ORACLE®

Oracle Labs : Leading the way in Hardware Software co-design

Nipun Agarwal, Senior Director, Oracle Labs

The following is intended to outline our general research direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions.
The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Our Technology Changes the World

Developed in Oracle Labs,
transferred to Oracle products and to the outside world

Agenda

- Hardware-software codesign : our view
- Power efficiency in computation
- RAPID – research project
- SPOT

Hardware Software co-design

- Long history of *leveraging* new hardware capabilities
- Last several years pushing requirements to hardware – Exadata
- Sun acquisition provided significant impetus to co-designing with hardware
- A specific instance of heterogeneous computing

In large enterprises, organization often determines architecture

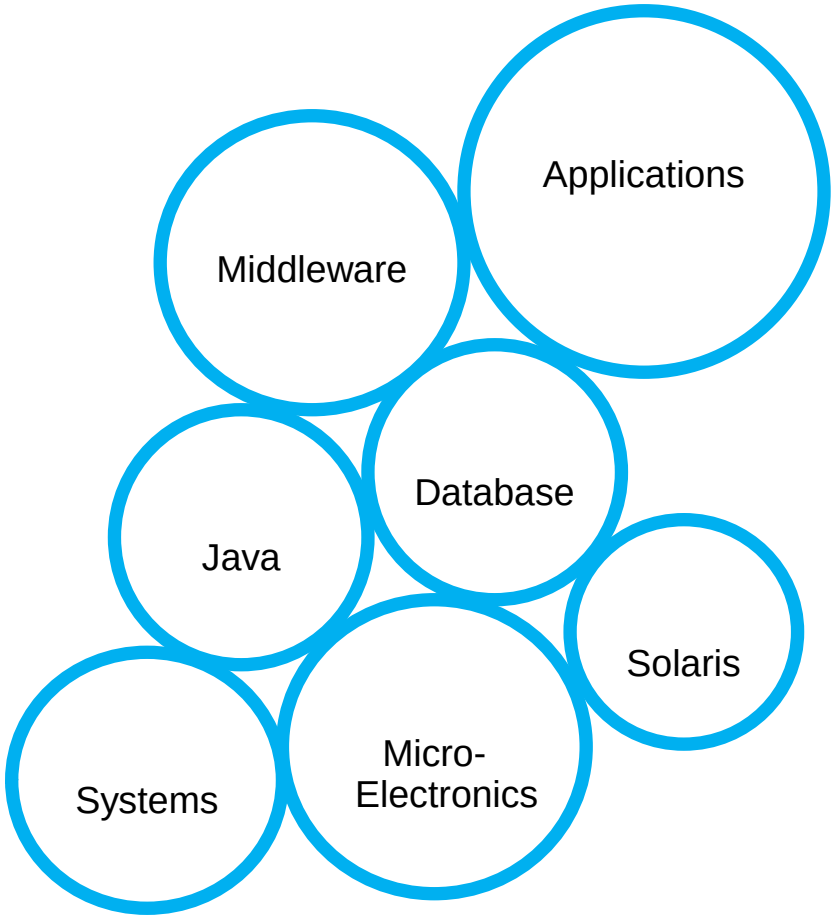

Different product areas need different kinds of support: building user interface is very different than building servers

