# RapidMind

## Accelerating Medical Imaging

May 13, 2009

# Outline

- Medical imaging software challenges
- Case studies
  - Elastography, registration, breast cancer screening
- Platform
  - System overview
  - Detailed example: *algebraic reconstruction*
- Medical imaging solutions
  - Imaging extensions
  - Registration component
  - Component roadmap


Copyright © 2009 RapidMind Inc.

# Challenges

## Need for application acceleration

- Growing data sizes
- Mainstreaming of volume imaging
- Need for efficient workflows and interactivity
- Need for advanced volume processing increasing

## Insufficient development resources

- GPUs and multi-core CPUs are capable but difficult to program
- Hardware choices in constant flux
- High costs associated with changing development tools
- Need to complete projects faster
- Need to reduce effort spent on non-core development
- Want to focus effort on value-added features


Copyright © 2009 RapidMind Inc.

# Need for Application Acceleration

## Growing data sizes

- Toshiba 320-slice dynamic volume CT imager
- 4D perfusion and cardiac studies

## Mainstreaming of volume imaging

- Frost and Sullivan: 3D studies now standard, 4D usage increasing

## Need for efficient workflows and interactivity

- Radiologists paid by case

## Need for advanced volume processing

- Automated cancer screening
- Automated Alzheimer's screening
- Multimodality studies
- Stroke evaluation
- Real-time registration for image-guided surgery
- Real-time registration for radiation therapy


Copyright © 2009 RapidMind Inc.

# Medical Imaging and Visualization Case Studies

## University of Waterloo

•2D/3D Registration of CT and X-Ray data for image-guided surgery

***With RapidMind …***
**• 100x speed up driving time to 3 seconds**
• Image analysis is rendered real-time
• Cancer screening and diagnosis is immediate, and can be used to guide non-invasive surgeries

## Dartmouth Medical School

•Quasi-static Elastography image analysis used to diagnose breast cancer from ultra-sound data

***With RapidMind …***
**• 10x speed up driving time to 1 second**
• Image analysis is rendered real-time
• Cancer screening and diagnosis is immediate, and can be used to guide non-invasive surgeries

## RapidMind Demonstration

- Algebraic reconstruction

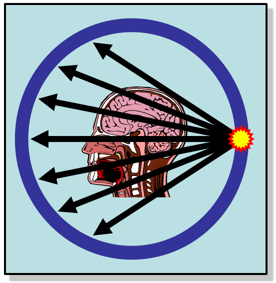

***With RapidMind …***
- **60x speedup**
- Reconstruction completed faster than data can be acquired
- Possible to use more flexible imaging geometries and obtain higher image quality than is possible with traditional cone-beam backprojection

## medipattern

•Pioneer in the development of medical software solutions that help improve workflow and productivity
•B-CAD™ V1 Computer-aided Detection of breast cancer from ultrasound images

***With RapidMind …***
**•9x speed up on multi-core CPU with GPU acceleration**
•A heterogeneous application that leverages multi-core with seamless GPU acceleration
•Increased accuracy of diagnosis


Copyright © 2009 RapidMind Inc.

## Medical Imaging Case Studies
# 2D/3D Registration

**100x**

**RapidMind implementation runs over 100 times faster than serial C++ baseline**

UNIVERSITY OF
Waterloo

- 2D/3D registration finds rotation of 3D volume so its projection aligns with given 2D image
- Combines volume rendering for projection with computation of registration metric and optimization
- Useful for contrast enhancement in image-guided surgery
- For more information see: Lin Xu and Justin Wan, *Real-Time Intensity-Based Rigid 2D-3D Medical Image Registration Using RapidMind Multi-Core Platform*.


Copyright © 2009 RapidMind Inc.

# Computer-aided Detection of Breast Cancer

**9x**

**RapidMind implementation runs 9 times faster than original design – a performance requirement for customer acceptance**

medipattern®

Medipattern's award-winning flagship B-CAD® solution helps characterize, evaluate, and document breast ultrasound images in a thorough, organized way that integrates easily into daily practice.

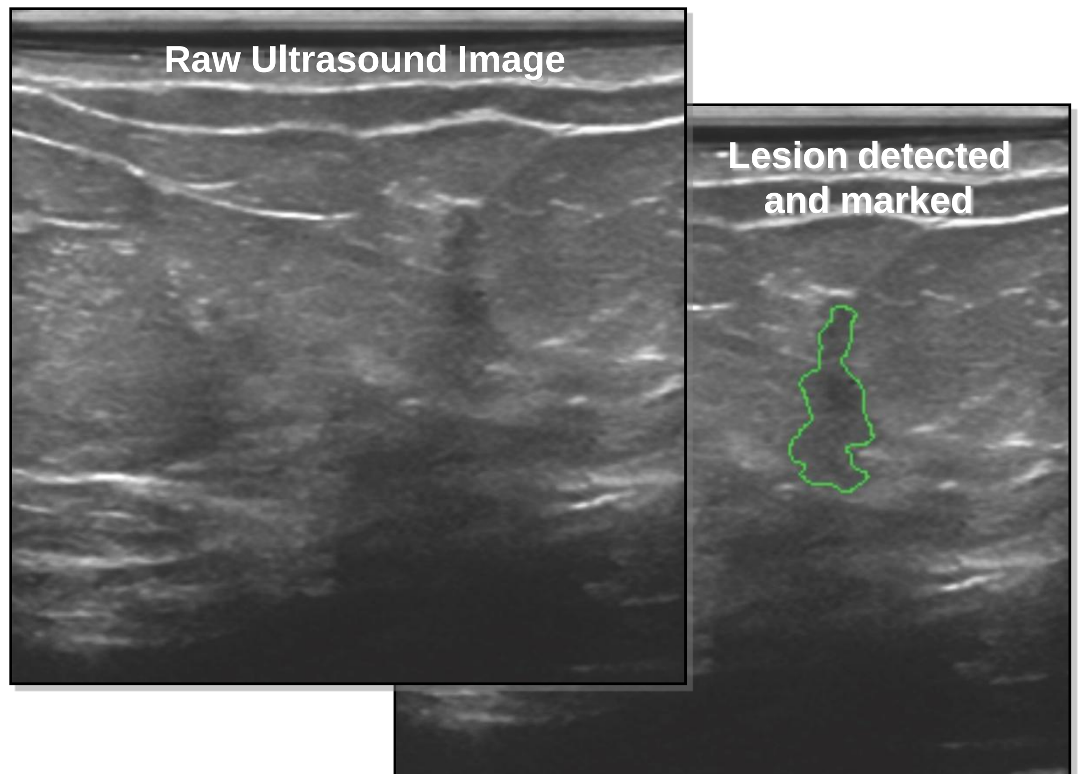



Copyright © 2009 RapidMind Inc.

