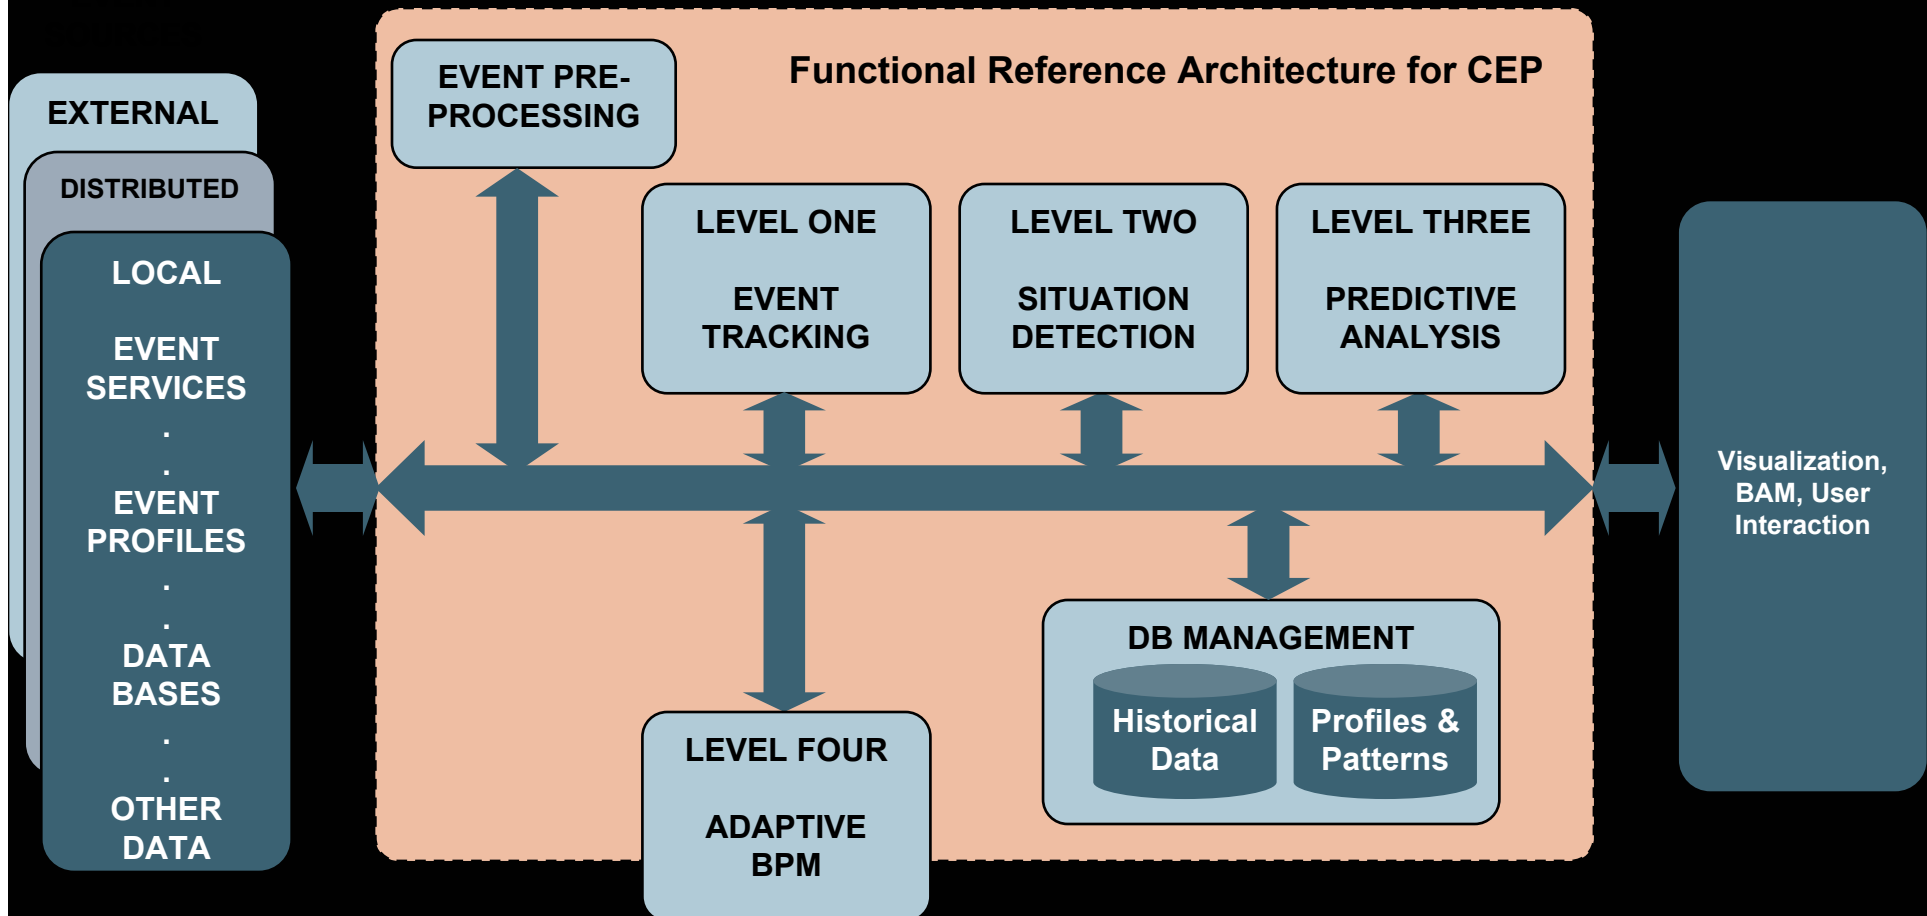
Correlation, State Estimation,
Classification

