

Cray XMT2

uRIKA

Overview

May 21, 2012

uRIKA == universal RDF integration Knowledge Appliance

Agenda



Big Data Analysis

The Cray XMT2

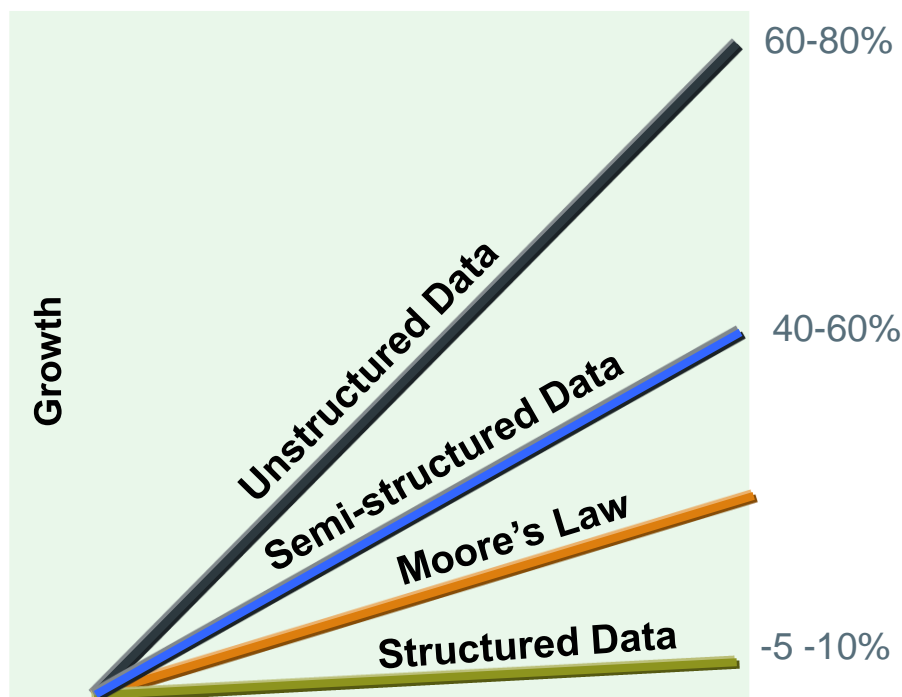
The uRIKA Graph Analytics Appliance

Customer Use Cases

Summary

Disruptive Change in Big Data Analysis

Moore's Law vs.
Growth in Dataset Size



- Exponential Growth in Overall Data Volume
- Variety of Data Types increasing
- Regulatory Requirements growing...
- *Unstructured and Semi-Structured Data becoming key!*
- Gartner: “Success goes to business which can leverage all available data... at the greatest Velocity”


- Structured: Databases, Spreadsheets...
- Semi-structured: XML, EDI, ...
- Unstructured: E-mail, Docs, Multimedia, Wikis, Social Media, ...

**Volume, Variety, Velocity:
New demands for Data Analytics**

Disruptive Change in Search Technology

Web 2.0:

- Hyperlinked Documents
- Keyword Search
- Standards:
 - HTML, XML
- Databases



Web 3.0:

- Semantically linked Documents
- Semantic queries
- Standards:
 - RDF, SPARQL
- Graphs



Web 3.0 allows...