# Elastography Image Analysis

**10x**

## RapidMind implementation runs 10 times faster achieving interactive user experience

### Raw Ultrasound Image

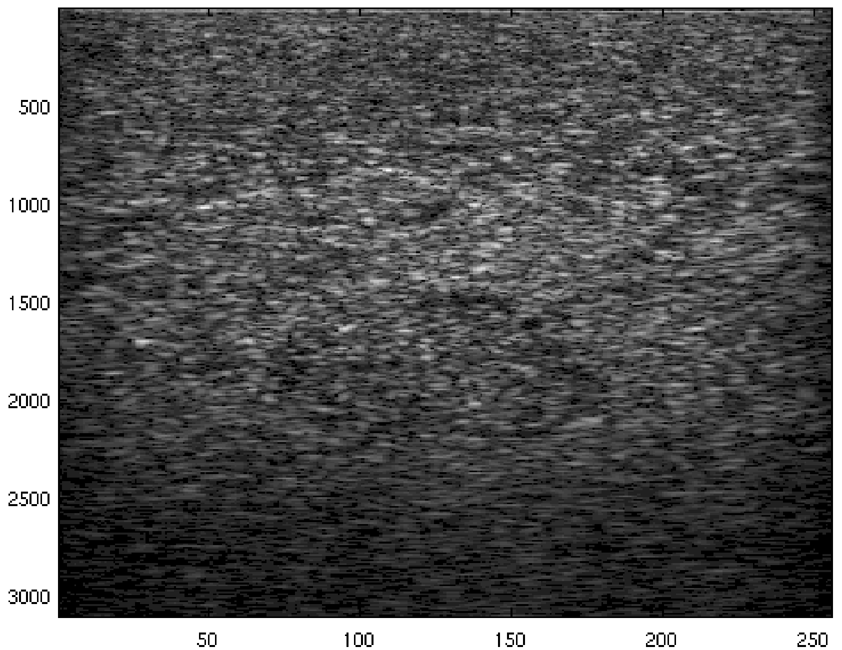


### Correlated Ultrasound Image

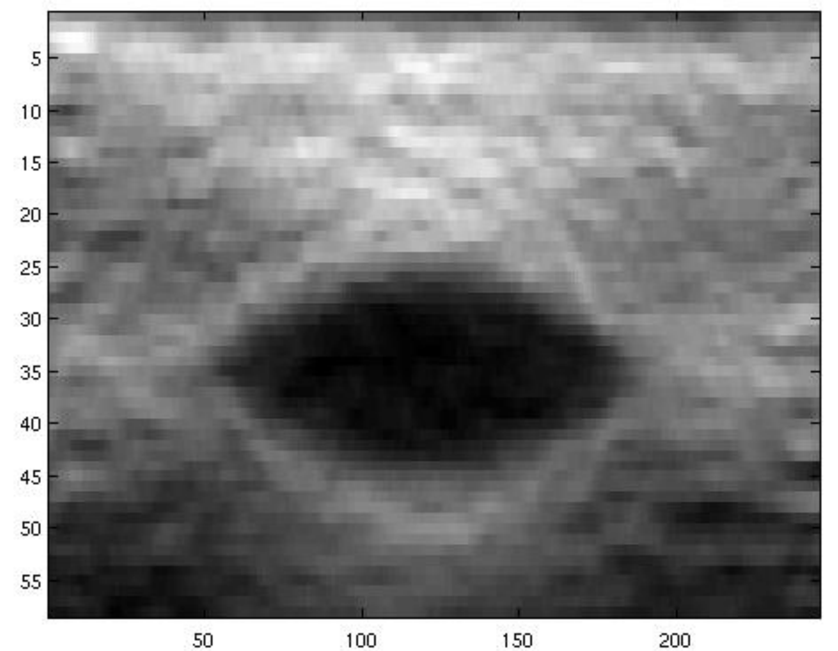


**A closer look: NVIDIA 8800 GPU**

Original MATLAB implementation: 240 seconds to render
C++ implementation: 10 seconds to render
RapidMind implementation: 1 second to render

**Source:**

Dr. Susan Swartz, Dartmouth Medical School, Dr. Marvin Doyley, Professor of Radiology and John Wallace, Academic Computing.


Copyright © 2009 RapidMind Inc.