Use of Distributed
Sensors for Estimations

Raw Sensor Data
(Passive and Active)

Adapted from: Waltz, E. & Llinas, J., Multisensor Data Fusion, 1990

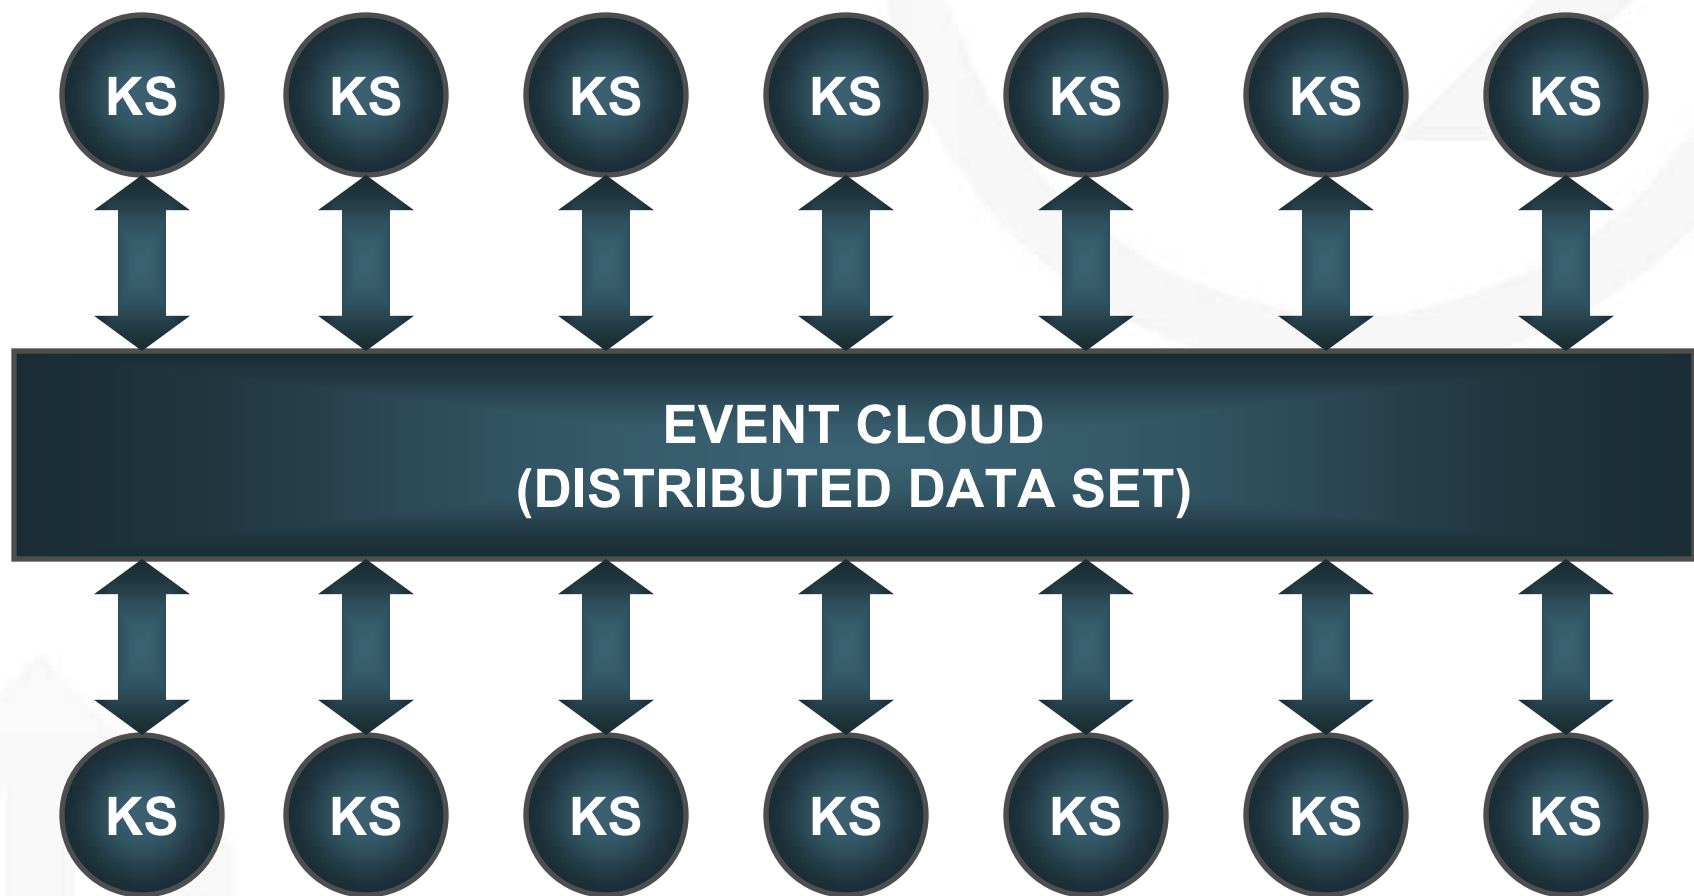
Functional Reference Architecture



Adapted from JDL:

Steinberg, A., & Bowman, C., Handbook of Multisensor Data Fusion, CRC Press, 2001

AI, Agents and Blackboards



Adapted from: Englemore, R. S., Morgan, A.J., & Nii, H. P., Blackboard Systems, 1988 & Luckham, D., The Power of Events, 2002

Event Processing and Data Fusion

- Multi-level inference in a distributed event-decision architectures
 - **User Interface (Dashboards, BAM, Visualization, Portals)**
 - Human visualization, monitoring, interaction and situation management

