




Research @ Databases and Distributed Systems Group

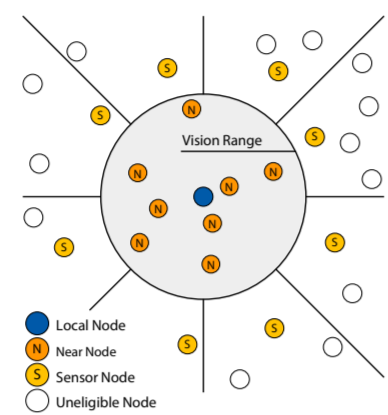


Planet PI4



- Massively multiplayer online game prototype for the evaluation of (P2P) overlays
- Parameterizable space scenario
- Configurable bots for workload generation

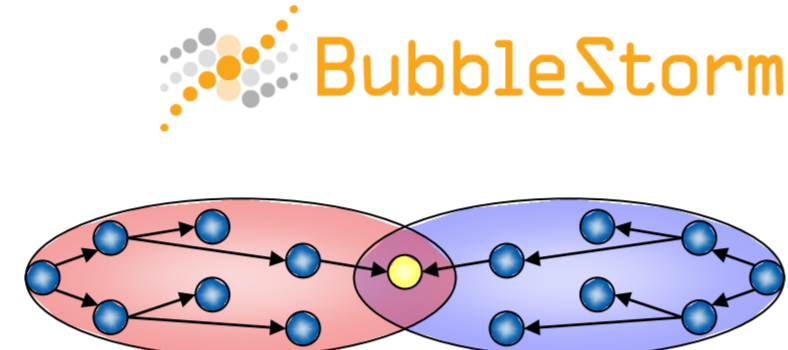
pSense



- Dynamic P2P overlay for virtual environments
- Low-latency communication within vision range

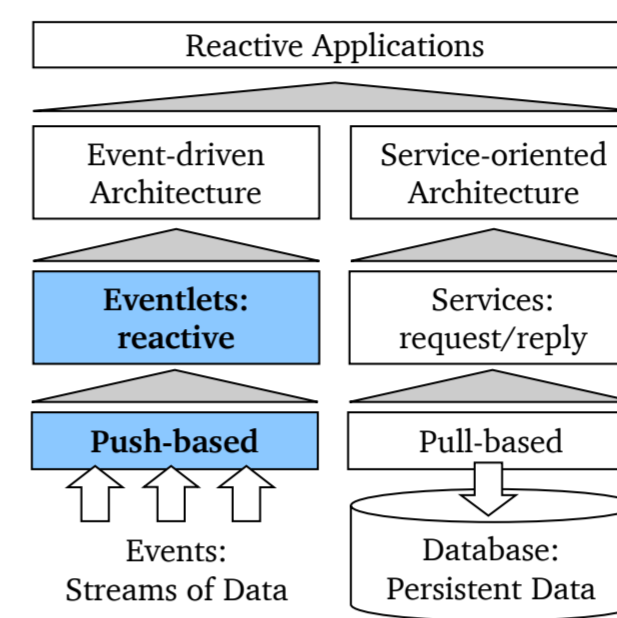
BubbleStorm

- P2P information management system
- Rendezvous search
- Different replication modes
- Store/retrieve, publish/subscribe, and more complex usage patterns




Reactive Enterprise Software Systems

- Integration of events and event streams with business applications
- Support for event streams from business process modeling to business process execution
- Eventlets as push-based equivalent to pull-based services



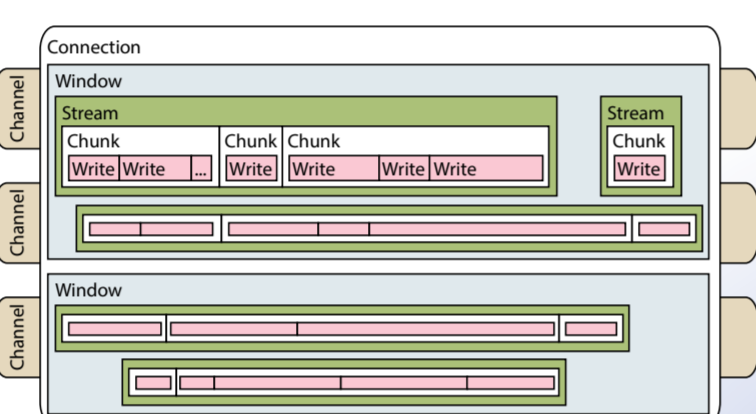
SPECjms2007

- Standard benchmark for performance evaluation of message-oriented middleware that implements the Java Message Service (JMS)
- Jms2009-PS: Modified version to test publish/subscribe communication
- Realistic workload: supply chain management of supermarkets