# Algebraic Reconstruction

**60x**

## RapidMind's algebraic reconstruction implementation can reconstruct an entire volume in under 3 seconds using a single GPU

### X-rays projecting through the volume

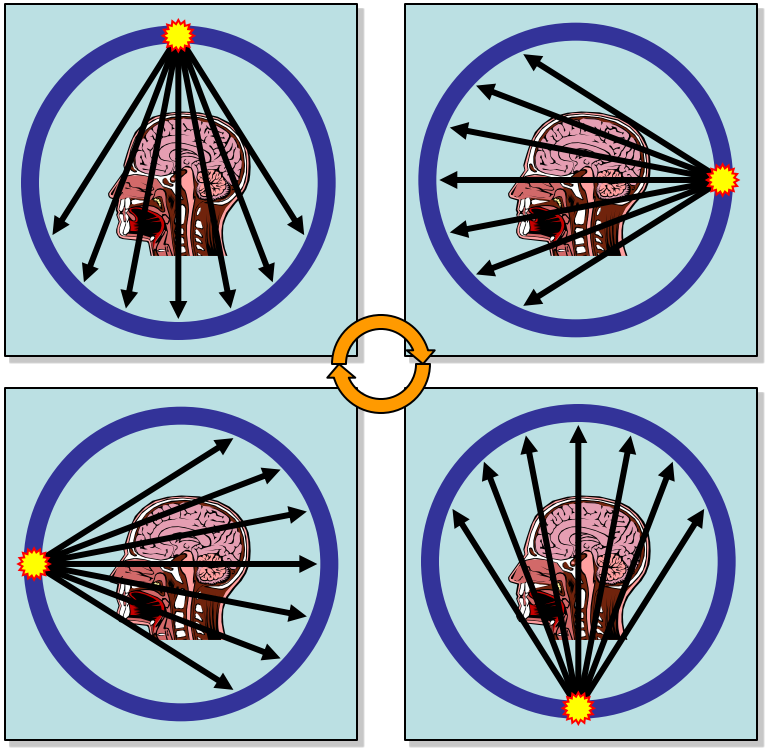

### Visualization of algebraic reconstruction of CT data

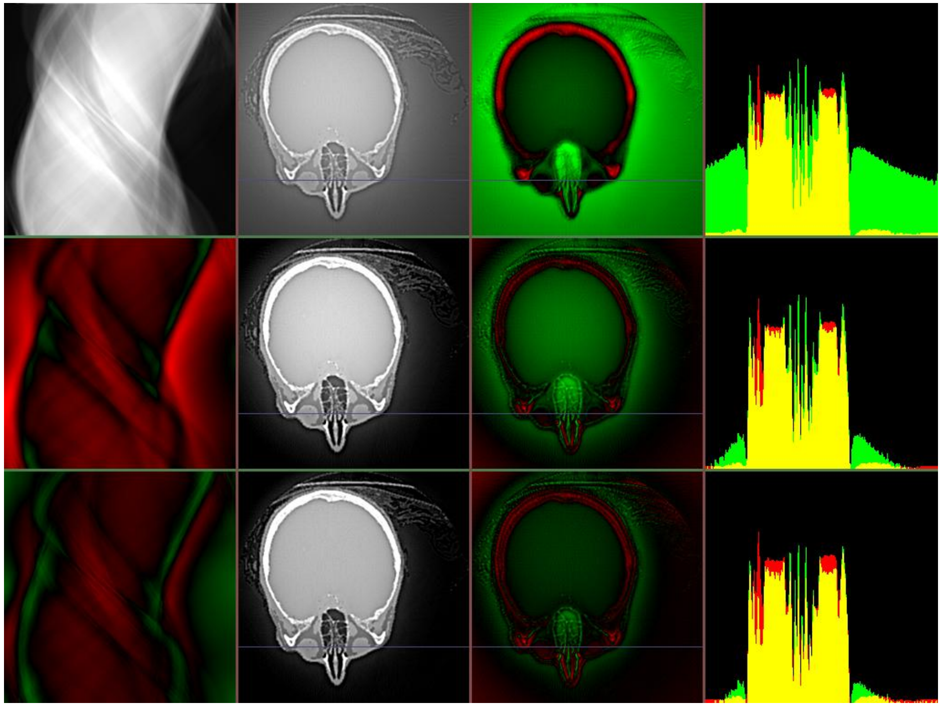

**A closer look: GPU**
AMD quad-core Opteron with ATI 1900 GPU
6 iterations used to reconstruct a 256x256x112 volume in 3s total


Copyright © 2009 RapidMind Inc.

RapidMind

# Product Overview

**RapidMind provides:**

1. A flexible ***platform*** that allows an arbitrary algorithm to be expressed and efficiently mapped to both multi-core CPUs and GPUs