Level of Inference

High

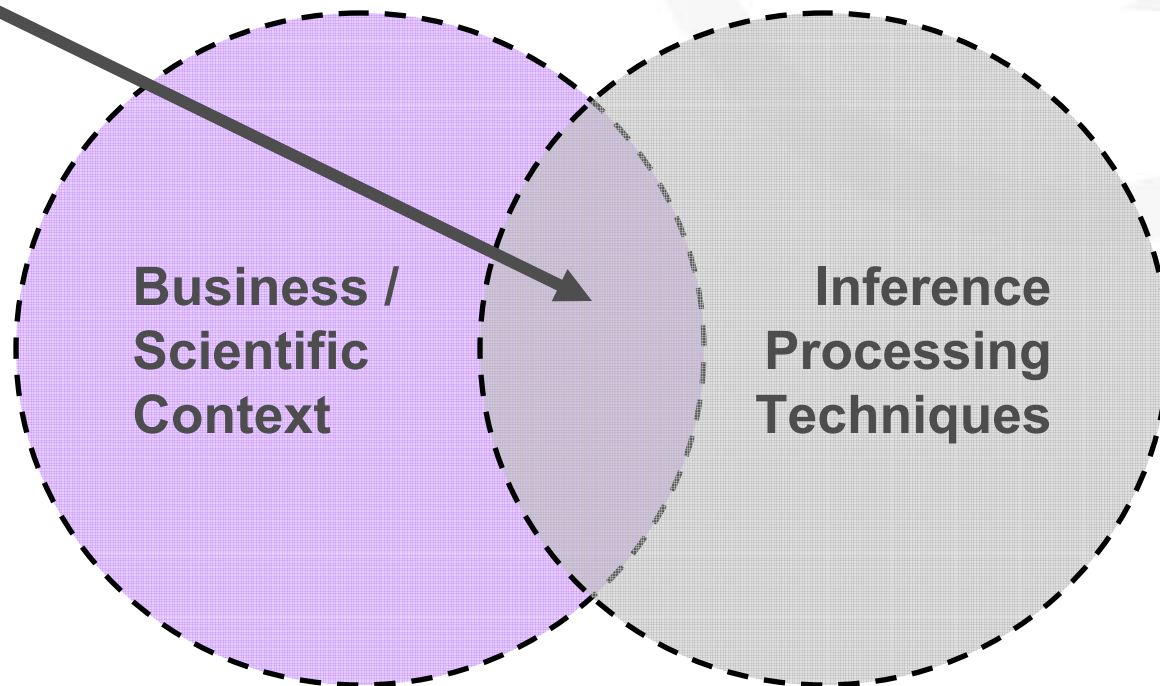
Med

Low

- **Level 4 – Process Refinement (Adaptive BPM)**
 - Decide on control feedback, for example resource allocation, sensor and state management, parametric and algorithm adjustment
- **Level 3 – Impact Assessment (Predictive Analytics)**
 - Impact assessment, i.e. assess intent on the basis of situation development, recognition and prediction
- **Level 2 – Situation Refinement (Situational Detection)**
 - Identify situations based on sets of complex events, state estimation, etc.
- **Level 1 – Event Refinement (Event Tracking)**
 - Identify events & make initial decisions based on association and correlation
- **Level 0 – Event Preprocessing**
 - Cleansing of event-stream to produce semantically understandable data