Different product areas often have different ways of measuring success

Academic research has similar problems

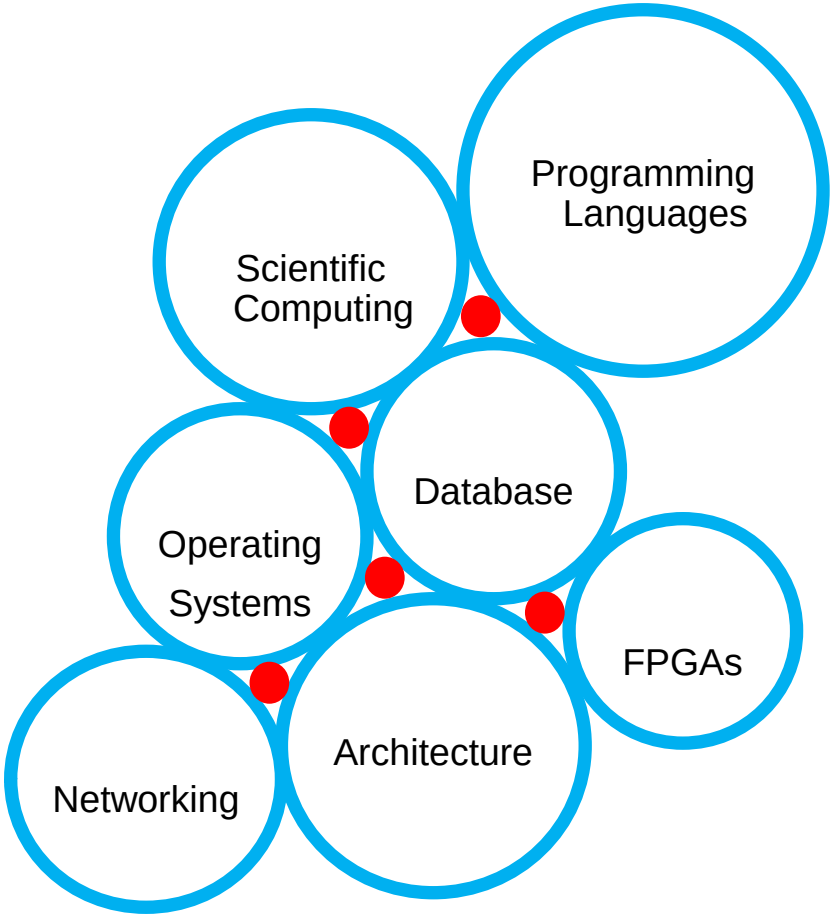

Faculty members & students with a track record & contacts in one community find it difficult to publish if they stray far from the core interests of the conferences serving that community

Oracle Labs is making an extra effort to start and fund cross-domain research

Heterogeneous Systems

- Already deployed in many areas:
  - Mobile devices
  - CPU + GPU in laptops and desktops
  - Java + SQL software stack in the data center
- They provide the opportunity to move to:
  - Lower-power hardware
  - More parallel software
- Heterogeneity allows us make these changes incrementally to the existing technology stack

Heterogeneous Hardware Challenge - Creates Software Complexity That Requires Management

- Programming a system with both GPUs & CPUs is extremely difficult
  - Especially optimizing the workload across the devices
  - Would be nice to compile for this automatically
  - A number of academic projects in this area
- Deciding what parts of the workload is appropriate for HW accelerators is very difficult
  - Consider an encryption algorithm written in C/Java
  - Could a compiler figure out from the C code that the implementation is doing AES and use a HW implementation?
- Higher-level languages that capture more of the application semantics are very helpful

Examples of hardware software co-design projects in Oracle Labs

- RAPID – Rapid Analytic Processing In DRAM
  - Lots of low power cores + fixed-function accelerators + lots of DRAM
  - Scans & other simple database operators pushed onto accelerators

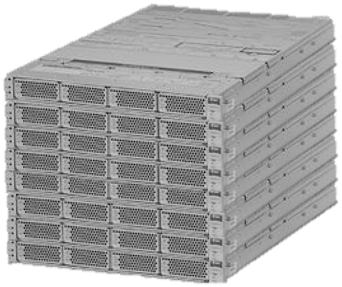

- Sun SPOTs – Sensor Platform
  - "Heavyweight" ARM9 processor: 180MHz 32-bit ARM920T
  - Lightweight 8-bit microcontroller: Atmel Atmega88 w/ 8K Flash
  - Microcontroller wakes up the ARM core when sensor data is interesting

Agenda

- Hardware-software codesign : our view
- Power efficiency in computation
- RAPID – research project
- SPOT

Power efficiency - where does the power go?

- We could make high performance compute + storage much smaller if we could power and cool it, so either:
  - Build more power plants + do advanced (i.e. liquid) cooling
  - Make computers more power-efficient
- Need to address power in compute, memory & IO subsystems
- Power considerations force more parallelism