2. Accelerated volume processing ***components*** that provide core building blocks for medical imaging applications

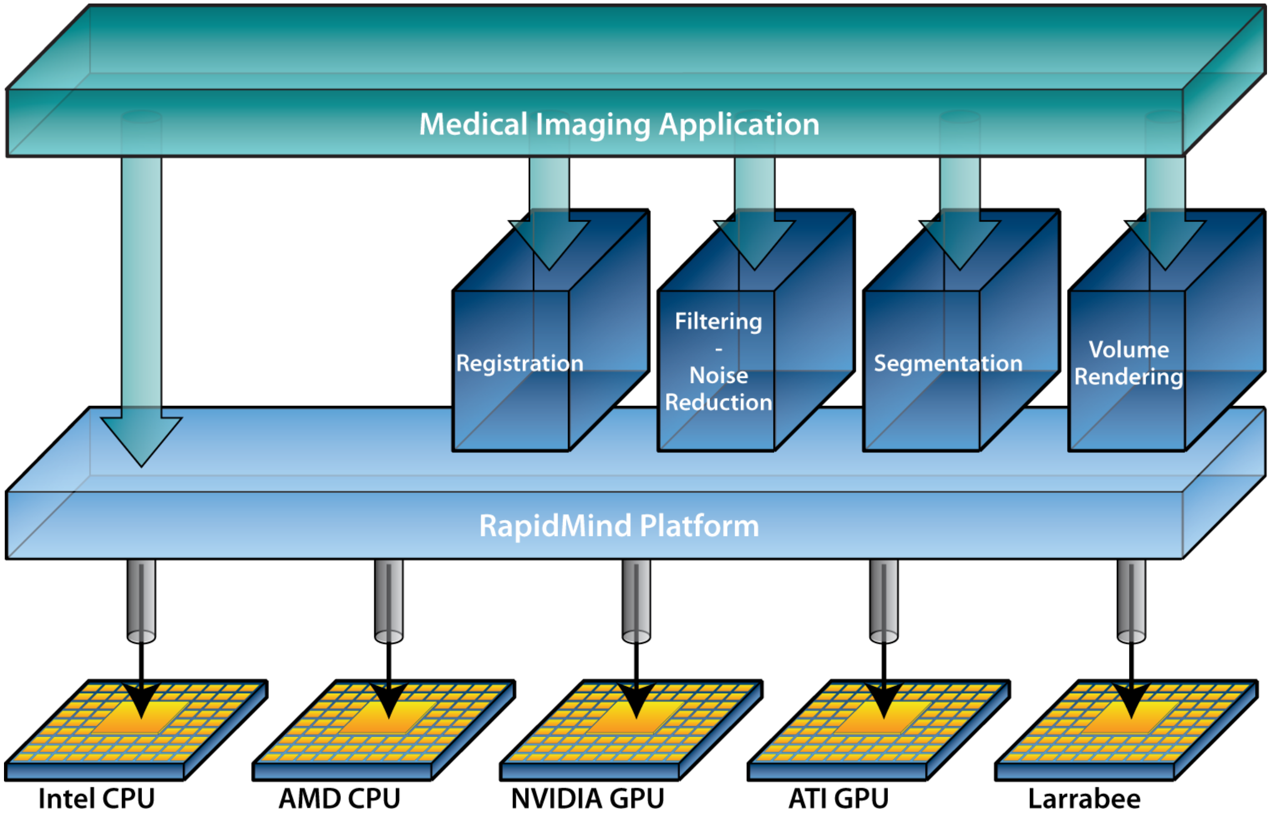



Copyright © 2009 RapidMind Inc.

## *The RapidMind Multi-Core Development Platform provides software organizations with:*

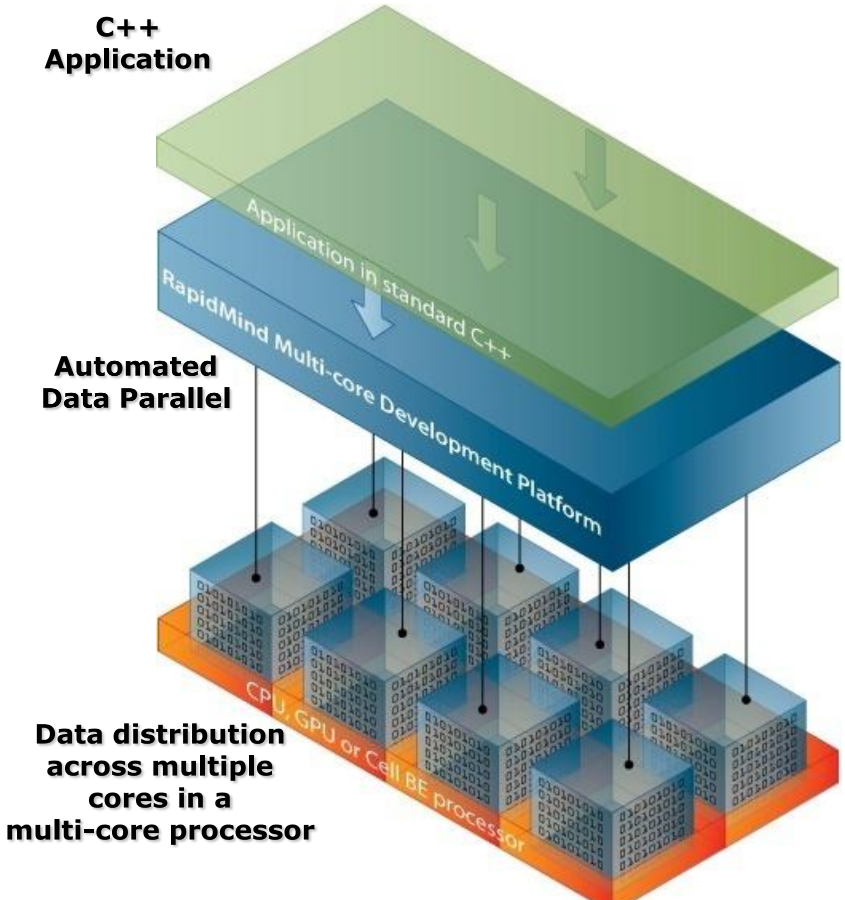


### Productivity

Build and deliver multi-core capable applications faster using existing practices, tools and compilers

### Performance

Significant performance gain on one core, multiple cores, *or* accelerators

### Portability

Applications are hardware independent and will automatically scale to additional cores and future multi-core processors and accelerators

Copyright © 2008, RapidMind, Inc.