Novel Transport Protocols - CUSP

- Generic transport protocol optimized for the demand of P2P applications
- Variable number of potentially short-lived, unidirectional streams on a common channel



Contextualization

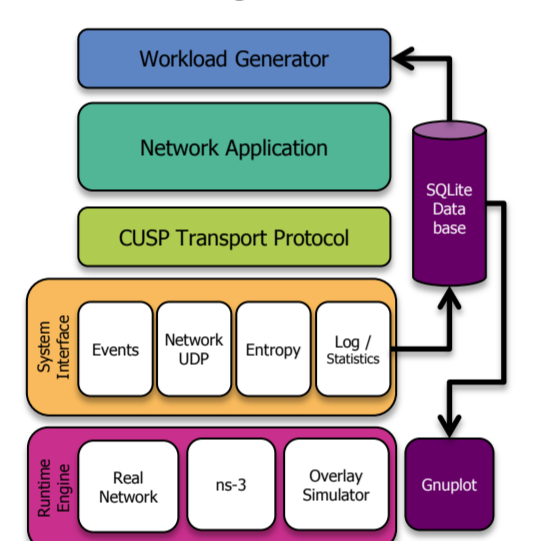
- Events need to be put into context to be usable:
- Dynamic integration of software systems across enterprises
- Transformation of events to cope with different semantics and data formats (ACTRESS)
- Contextualize by adding external information on demand