Saving power requires trade-offs: Pick two

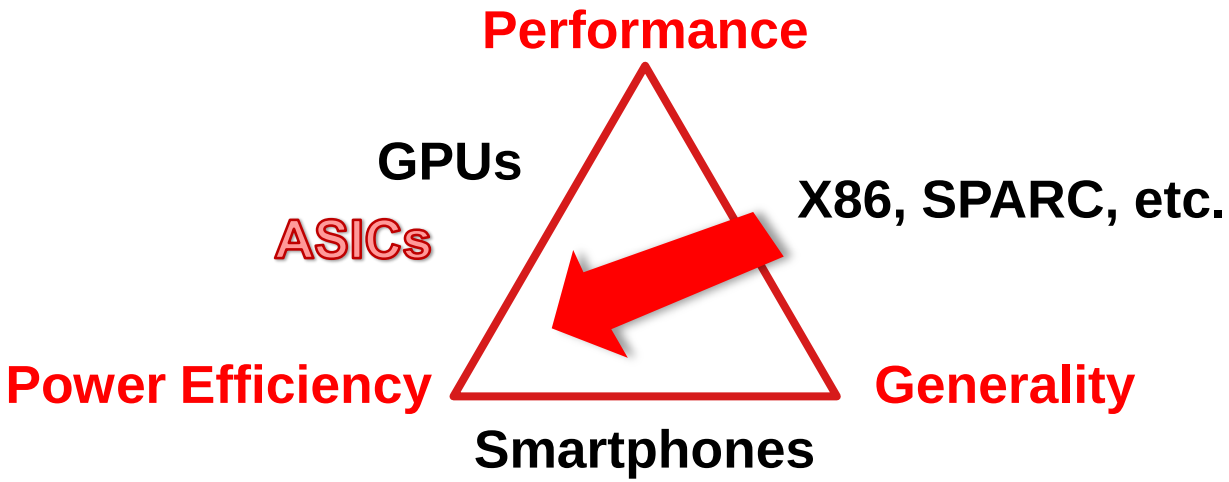

- GPUs optimize for orders of magnitude more compute
  - Bandwidth not improved much over X86 CPUs
- Oracle Labs wants to optimize for orders of magnitude more bandwidth per unit power than CPUs
  - Database applications prefer bandwidth over arithmetic

Big power savings in the processor with application-specific hardware

- H.264 encode study

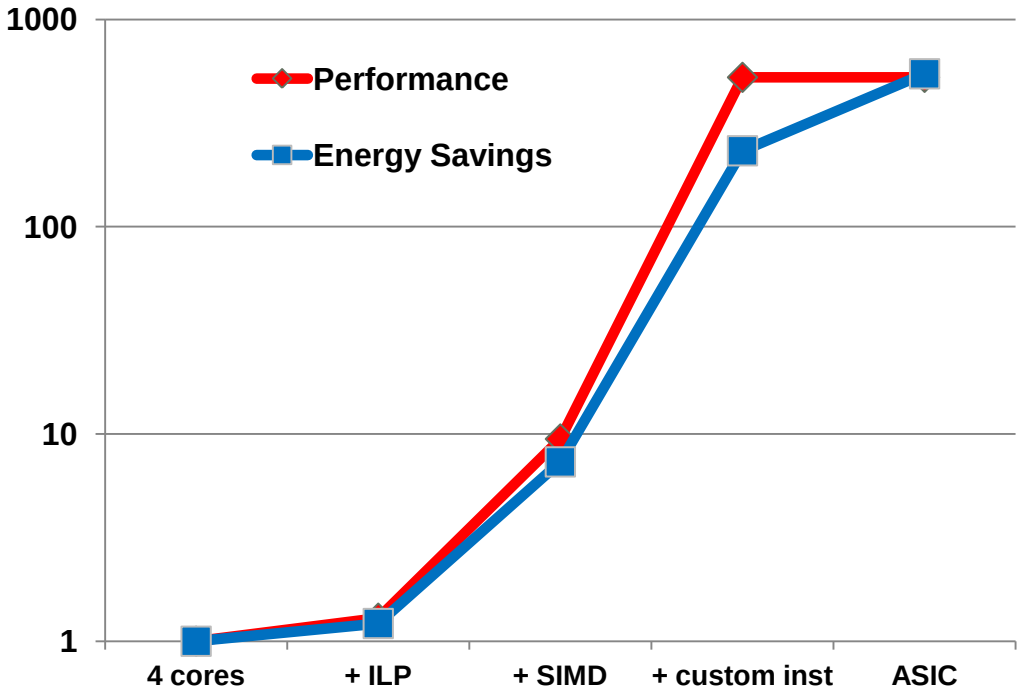

Horowitz et al. *Understanding Sources of Inefficiency in General-Purpose Chips* (ISCA 2010)

- Anton molecular dynamics computer

- **400 MHz, 100x power savings**
- **1000x performance improvement**
- **SC 2009 Best Paper**

Software Development Has Similar Tradeoffs Around Generality

- Why can't I have it all?
  - Restrictions on language semantics allow many more optimizations
  - General, committee-designed languages accumulate a kitchen sink of features with bizarre interactions making optimization harder
    - Programming at scale requires "coding standards" to outlaw sets of features and simplify style in general languages