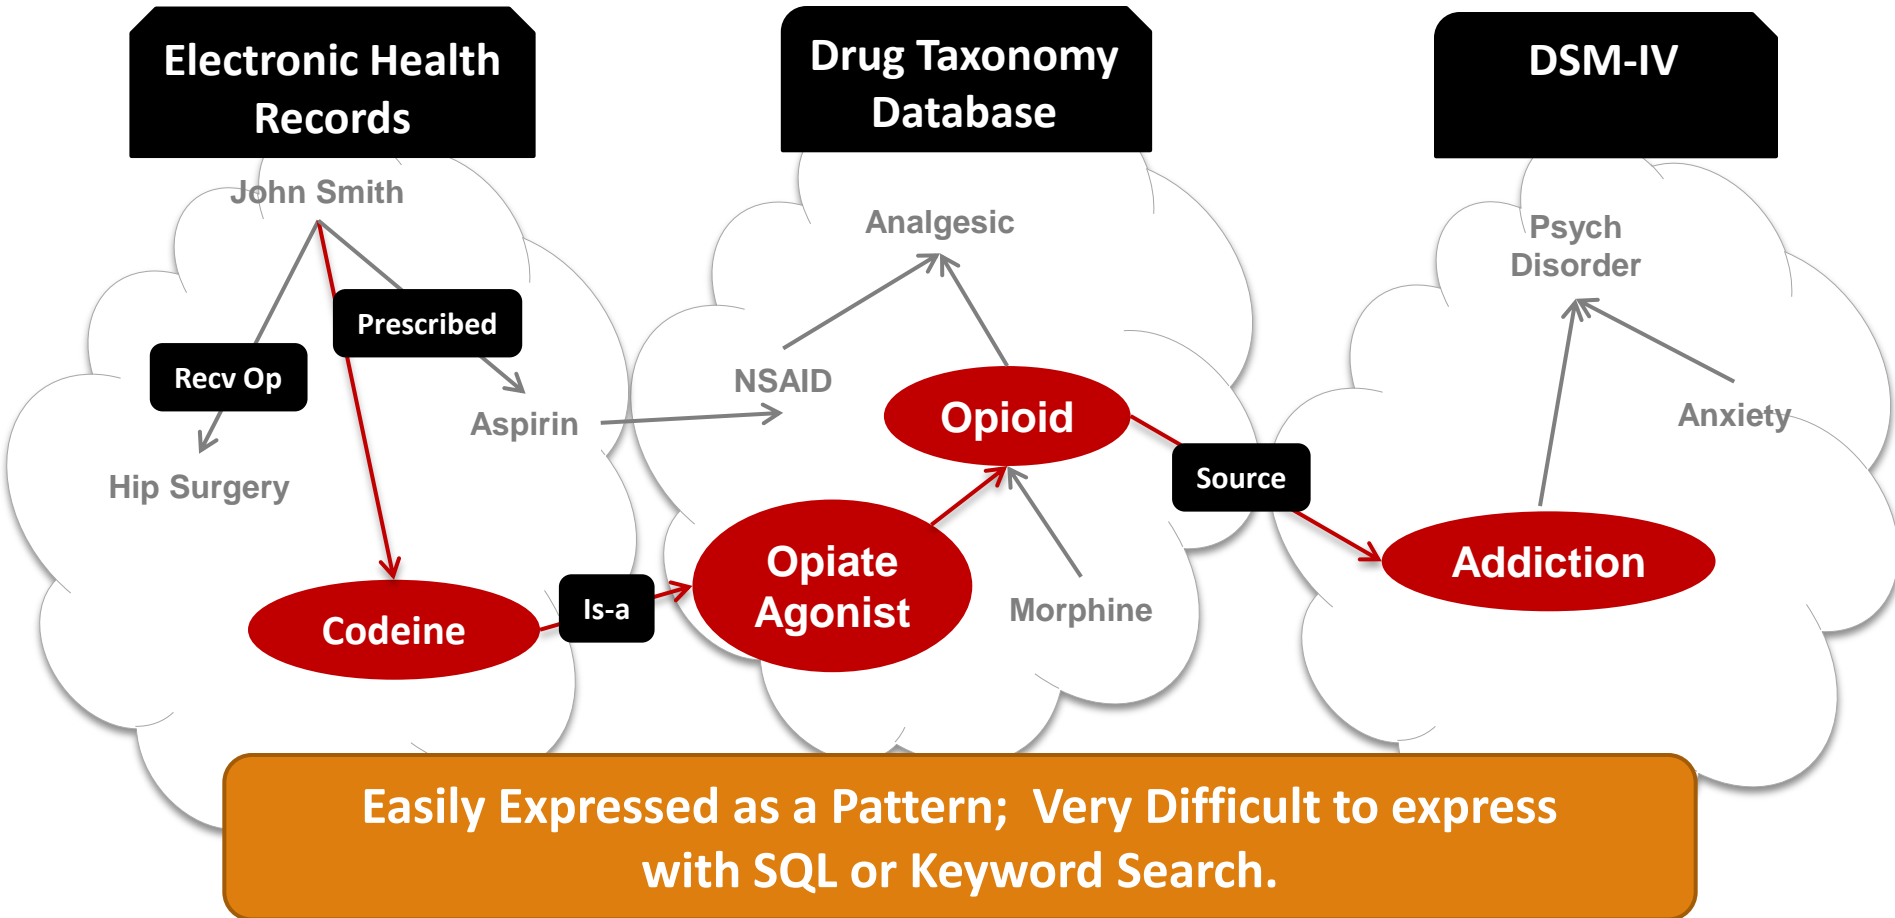
Merging
data
sources

Pattern
based
queries

Much More
Meaningful
Results

An Example from the Health Industry...

Query: What drugs are causing post-operative addictions in Hip Surgery patients?



Analytics Today

Big Data

- Structured
- Semi-structured
- Unstructured

Scale-out Analytics

- Unstructured Data
- Partitionable Datasets
- Keyword Search
- Example: Hadoop, MapReduce

In-memory Analytics

- Structured AND Unstructured Data
- Non-partitionable
- Complex Queries (“Pattern Matching”)
- Example: YarcData uRiKA

Data Warehouses BI Tools

- Structured Data
- OLAP Cubes
- Regular Queries, Known Variables
- Example: Oracle Exadata

Agenda



Big Data Analysis

The Cray XMT2

The uRIKA Graph Analytics Appliance

Customer Use Cases

Summary


Cray Is Supercomputing

We build the world's largest and fastest supercomputers for the highest end of the HPC market



Targeting the growing capability needs of government agencies, research institutions and large enterprises


We help solve the “Grand Challenges” in science and engineering that require supercomputing