Aggregation for Pub/Sub Systems

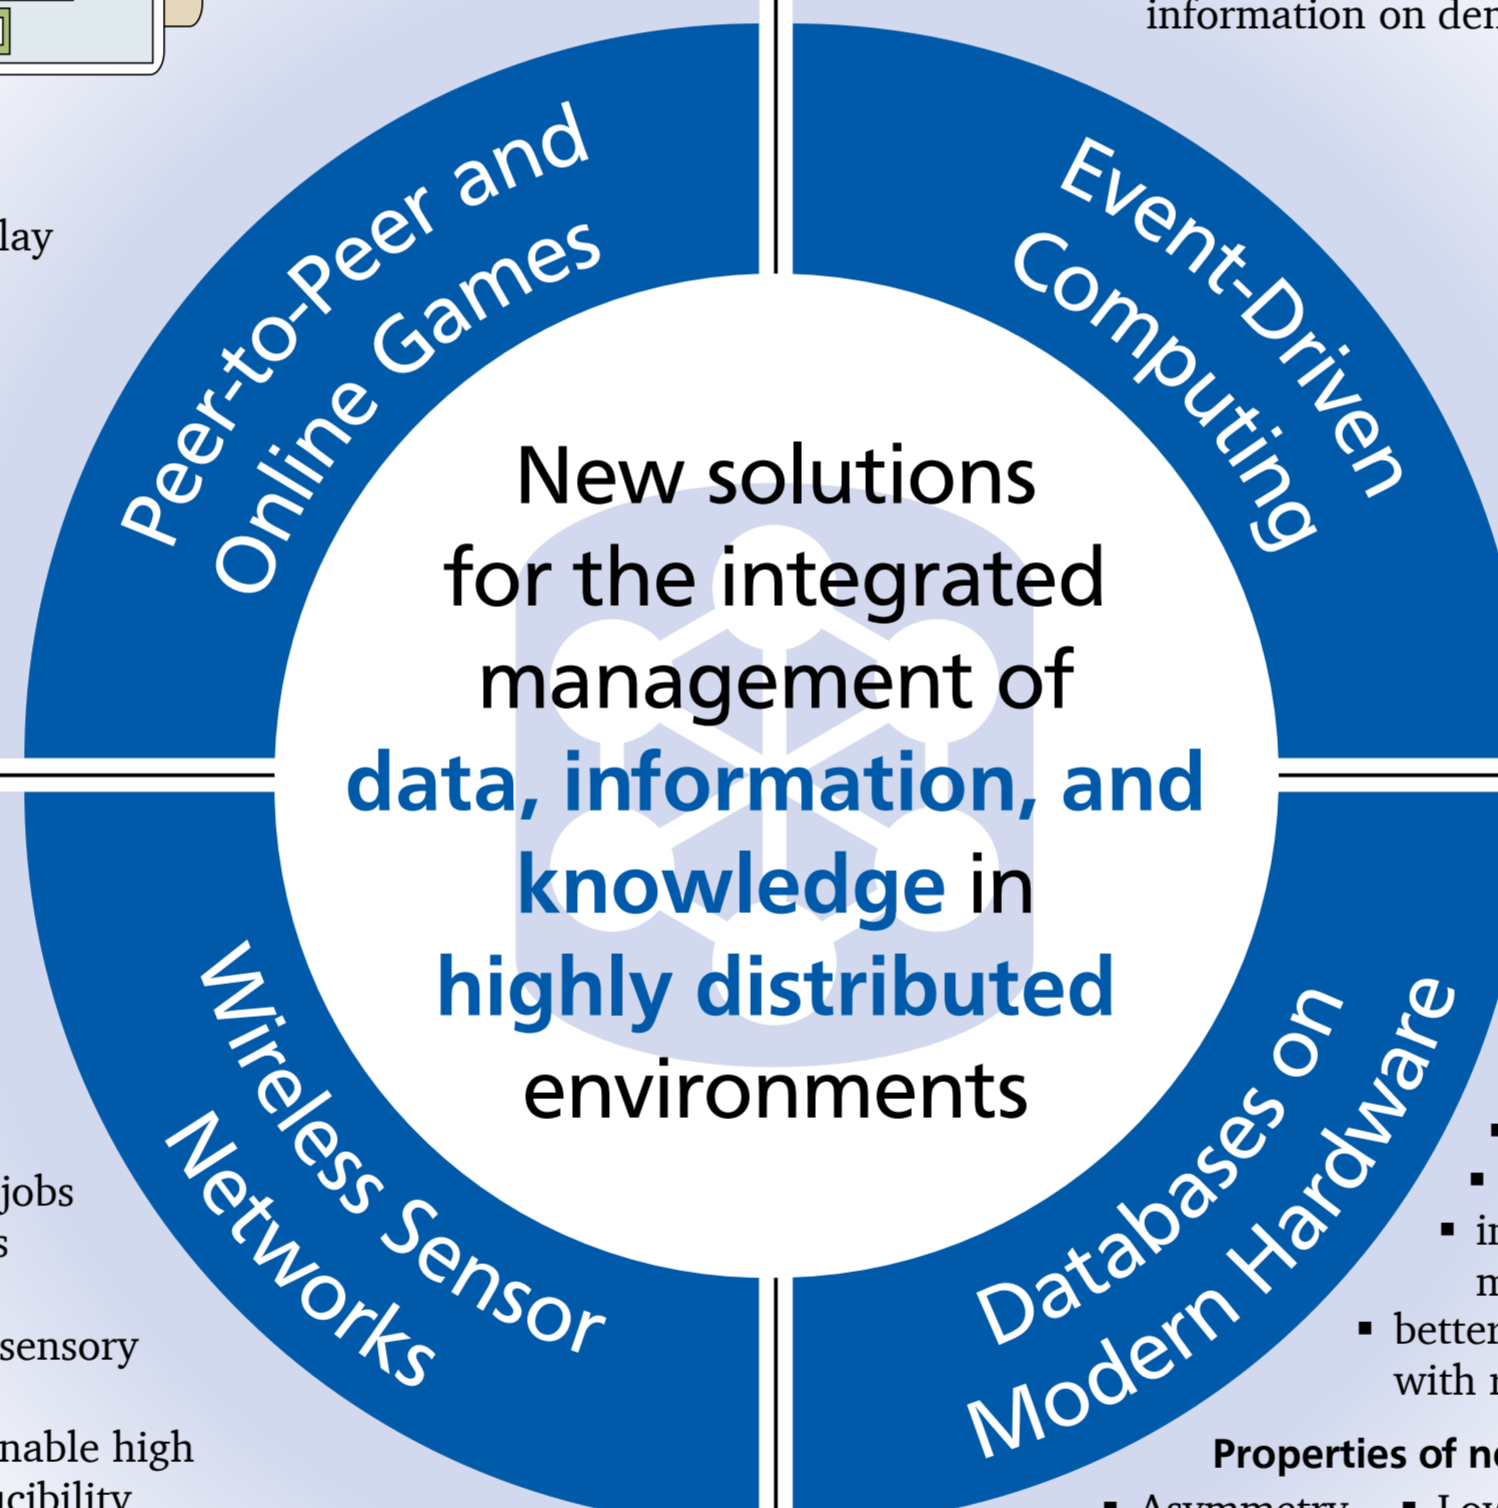
- Lightweight support for in-network aggregation of information in content-based publish/subscribe systems, e.g., number of subscribers
- International cooperation between TU Darmstadt, Purdue University, Imperial College London, University of Otago, and Vrije Universiteit Amsterdam