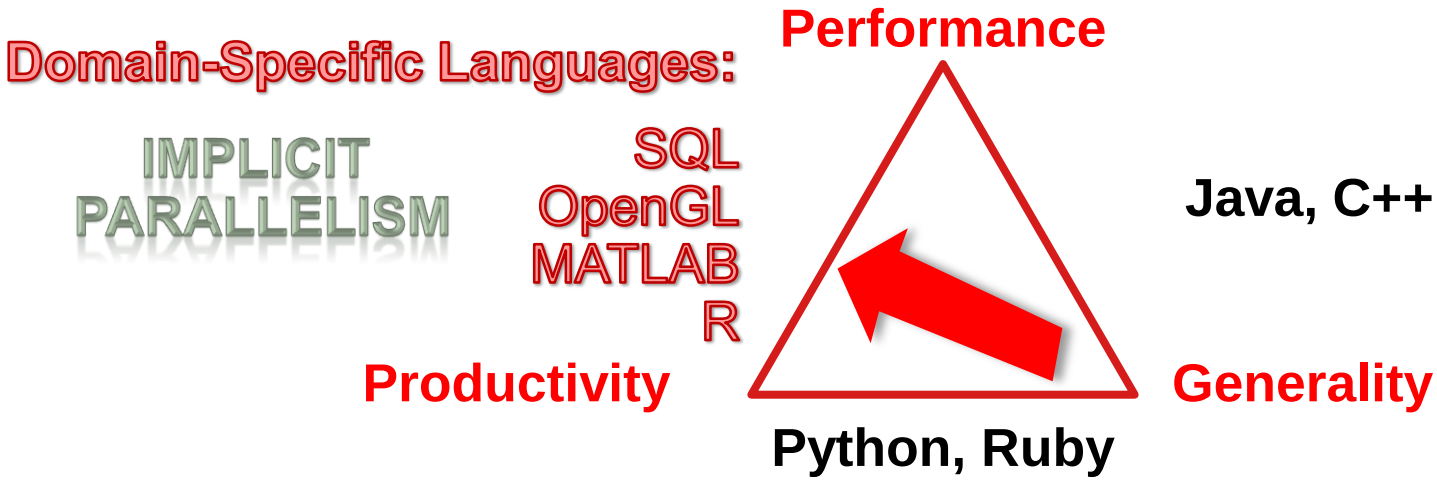

Agenda

- Hardware-software codesign : our view
- Power efficiency in computation
- RAPID – research project
- SPOT

Case Study on Power: Analytics on "Big Data"

- Hot technology space
  - Analytic "Scan" Appliances
  - Map/Reduce workloads
- What are the limits?
  1. Power
  2. Space
  3. Need new software
  4. Cost
- Example problem:
  - Scan 10 physical TB in < 1s (40 TB compressed)
  - How much HW is needed?

Commodity Hardware

- Choose most *power-efficient* Xeon configuration that can solve our problem
- State of the art extrapolated
  - 288GB DRAM
  - 100GB/sec peak DRAM bandwidth
  - ⇒ Scan performance is DRAM-bandwidth bound
- For our problem: scan 10TB in 1 second

  Need 120 of these servers

  … that's 3 full racks of servers

  … and minimally 30kW of power

RAPID Database Accelerator is an interesting exercise in trading off generality issues

- RAPID is a less general hardware platform in that:
  - Low single thread performance
    - Not going to be good at running most existing code
    - Requires architecture-specific code
  - Bad at chasing pointers across large memory heaps
  - Limited interconnect
    - Poor support for inter-module communication
    - Requires high selectivity in RAPID modules
- RAPID helps provide a more general DB experience
  - Depending on indexes for query performance requires casting query patterns in stone ahead of time
  - Scans are more general than using indexes
  - Table scans work for any F(row) a user might decide to do!