Earth Sciences
CLIMATE CHANGE &
EARTHQUAKE PREDICTION




Life Sciences
PERSONALIZED MEDICINE




Defense
AIRCRAFT DESIGN



Computer-Aided Engineering
CRASH SIMULATION

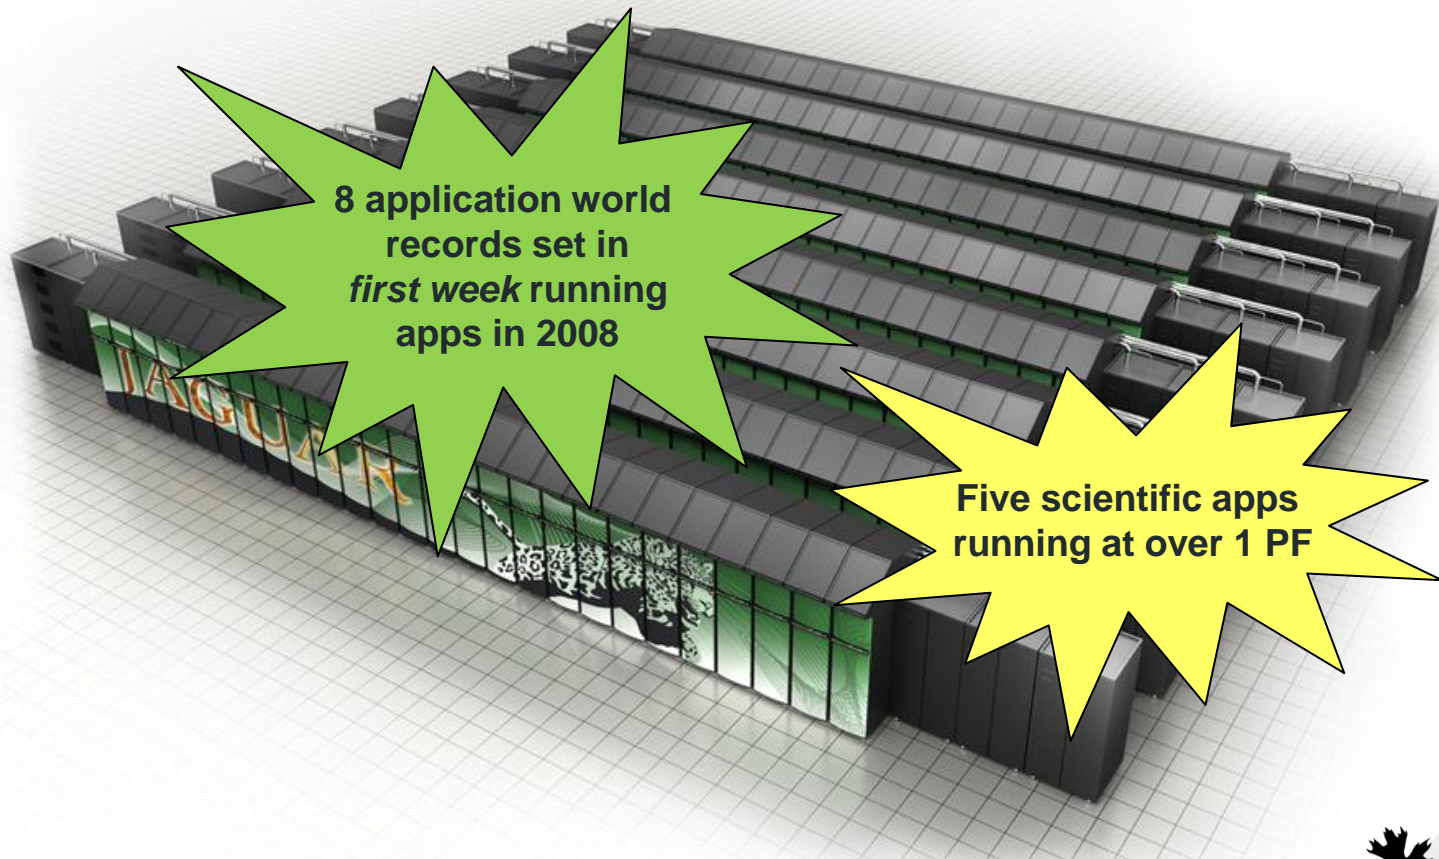


National Security
THREAT PREDICTION



Scientific Research
NEW ENERGY SOURCES &
NANOFUEL DEVELOPMENT

Sustained Petaflop Computing



8 application world records set in *first week* running apps in 2008

Five scientific apps running at over 1 PF

Cray has four deployed PF systems (out of 10 worldwide, Nov'10)
Cray remains (for over 2 years) the *only* system that has sustained an application PF



Standard Processors Can Constrain Algorithm Programming

To achieve high performance, you must...

- Place data near computation
- Access data in order and reuse data
- Partition program into independent, balanced computations (load balancing)
- Minimize synchronization and communication operations
- Avoid modifying shared data
- Avoid adaptive and dynamic computations

But what if your algorithm or application can't take advantage of these techniques?



What if the Application...

- **References very large data sets with very little locality?**
 - Caches don't work
 - Communication overhead can be overwhelming in clusters
 - Even in shared memory machines, translation hardware falls over
- **Has abundant thread-level parallelism, but very little concurrency per thread?**
 - Access pattern is data dependent
 - No computation to hide latency
 - ⇒ Threads spend most of their time waiting on global memory refs

You need a machine that...

....can efficiently reference into a large, shared, global memory

....and can tolerate long memory latencies without losing efficiency.


This motivates the design of the Cray XMT



Cray XMT System

Purpose-Built for Large-Scale, Graph-Based Data Analysis

- **Specialized performance**
 - Not designed for general-purpose HPC apps
 - *Outstanding* performance on graph analytics
- **Large, globally shared memory**
 - Architecture supports up to 512 TB of memory
 - Address translation supports sparse references across entire memory
- **Massive multithreading**
 - 128 simultaneous threads per processor
 - Tolerates long global latencies
- **Network support for single-word accesses**
 - Allows high rate of global references
- **Tagged memory (full/empty bits)**
 - Efficient lightweight synchronization
- **Sophisticated runtime to manage parallelism**
 - Parallelism grows naturally from algorithms
 - Runtime manages threads and load balancing



Programming Implications

- No longer need to place data near computation
- No longer need to access data with stride one
- No longer need to partition programs into balanced computations
- No longer need to minimize communication or synchronization events
- Adaptive and dynamic methods are okay
- Graph algorithms and sparse methods are okay
- Recursion, dynamics programming, branch-and-bound, dataflow are okay