Processing Patterns for Event Processing

Processing Patterns

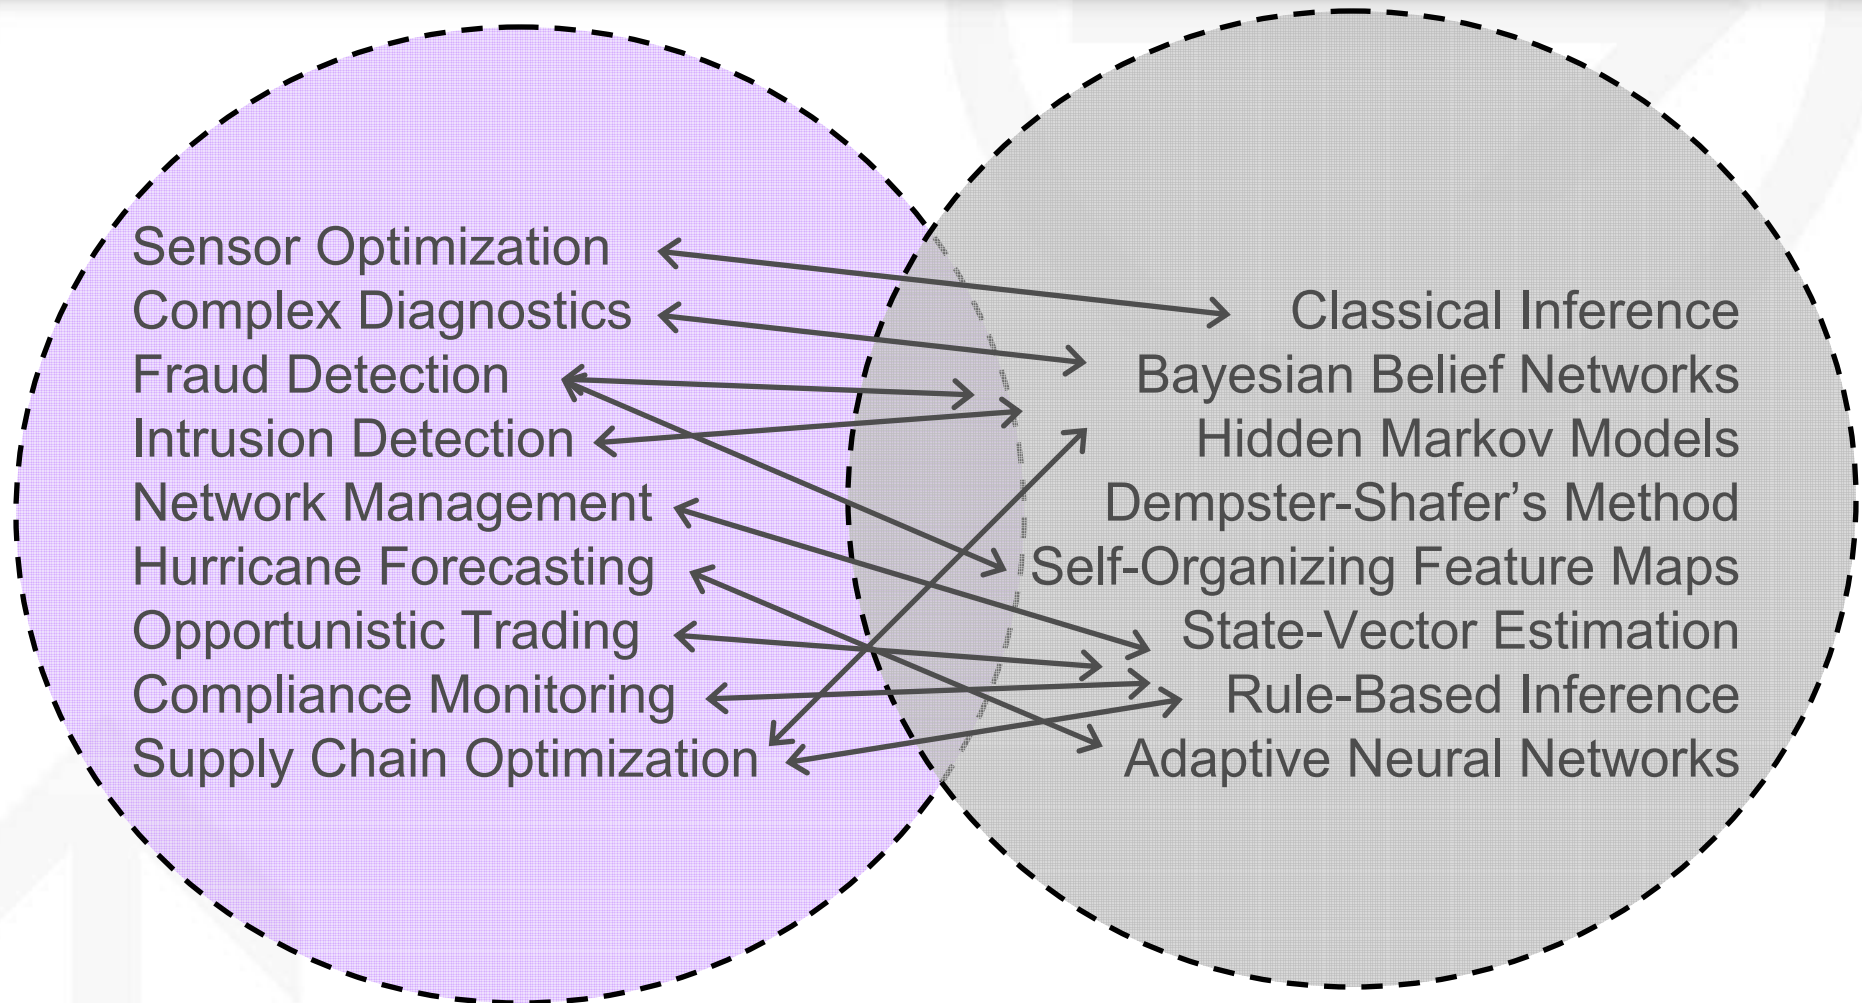


Inference Algorithms for Event Processing

- **A sample of event processing algorithms relevant to CEP:**
 - Rule-Based Inference
 - Bayesian Belief Networks (Bayes Nets)
 - Dempster-Shafer's Method
 - Adaptive Neural Networks
 - Cluster Analysis
 - State-Vector Estimation
- **Key Takeaway:**
 - Analytics for CEP exist in the art & science of many disciplines - these analytics can be mapped to recurring business patterns.
 - Rules are only one of many methods for CEP.

Map Business Context to Classical Methods

Note: For Illustrative Purposes Only



Business / Science Context

Inference Processing Techniques

Example Inference Patterns for Event Processing

- **Bayesian Techniques:**


- SPAM Filtering
- Telecommunications Fraud Detection
- Fraud & Intrusion Detection
- Financial Risk Management
- Credit Approval and Credit Limit Automation
- Medical Diagnosis
- Military ID, Command and Control

- **Rules**

- Algorithmic Trading
- Routing and Scheduling
- Fraud Detection

ESP and CEP are One in the Same
Rules are the Brains of CEP

Mythbusted





Thank You!

Tim Bass

**The Complex Event Processing Blog -
eventprocessing.wordpress.com**

www.silkroad.com

Appendix – Backup Slides

- **Example Stream Processing Challenges**
- **Posets and Tosets**

Example Stream Processing Challenges

- Noise from the data sources
- Resource management of the system resources
- Evolutionary changes in data trends
- Approximate query answering
- Limited memory
- Limited storage
- Limited bandwidth
- Limited processing
- Out of order data



Partially Ordered Sets of Events

■ POSET