Agenda

- Hardware-software codesign : our view
- Power efficiency in computation
- RAPID – research project
- SPOT

SPOT Device

- Sensor Platform for Oracle Technology
- Device has three layers
  - Sensor Board (light, humidity, temperature, shock)
  - Processor board, mesh communication
  - Battery, storage
- User programs the device entirely in Java
- 25,000 devices being used

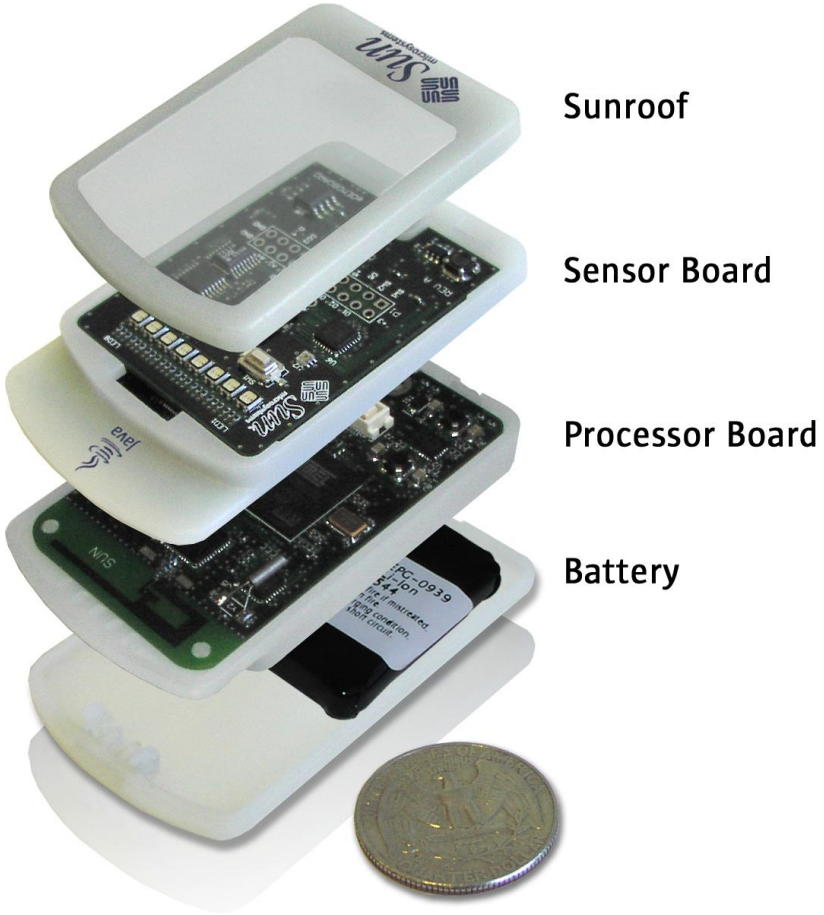

SPOT Capabilities

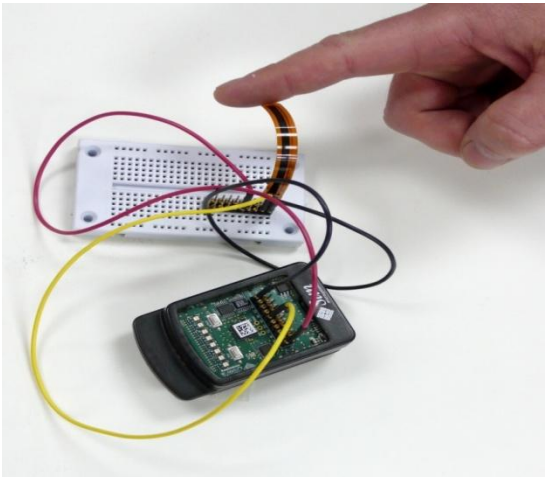

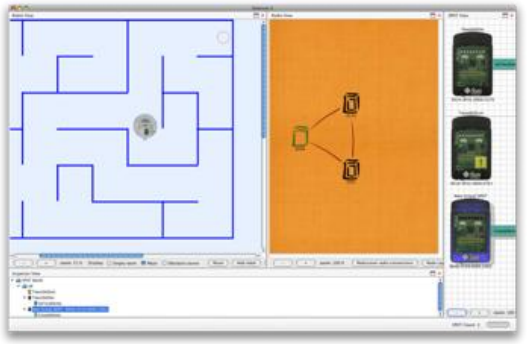

- Embedded Development Platform
- Easy to program - Java top to bottom
- Connected - Wireless Communication
  - Mesh Networking
  - Over the Air Programming
- Mobile
- Aware and Active
- Secure
  - Built-in high grade ECC public key cryptography
- Capable of remote software deployment and re-configuration

Oracle's Enterprise Sensor Platform

- Integrated wireless and satellite communicatior
  - Combines redundant GPRS and satellite services
  - GPS location as close as 12 feet
- Advanced Intrusion Detection
  - Specialized door and motion sensors
  - 6-wall sensor using ultra-wideband radar
- Option for adding other sensors
  - It's a platform
  - Chemical, radiation, explosives, etc.
- Sensor Fusion
  - On-board processing distills sensor data
    - Bayesian network-based probability functions provide self-calibrating feedback to reduce false alarms