History of the XMT Architecture

- **MTA-1 (Multi Threaded Architecture) launched in 1998**
 - 18 GaAs chips per processor blade, with custom memory
- **Cray MTA-2 launched in 2002**
 - 5 CMOS chips per processor on 1 large PC board with custom DIMMS
- **Cray XMT launched 2008**
 - Processor reduced to single CMOS chip in Opteron socket
 - 4 processors per PC board, standard DIMMS
 - Cray XT network, packaging, cooling and RAS features
- **First Next Generation XMT2 delivered to CSCS in 2011**



Agenda



Big Data Analysis

The Cray XMT2


The uRIKA Graph Analytics Appliance

Customer Use Cases


Summary

Many Big Data problems are based on Graphs


Social Networking





Intelligence/Security




Telecom/Mobile



Life Sciences/Biology



Healthcare/Medicine




Supply Chain



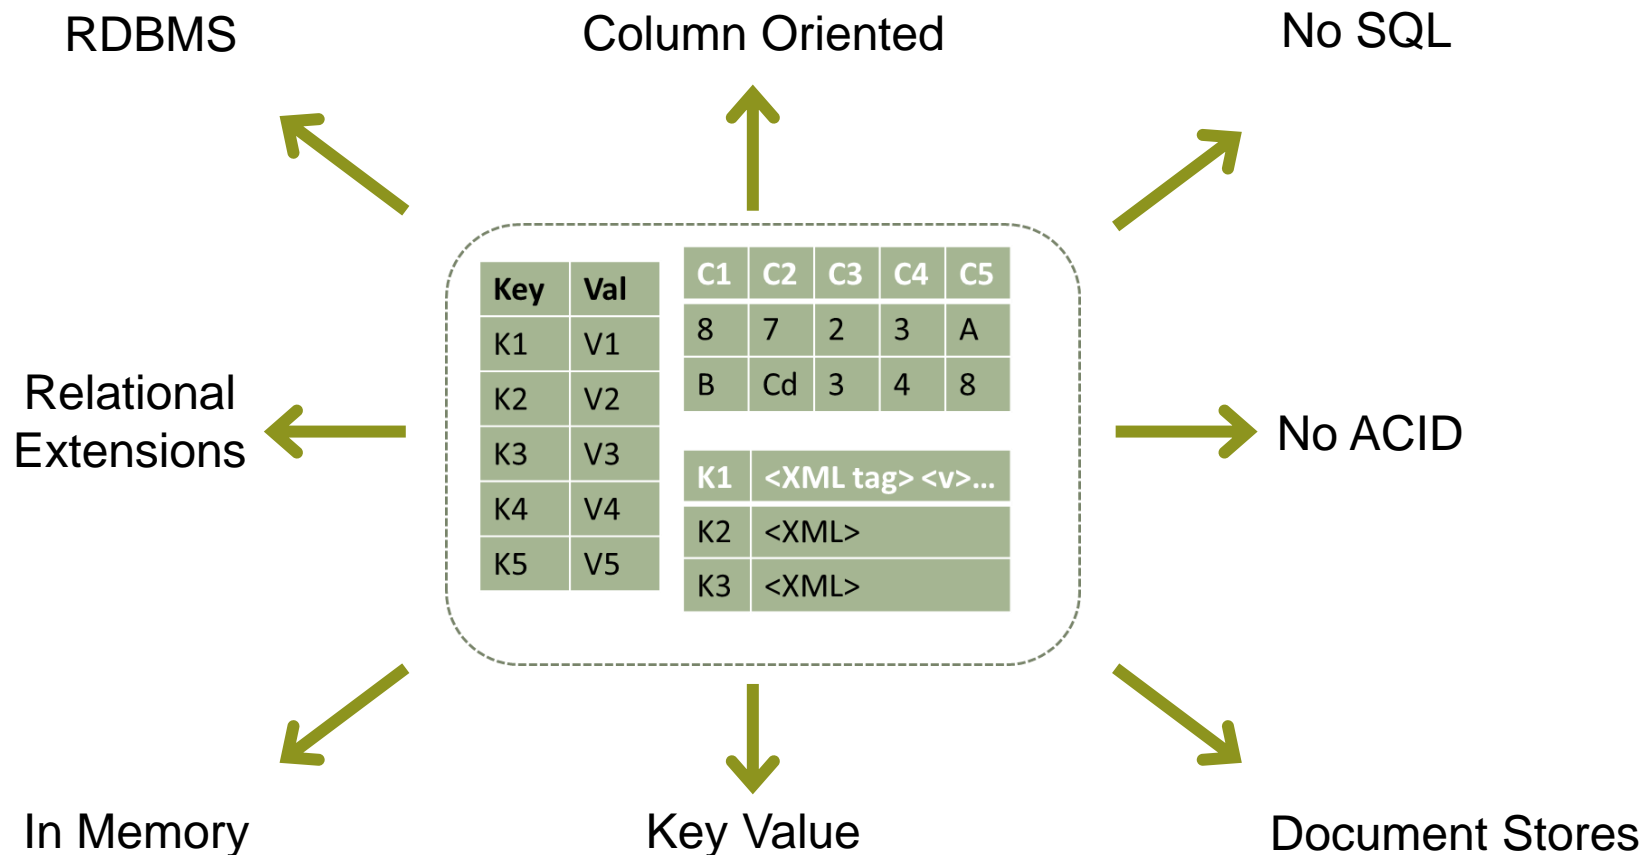
Targeted Marketing



Finance




But most Big Data solutions are based on partitioned data structures that scale out on clusters



Current Big Data approaches (including graph databases) result in low performance on Graphs...