# Portability
# Now and Into the Future

## Trends: Massive parallelism, heterogeneity, and hybrid computing

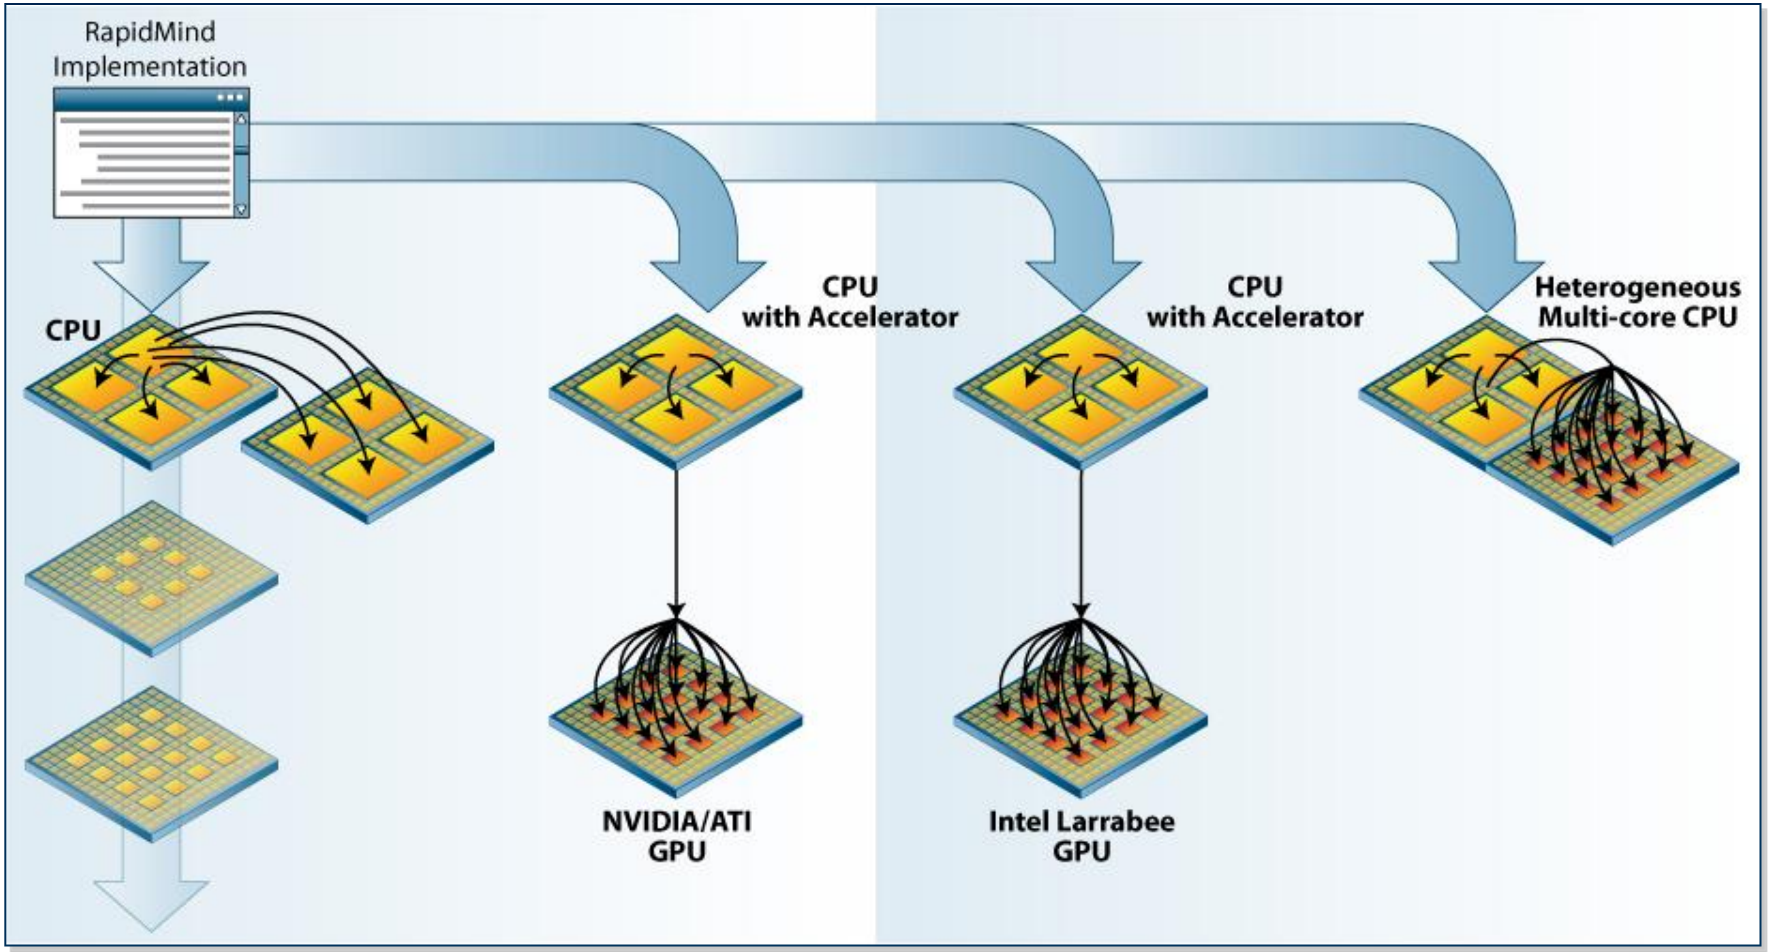


## RapidMind provides portability, scalability and future-proofing


Copyright © 2009 RapidMind Inc.