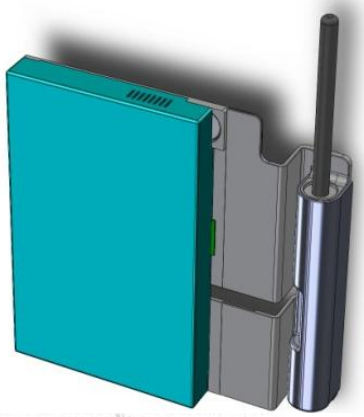

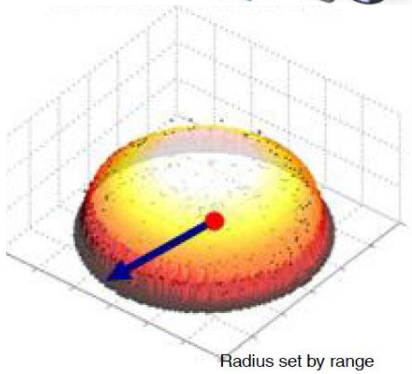

Sensor Fusion Architecture

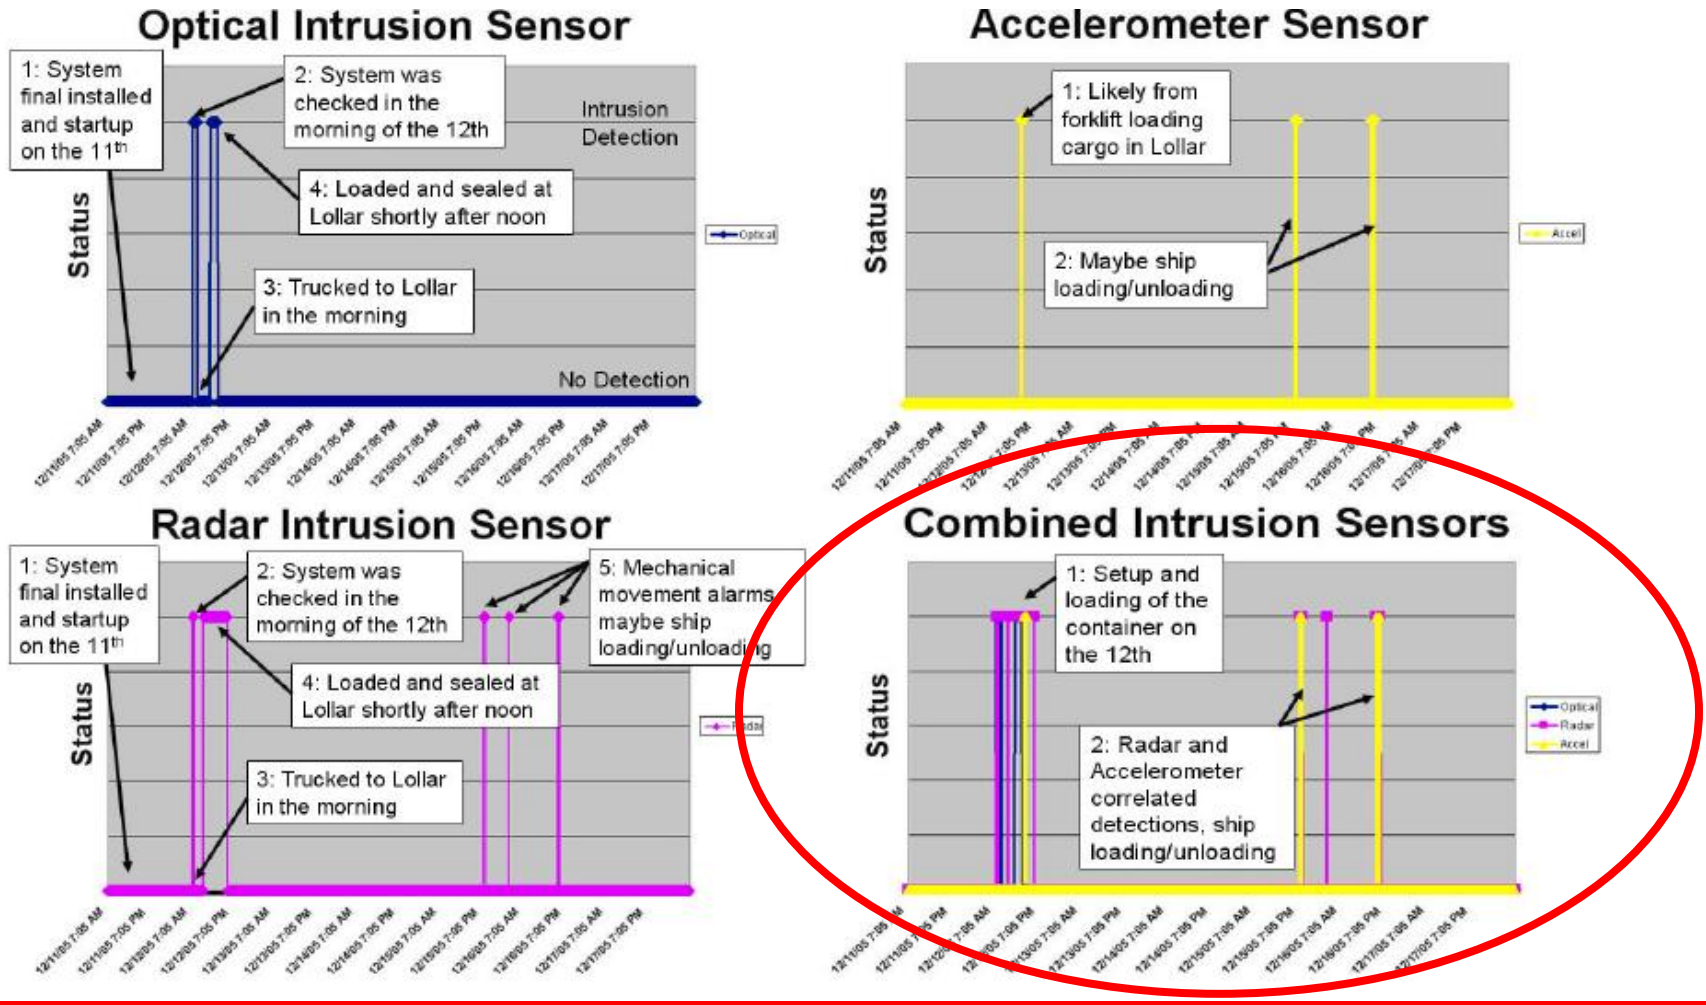

Tracking & Monitoring Portal

Refresh | Clear

Container Monitoring Dashboard

ORACLE

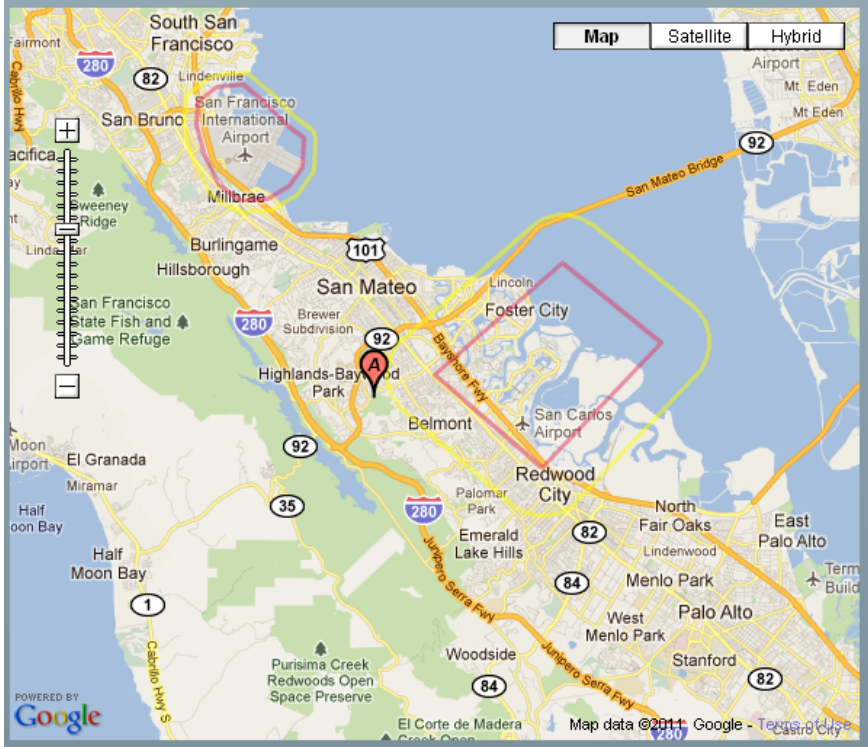

Alert Messages

SourceID	Alert Time	Message
A	23:22:01 on 28,Jun,2011	Temperature threshold exceeded - possible cargo spoilage!
A	23:14:01 on 28,Jun,2011	Container tilted - possible cargo damage!
A	23:47:01 on 28,Jun,2011	Container has been opened - possible cargo theft!
A	23:22:01 on 28,Jun,2011	Temperature threshold exceeded - possible cargo spoilage!