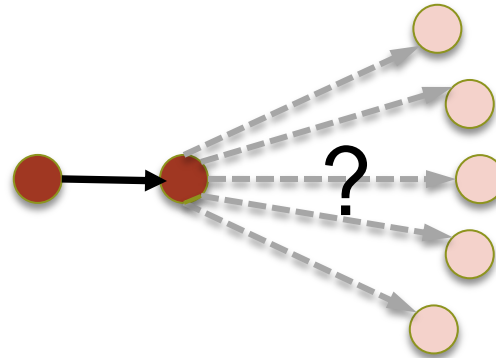
Graphs are hard to Partition



High cost to follow relationships that span Cluster Nodes

Network is 100 times SLOWER than Memory*


Graphs are not Predictable



High cost to follow multiple competing paths which cannot be pre-fetched/cached

Memory is 100 times SLOWER than Processor*

Graphs are highly Dynamic



High cost to load multiple, constantly changing datasets into in-memory graph models

Storage I/O is 1000 times SLOWER than Memory I/O*

uRiKA: Purpose-built Graph Appliance...

Graphs are hard to Partition



Large Shared Memory
Up to 512 TB

Graphs are not Predictable




Massively Multi-threaded
128 threads/processor

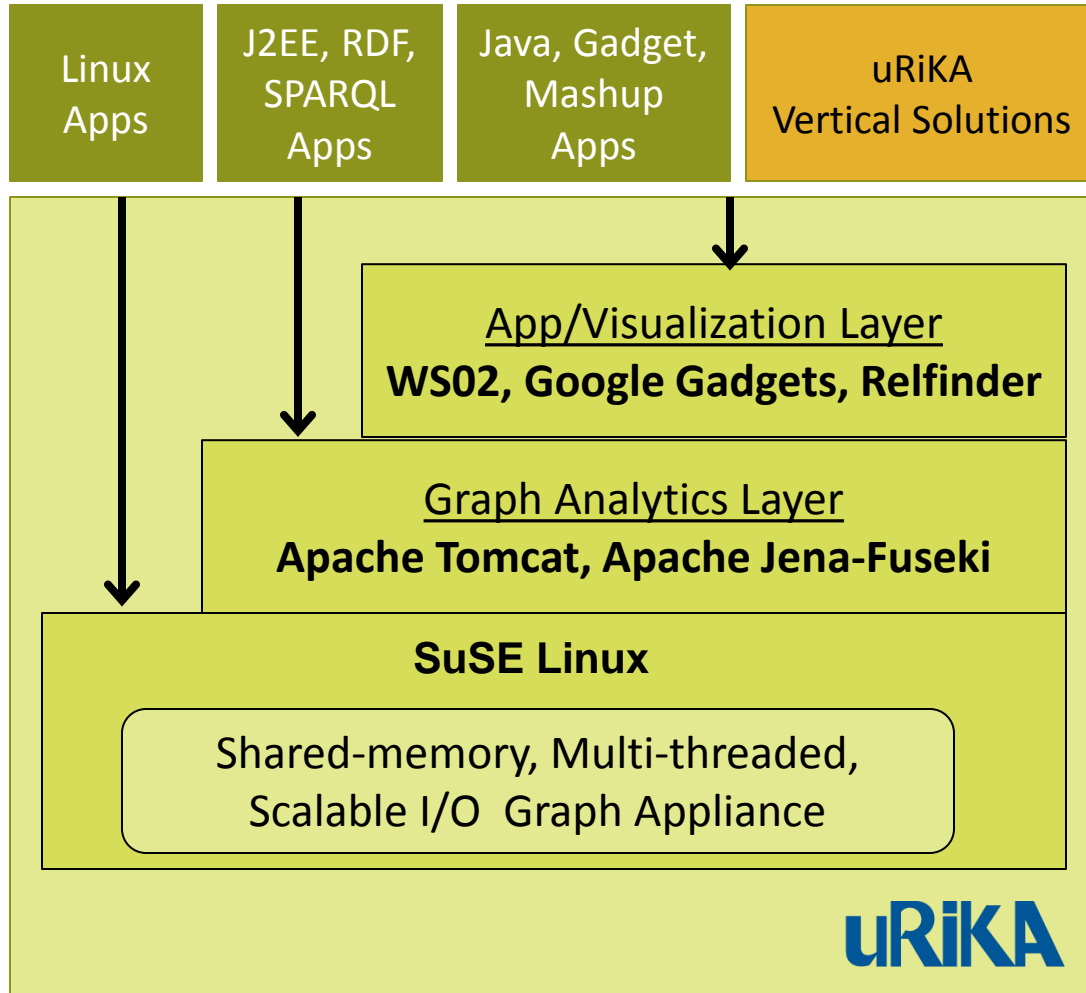
Graphs are highly Dynamic



Highly Scalable I/O
Up to 350 TB/hr



uRiKA: ... optimized for Enterprise Applications



- **Industry-standard, Open-source Software Stack**
 - Linux, Java, Apache, WS02, Gadgets, Mashups...
- **Reusable Existing Skillsets**
 - OSGI, App Server, SOA, ESB, Web toolkit...
- **No Lock-in**
 - All applications and artifacts built on uRiKA can be run on other platforms
- **Subscription Pricing model**