Overlay Network Testbed

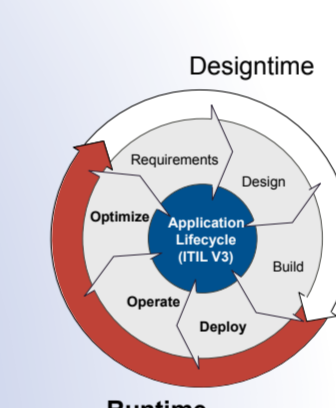


- Testbed for the evaluation of overlay networks and their applications
- Deterministic simulation and real network mode using the same application code
- Configurable workloads
- Central database of results, analysis tools



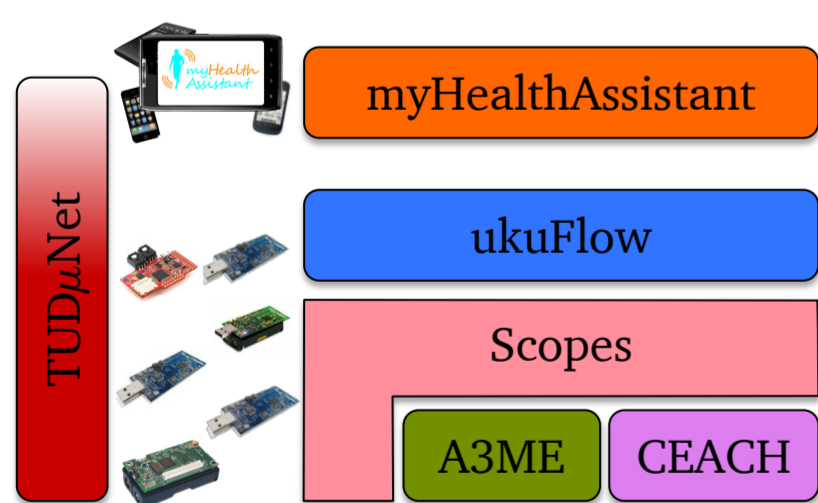
Governance in Event-based Systems

- Service level agreements (SLA) are an essential mechanism in service-oriented architectures
- Integration of event-based components in enterprise applications requires SLA concepts suitable for the push-based nature of event integration in enterprises



myHealthAssistant

- Smartphone middleware for body sensor networks (BSNs)
- Handles sensor communication
- Provides data abstraction to application
- Transforms sensor data into required data abstraction
- (Preventive) health care applications built on-top for vital sign, fitness workout and daily activity monitoring



TUDµNet

- Federation of multiple, autonomous WSN testbeds:
 - multiple, application-specific domains
 - centralized coordination, hierarchical zones, parallel jobs
- Combine best from simulators and real-world deployments
 - real platform, wireless and sensory environment
 - simplify experimentation, enable high level of experiment reproducibility