Geofence Monitoring Alerts

SourceID	Time	Area	Type	Action
A	23:14:01 on 28,Jun,2011	ORACLE	IN	Exit
A	23:13:01 on 28,Jun,2011	ORACLE	NEAR	Exit
A	23:13:01 on 28,Jun,2011	ORACLE	IN	Enter
A	23:08:01 on 28,Jun,2011	ORACLE	IN	Exit

Geofenced Areas of Interest

Id	Updated	Name	Version	Buffer	Sides	Lat	Lng
P1	21:33:28	Parliment	2	100	19	51.5009	-0.1261
SFO	18:20:18	SFO Airport	1	500	14	37.6372	-122.4052
ORACLE	00:37:35	Oracle HQ	1	2000	4	37.5350	-122.2932
L25	16:51:35	L25	1	250	4	51.5156	-0.1000
L56	11:22:01	L56	1	100	4	51.5072	-0.1295

Container Monitoring Updates

SourceID	MessageID	Time	Latitude	Longitude	Luminance	Temperature	Accelerometer			Battery
							X	Y	Z	
A	510	23:23:01 on 28,Jun,2011	37.5265	-122.3220	20	51.2500	0.0000	-0.0937	0.9844	100
A	509	23:22:01 on 28,Jun,2011	37.5265	-122.3220	15	51.0500	0.1719	0.2031	1.0469	100
A	508	23:21:01 on 28,Jun,2011	37.5265	-122.3220	3	44.5500	0.0000	-0.0937	1.0156	100
A	507	23:20:01 on 28,Jun,2011	37.5265	-122.3220	3	36.5500	-0.0156	-0.0781	1.0313	100
A	506	23:19:01 on 28,Jun,2011	37.5265	-122.3220	5	27.4500	-0.0156	-0.0625	1.0156	100

ORACLE

Other Research Projects in Oracle Labs

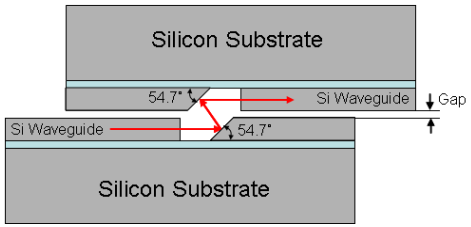

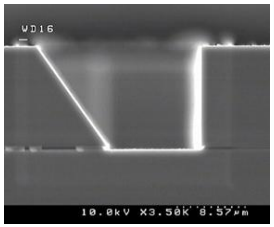

OPxC optical

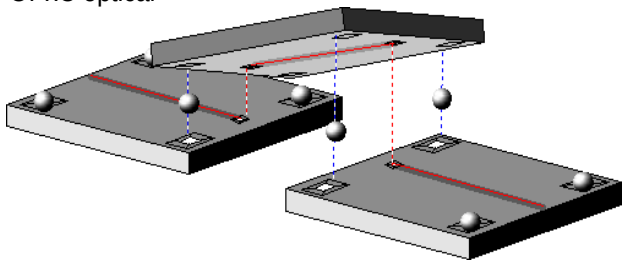

- Silicon photonics
  - $50m DARPA project “UNIC” to build on-chip and system-level interconnect using optical waveguides: very low power (>10x) interconnect
- Program analysis
  - Look at source code to find bugs, security problems, concurrency issues, etc. (static & dynamic approaches)
- Titan – Enterprise “Bit-Torrent”-style provisioning
- Machine learning & information retrieval
- Adaptive optimization
- Large-scale system simulations

Oracle Labs beyond Oracle

- Oracle Labs **extends beyond Oracle** through the Labs' "External Research Organization" or **ERO**
- External Research Collaborations
  - Involve active participation with Oracle Principal Investigator
    - Test for Oracle business relevance
  - Have defined deliverables, timeframes & staffing
  - Typically oriented around funding graduate students
- Benefits include: learning & exploration for Oracle Pis, ecosystem improvements, recruiting, tech transfer

Hardware and Software
Engineered to Work Together