Agenda



Big Data Analysis

The Cray XMT2


The uRIKA Graph Analytics Appliance

Customer Use Cases


Summary

The uRiKA Moment: Discovery of Unknown/Hidden Relationships in Big Data


1. Aggregate data and relationships from multiple sources



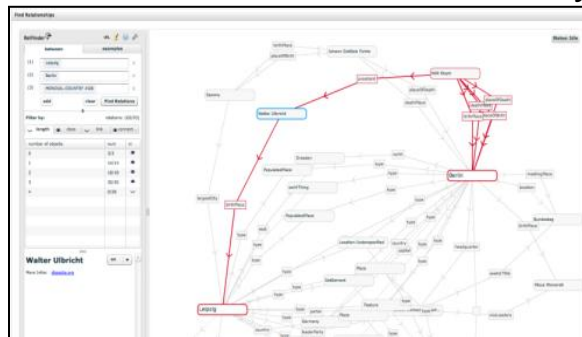
3. Build a Dynamic Relationship Warehouse



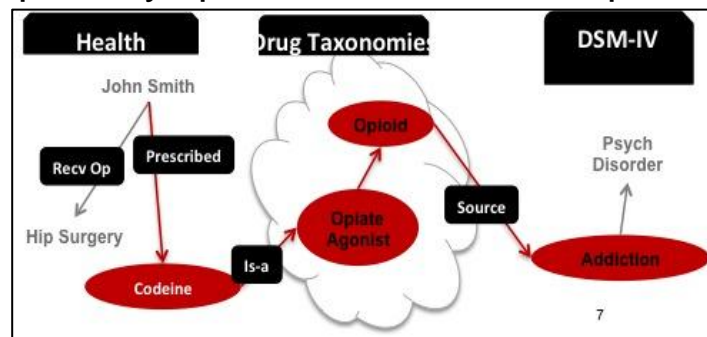
2. Augment Relationships through automated inference and deduction



Visualize relationships for real time, interactive Discovery



Search for relationships based on partially specified Patterns/Templates



uRiKA Customer Use Case: Government Organization

“Connecting the dots” to identify Persons of Interest

The Challenge

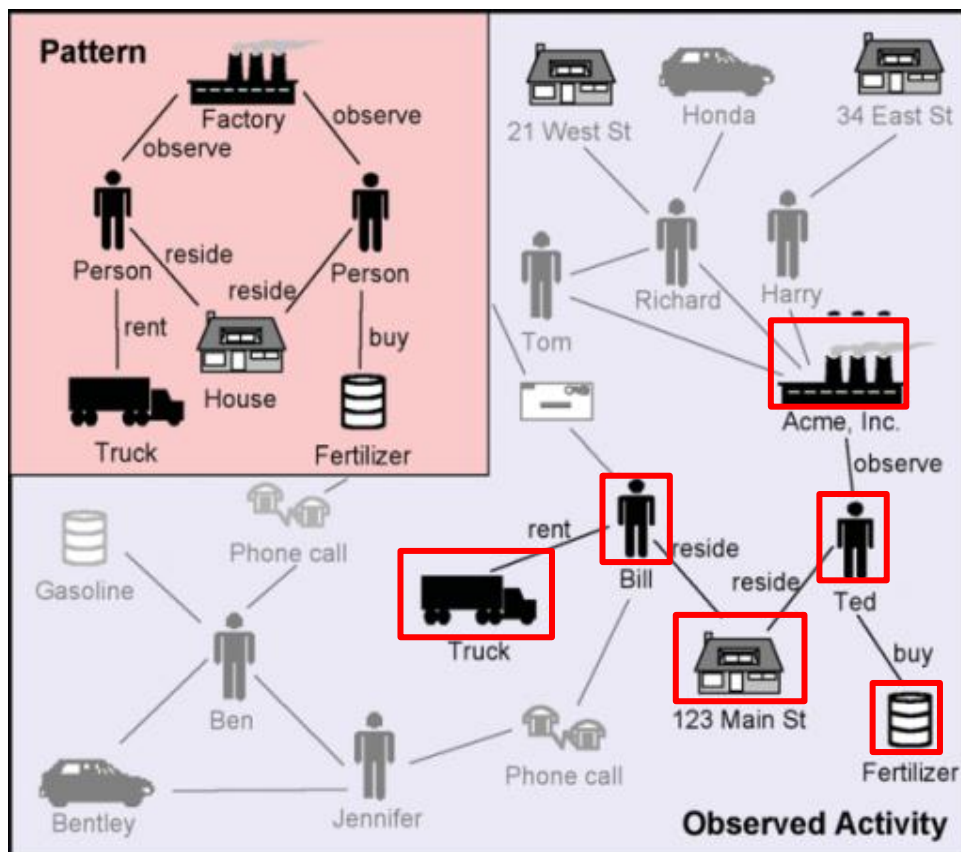
- Massive data stores of multiple data types from multiple sources
- Inaccurate, Incomplete and Falsified data
- Continuous stream of incoming data

uRiKA Solution

- uRiKA holds entire relationship graph in memory – updated constantly
- Search for Patterns of suspicious behavior and activities
- Graphical interactive exploration of relationships between people, places, things, organizations, communications, etc.