Scopes


- Multiple, parallel apps. on sensor networks
- Declarative creation of node groups, called scopes:
 - CREATE SCOPE Room6Motion (EXISTS SENSOR Motion) AS SUBSCOPE OF Room22;
- Bidirectional comm. link between members and root node
- Automated maintenance: support for network dynamics (node churn)

A3ME

- Enable interoperability between devices: sensor nodes, mobile phones, etc.
- Device agents to model entities
- Common, extensible agent ontology to classify device agent properties
- Query language, encoded in ASN.1 Packet Encoding Rules

ukuFlow

- Workflows as mechanism to define WSN application logic
- Extensions to Business Process Modeling Notation (BPMN) v2.0 to define app. logic:
 - imperative and event-driven macro-programming
 - abstracts away internal complexities of routing, grouping, data collection, event detection and action execution, among other
- entirely in-network workflow engine for low power, 8/16-bit embedded microcontrollers
 - virtual machine (interpreted)



CEACH

- Clustering service for WSN traffic
- Considers different aspects for cluster creation:
 - dynamic properties (e.g., residual energy)
 - link quality metrics (both HW and SW)
 - node degree (size of 1-hop neighborhood)
 - network aspects (hop distance)
 - use EWMA for cluster quality assessment
- Adaptive node switching between clusters