# System Architecture

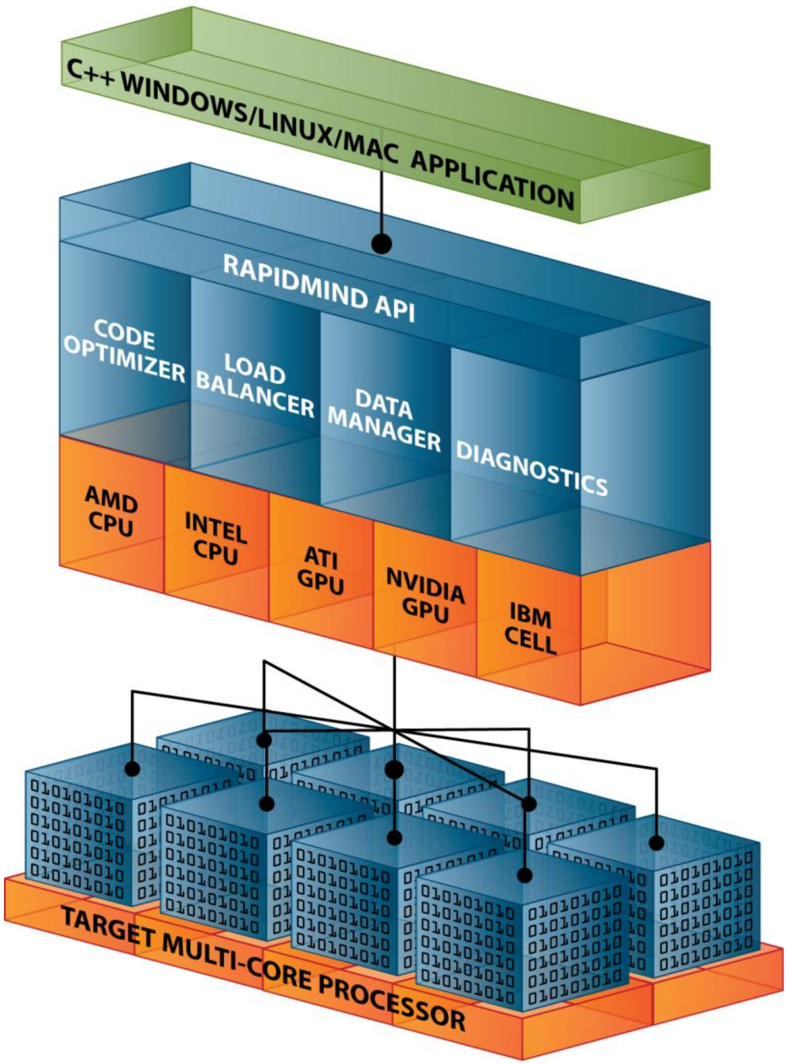


## API

- Intuitive, integrates with C++, and requires no new tools or workflow

## Platform

- **Code Optimizer** analyzes and optimizes computations to remove overhead
- **Load Balancer** plans and synchronizes work to keep all cores fully utilized
- **Data Manager** reduces data bottlenecks
- **Logging/Diagnostics** detects and reports performance bottlenecks

## Processor Support Modules

- x86 processors from AMD and Intel
- ATI/AMD and NVIDIA GPUs
- Cell Blade, Cell Accelerator Board, PS3

Copyright © 2008, RapidMind, Inc.

# Programming Model

- Map function over array:
  - Parallel application: **`C = f(A,B)`**
  - Can use control flow (SPMD model)
  - Can random access other arrays
  - Can declare and use local arrays
  - Can call sub-programs
  - Can read and write to subarrays
- Collective operations:
  - Reduce: **`a = reduce(p,A)`**
  - *Others…*

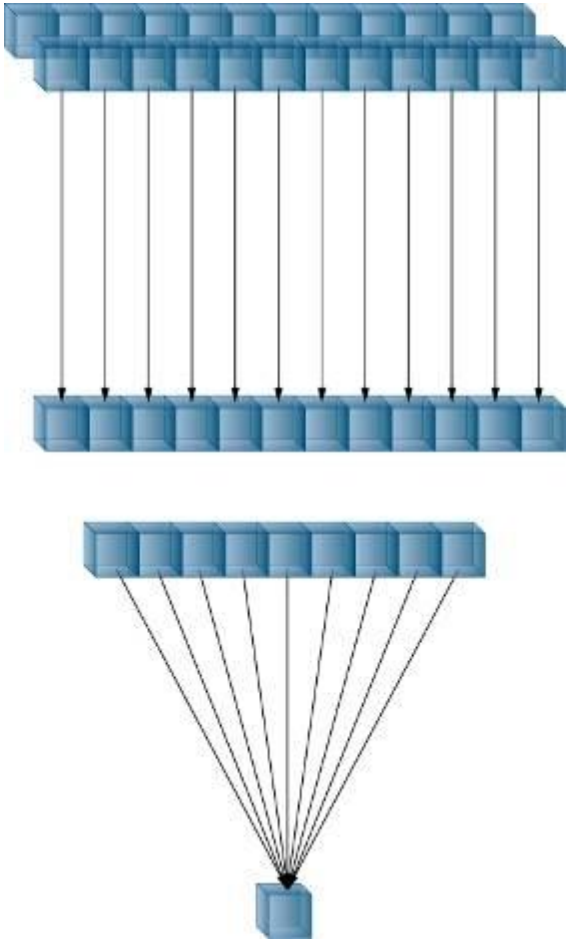

Copyright © 2008, RapidMind, Inc.

# Interface Summary

## Usage:

- Include platform header
- Link to runtime library

## Data:

- Values
- Arrays
- *Data abstraction*

## Programs:

- Defined dynamically
- Execute on coprocessors
- *Code abstraction*

```
#include <rapidmind/platform.hpp>
using namespace rapidmind;

Value1f f = 2.0f;
Array<2,Value3f> a(512,512);
Array<2,Value3f> b(512,512);

Program prog = BEGIN {
  In<Value3f> r, s;
  Out<Value3f> q;
  q = (r + s) * f;
} END;

a = prog(a,b);
f = 3.0f;
stride(a,2,2) = prog(
   slice(a,0,255,0,255),
   slice(b,256,511,0,255));
```

Copyright © 2008, RapidMind, Inc.

# Algebraic Reconstruction

**60x**

## RapidMind's algebraic reconstruction implementation can reconstruct an entire volume in under 3 seconds using a single GPU