Business Value

- Proactive identification of terrorists, criminals and plots



uRiKA Customer Use Case: Personalized Treatment

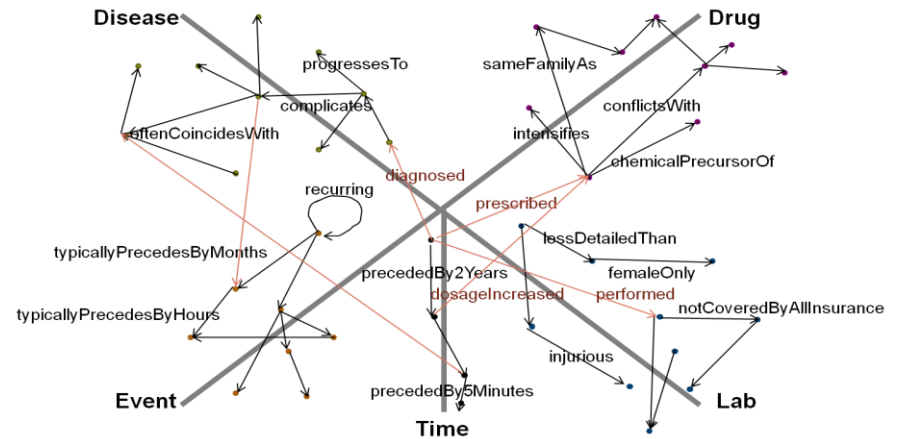
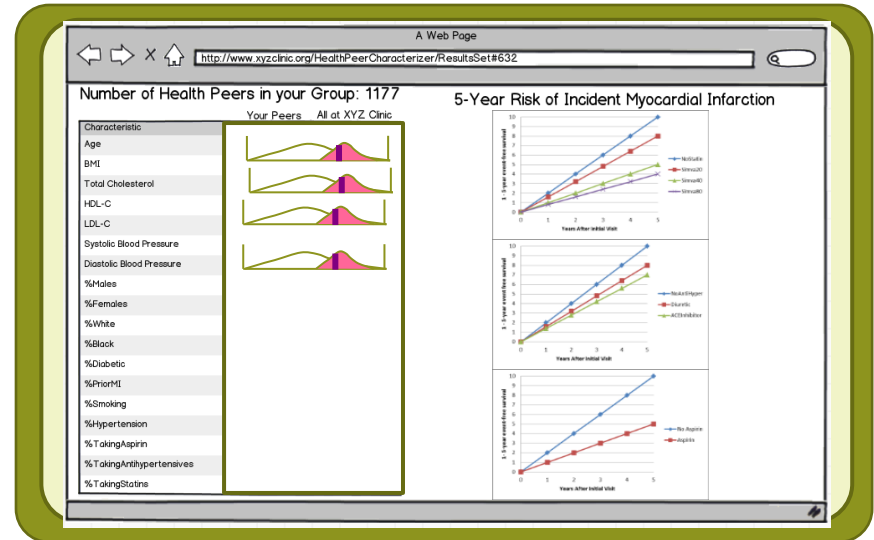
Identify “similar” Patients to optimize Treatment

The Challenge

- Longitudinal, historical data spanning all events, symptoms, diagnoses, diseases, treatments, prescriptions, etc of 10M patients including genetics and family history
- Ad-hoc, constantly changing definition of “similarity” based on thousands of parameters
- Interactive, real-time response during consultation

uRiKA Solution

- uRiKA holds entire relationship graph in memory – updated constantly
- Identify “similar” patients based on ad-hoc physician specified patterns
- Interactive, real-time access by entire physician community



uRiKA Customer Use Case: Cancer Research Institute

Integrate information across species, tumor types, sub-specialties
to see fuller picture of cancer

The Challenge

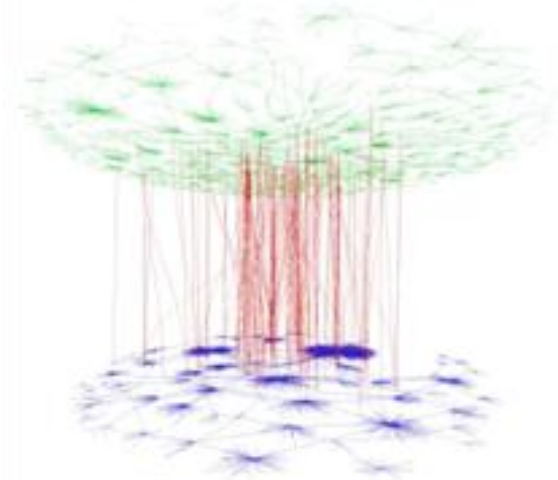
- Multiple massive datasets describing biological network graphs in cancer cells from published literature and experimental data, constantly updated
- Non-partitionable, densely and irregularly connected graphs
- Multiple researchers concurrently searching for relationships not found in published literature

uRiKA Solution


- uRiKA holds un-partitioned fused cell network graph in memory, combined with data from Medline
- Contrast experimental models and theories with published results to discover previously unknown relationships
- Interactive, real time access by multiple researchers