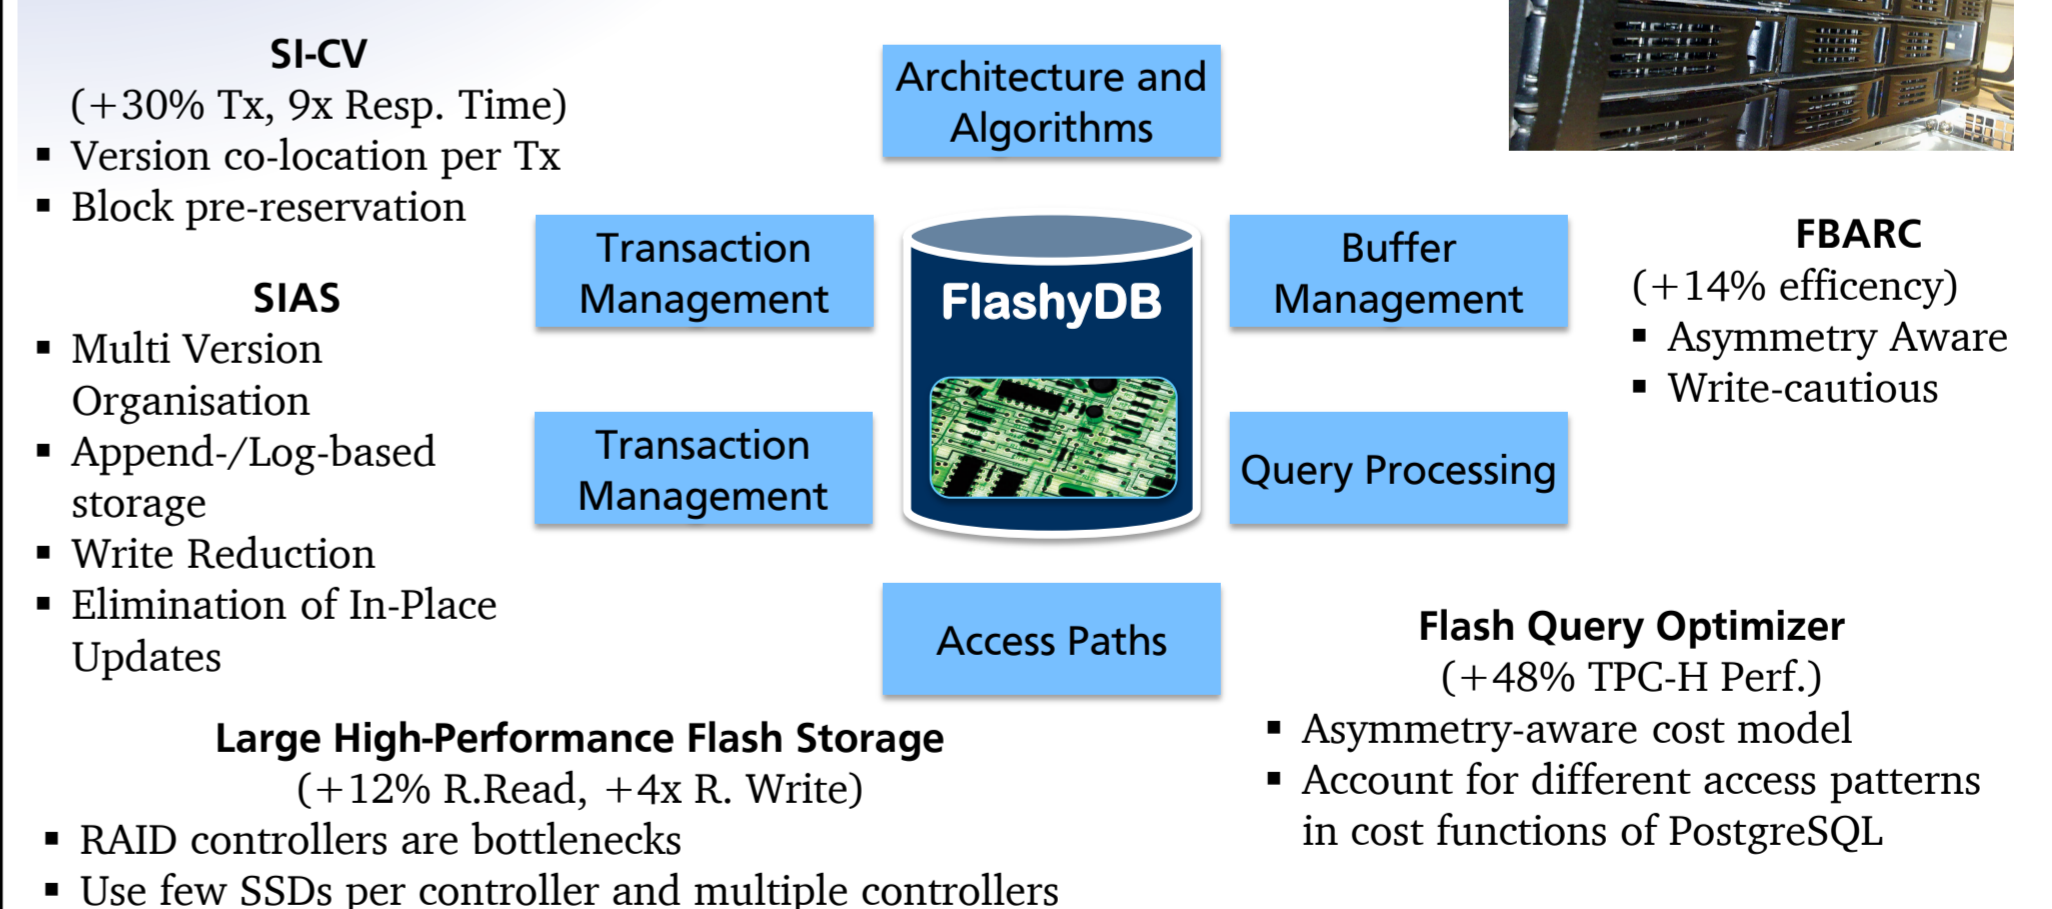
The architecture and algorithms of database systems are designed around the characteristics of legacy hardware:

- Access gap | Low Memory/Disk volume ratios | few CPUs
- Spinning Disks - HDD
- Trends in Hardware [1]:
 - many-core CPUs - 1000 cores per chip by 2022
 - growing RAM volumes - more than 128 TB RAM per server by 2022
 - increasing interconnect bandwidth
 - memory: 2.5 TB/s, I/O: 250 GB/s by 2022
 - better I/O technologies - fast 1TB Flash chips and Non-Volatile Memories with read performance similar to that of DRAM

Properties of new IO Technologies:

- Asymmetry
- Low latency → shrinking access gap
- Endurance
- Parallelism
- Addressability (Byteaddress vs. Blockaddress)
- Access Patterns (random, sequential | read, write)

What is the impact on data-intensive systems?



Efficient use of new hardware is a dominating trend in data-intensive computing, influencing the architecture, and algorithms

[1] Andreas von Bechtolsheim. Technologies for Data-Intensive Computing. HTTPS 2009.