- A relation R on a set A is called a partial order if R is reflexive, antisymmetric, and transitive. The set A together with the partial order R is called a partially ordered set or poset, and is denoted (A, R) .
- R is a partial order on A if it has:
 - Reflexivity: $a \leq a$ for all $a \in A$
 - Antisymmetry: $a \leq b$ and $b \leq a$ implies $a = b$
 - Transitivity: $a \leq b$ and $b \leq c$ implies $a \leq c$

Ref: <http://mathworld.wolfram.com/>

Linearly or Totally Ordered Sets (TOSETS)

In totally ordered sets of events, all events are comparable to the others.

■ Linearly or Total Ordered Set of Events





- If (A, R) is a poset, we say A is totally ordered if for all $x, y \in A$ either $x R y$ or $y R x$. In this case R is called a total order
- R is a total order on A if it has:
 - Reflexivity: $a \leq a$ for all $a \in A$
 - Antisymmetry: $a \leq b$ and $b \leq a$ implies $a = b$
 - Transitivity: $a \leq b$ and $b \leq c$ implies $a \leq c$
 - Comparability: for any $a, b \in A$, either $a \leq b$ or $b \leq a$.
- The first three are the axioms of a partial order, while addition of the trichotomy law defines a total order.

Ref: <http://mathworld.wolfram.com/>



Relations

- **What does the relation mean?**
 - Anything consistent with the properties of an order relation may be considered a partial order.
- **Examples of relations include:**
 - geometrical containment (e.g., in the work by W. Kainz)
 - membership among building parts (my work about scene space)
 - being earlier in a sequence (e.g., in Kuipers 1979)
- **The first two are specific instances of a general *part-of* relation.**

Gallery of Posets Produced by the package Posets.nb



Posets Illustrated



<http://www.math.unc.edu/Faculty/rap/DCLook.html>