Business Value

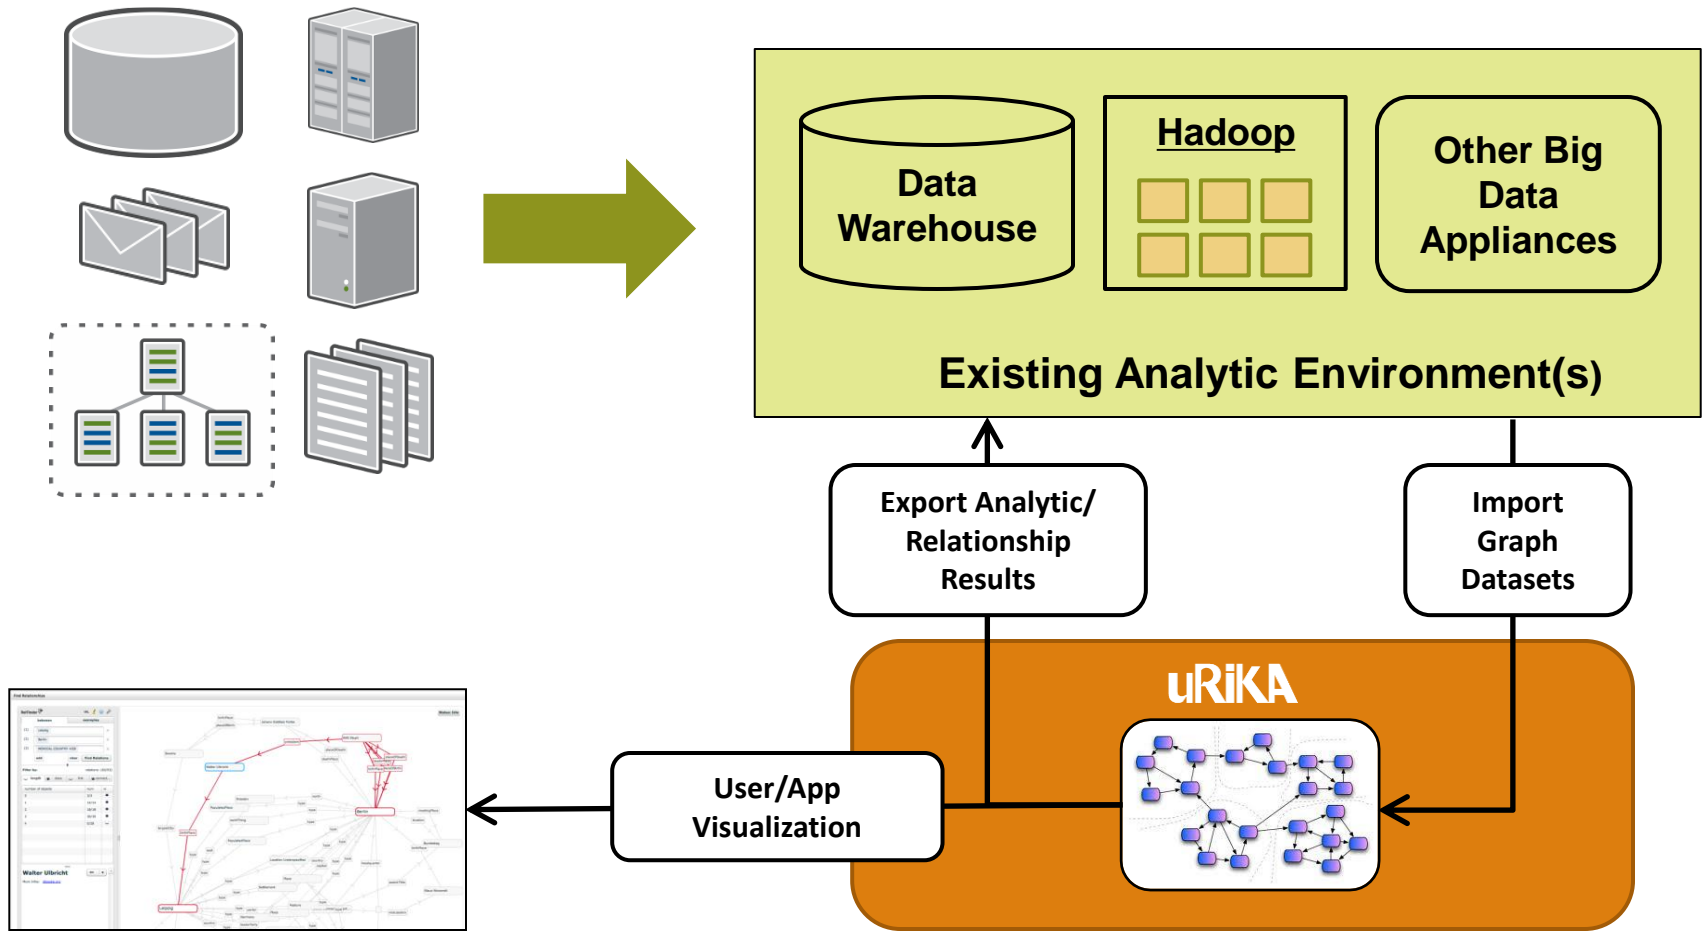
- Identify new pathways in cell models to refine cancer treatments



Confirmation of elevated VEGF levels by tissue microarray:



uRIKA complements existing Data Warehouse/Hadoop environment by offloading Graph Analytics



Agenda



Big Data Analysis


The Cray XMT2

The uRIKA Graph Analytics Appliance

Customer Use Cases

Summary

uRIKA: Big Data Graph Appliance for Relationship Analytics



Discover Unknown and Hidden Relationships in Big Data

- *Relationship Warehouse supporting Inferencing/Deduction, Pattern-based queries and Intuitive Visualization*

Perform Real-time Analytics on Big Data Graphs

- *High-performance, Graph Appliance with large shared-memory, massive multi-threading and scalable I/O*

Realize Rapid Time to Value on Big Data Solutions

- *Ease of Enterprise adoption with industry-standards, open-source software stack enabling reuse of existing skillsets and no lock-in*

Thank you!

James D. Maltby, Ph.D
jmaltby@yarcdata.com