### X-rays projecting through the volume

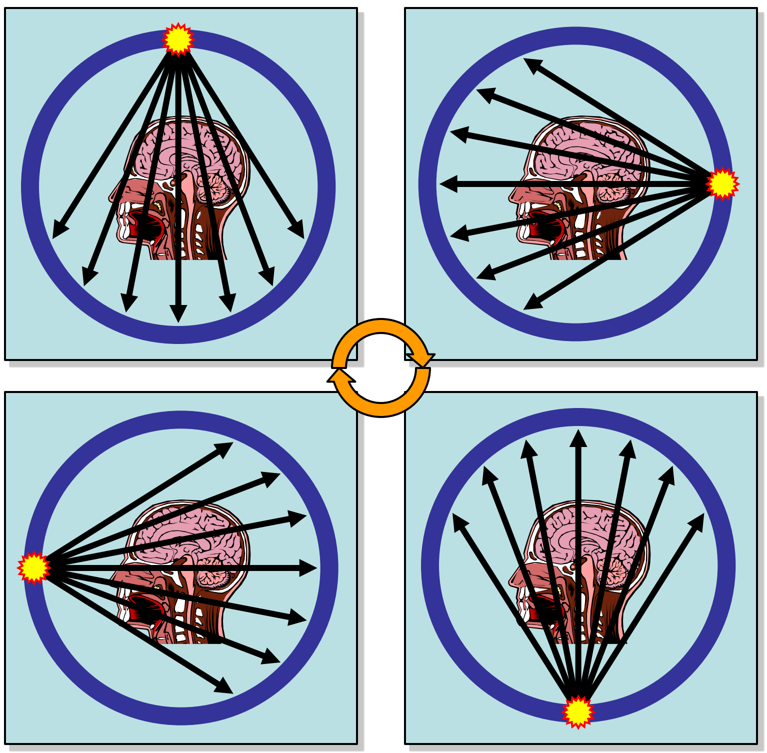

### Visualization of algebraic reconstruction of CT data

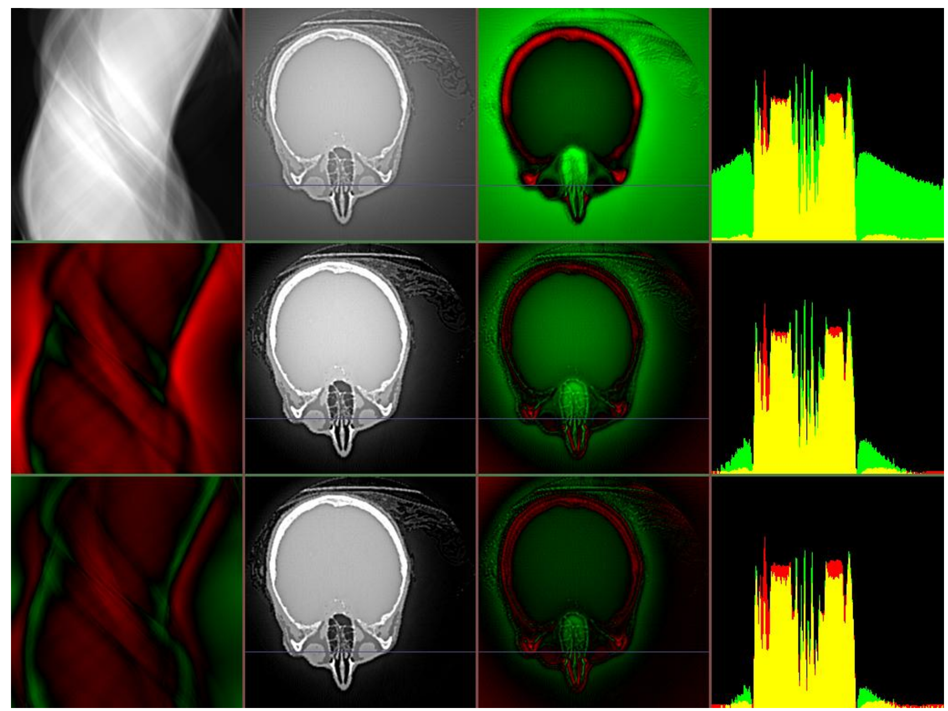

**A closer look: GPU**
AMD quad-core Opteron with ATI 1900 GPU
6 iterations used to reconstruct a 256x256x112 volume in 3s total


Copyright © 2009 RapidMind Inc.

# Kaczmarz Iterative Method

- Let y be the measured projections
- Let z be the (unknown) volume
- Let A be the volume-to-projection mapping
- Need to solve $y = Az$ for z given y
- Kaczmarz iteration:

$$z_{(i+1)} \leftarrow z_{(i)} + A^T R (y - A z_{(i)})$$

- Interpretation:
  - A is forward projection
  - $y - A z$ is correction
  - $A^T$ is backprojection
  - R is filtering (preconditioner; sharpening)

Copyright © 2008, RapidMind, Inc.

# Kaczmarz Iterative Method

```
Array<2,Value4f>
reconstruct (
  Array<2,Value4f> y
) {
  sharp = 100.0f;
  beta = 3.0f;
  Array<2,Value4f> z;
  z = backproject(y);
  for (int i = 1; i < n_iterations; i++) {
     Array<2,Value4f> c = project(y,z);
     z = backproject(z,c);
     sharp *= 0.05f; // converge to zero
     float blend = 0.7f;
     beta = blend*beta + (1.0f-blend)*0.9f;
  }
  return z;
}
```

Copyright © 2008, RapidMind, Inc. 

# Algebraic Reconstruction
# Kaczmarz Iterative Method

```
prog_backproj = BEGIN {
     In<Value2i> u;
     In<Value4f> base;
     Out<Value4f> r = 0;
     Value1f theta = 0;
     Value1f theta_inc = float(PI/data_height);
     Value2f p = Value2f(u)*Value2f(1.0/data_height,1.0/data_width)
               + Value2f(-0.5,-0.5);
     Value2f c = Value2f(-sharp,1.0f+2.0f*sharp);
     Value1i i = N;
     DO {
          Value2f d = Value2f(sin(theta),-cos(theta));
          Value1f t = dot(p,d)+0.5f;
          Value2f v = Value2f(t,theta/PI);
          Value2f w = floor(v);
          Value2f f;
          f(0) = v(0) - w(0);
          f(1) = 1.0f - f(0);
          Value4f ww = w(0,0,0,1) + Value4f(-1,1,2,0);
          Value4f s1 = data[w];
          Value4f s2 = data[ww(1,3)];
          Value4f s0 = data[ww(0,3)];
          Value4f s3 = data[ww(2,3)];
          Value4f h0 = c(0)*s0 + c(1)*s1 + c(0)*s2;
          Value4f h1 = c(0)*s1 + c(1)*s2 + c(0)*s3;
          r += f(1)*h0 + f(0)*h1;
          theta += theta_inc;
     } UNTIL (0 == --i);
     r = base + beta * r / float(N);
} END;
```

Copyright © 2008, RapidMind, Inc.

# Extensions

- Basic libraries and platform support for imaging
- Common patterns of parallel computation in imaging are built into the platform
  - Point operations and stencil-based filtering operations are “native” to the platform
  - Category reductions for computations on irregular “blobs” are built-in operations (also an implementation of “map-reduce”)
- Libraries for common operations:
  - K-means quantization
  - Flood-fill labeling of connected components
  - Noise reduction


Copyright © 2009 RapidMind Inc.

RapidMind
# Imaging Extensions

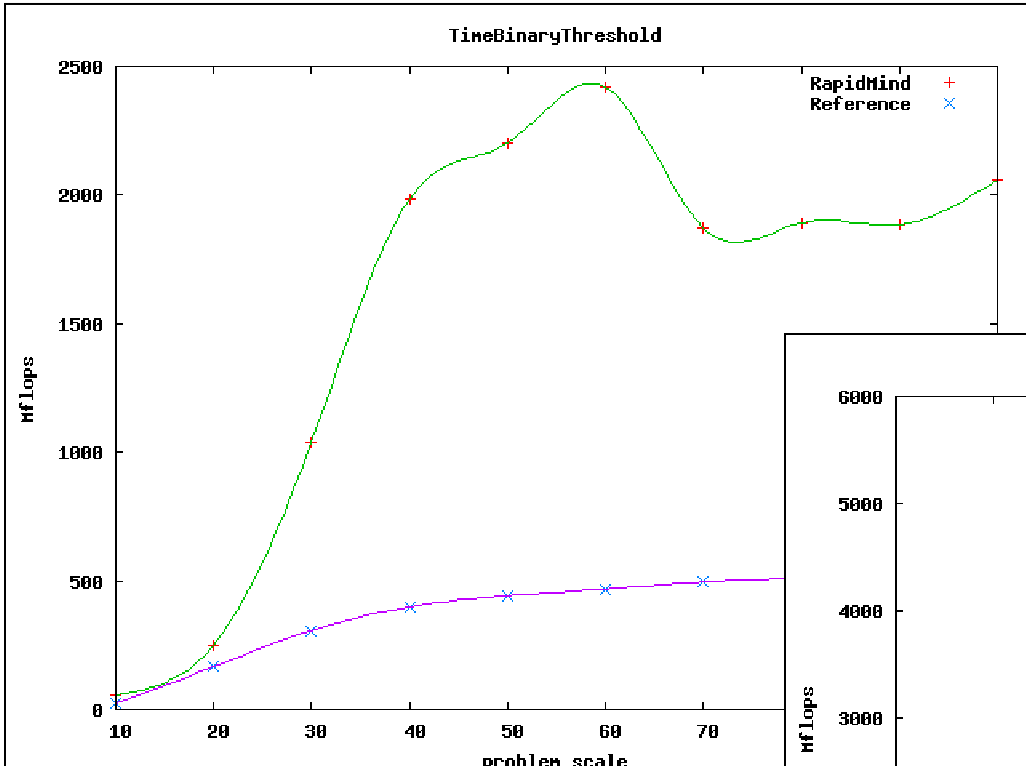


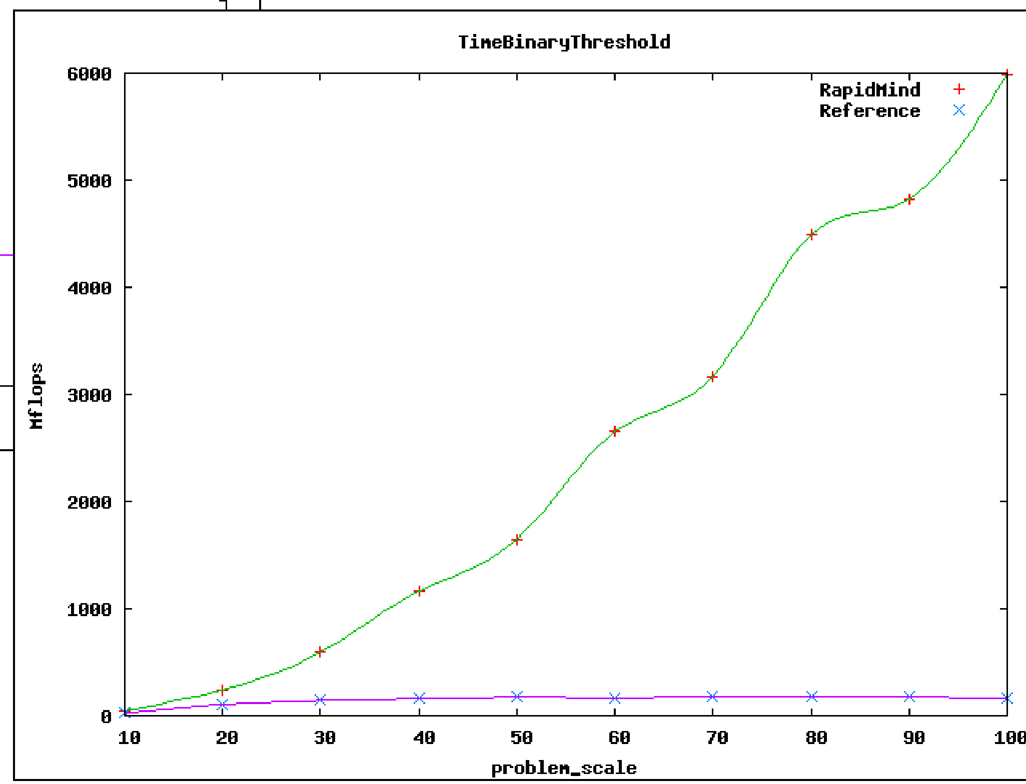



Copyright © 2009 RapidMind Inc.

RapidMind

# Imaging Extensions

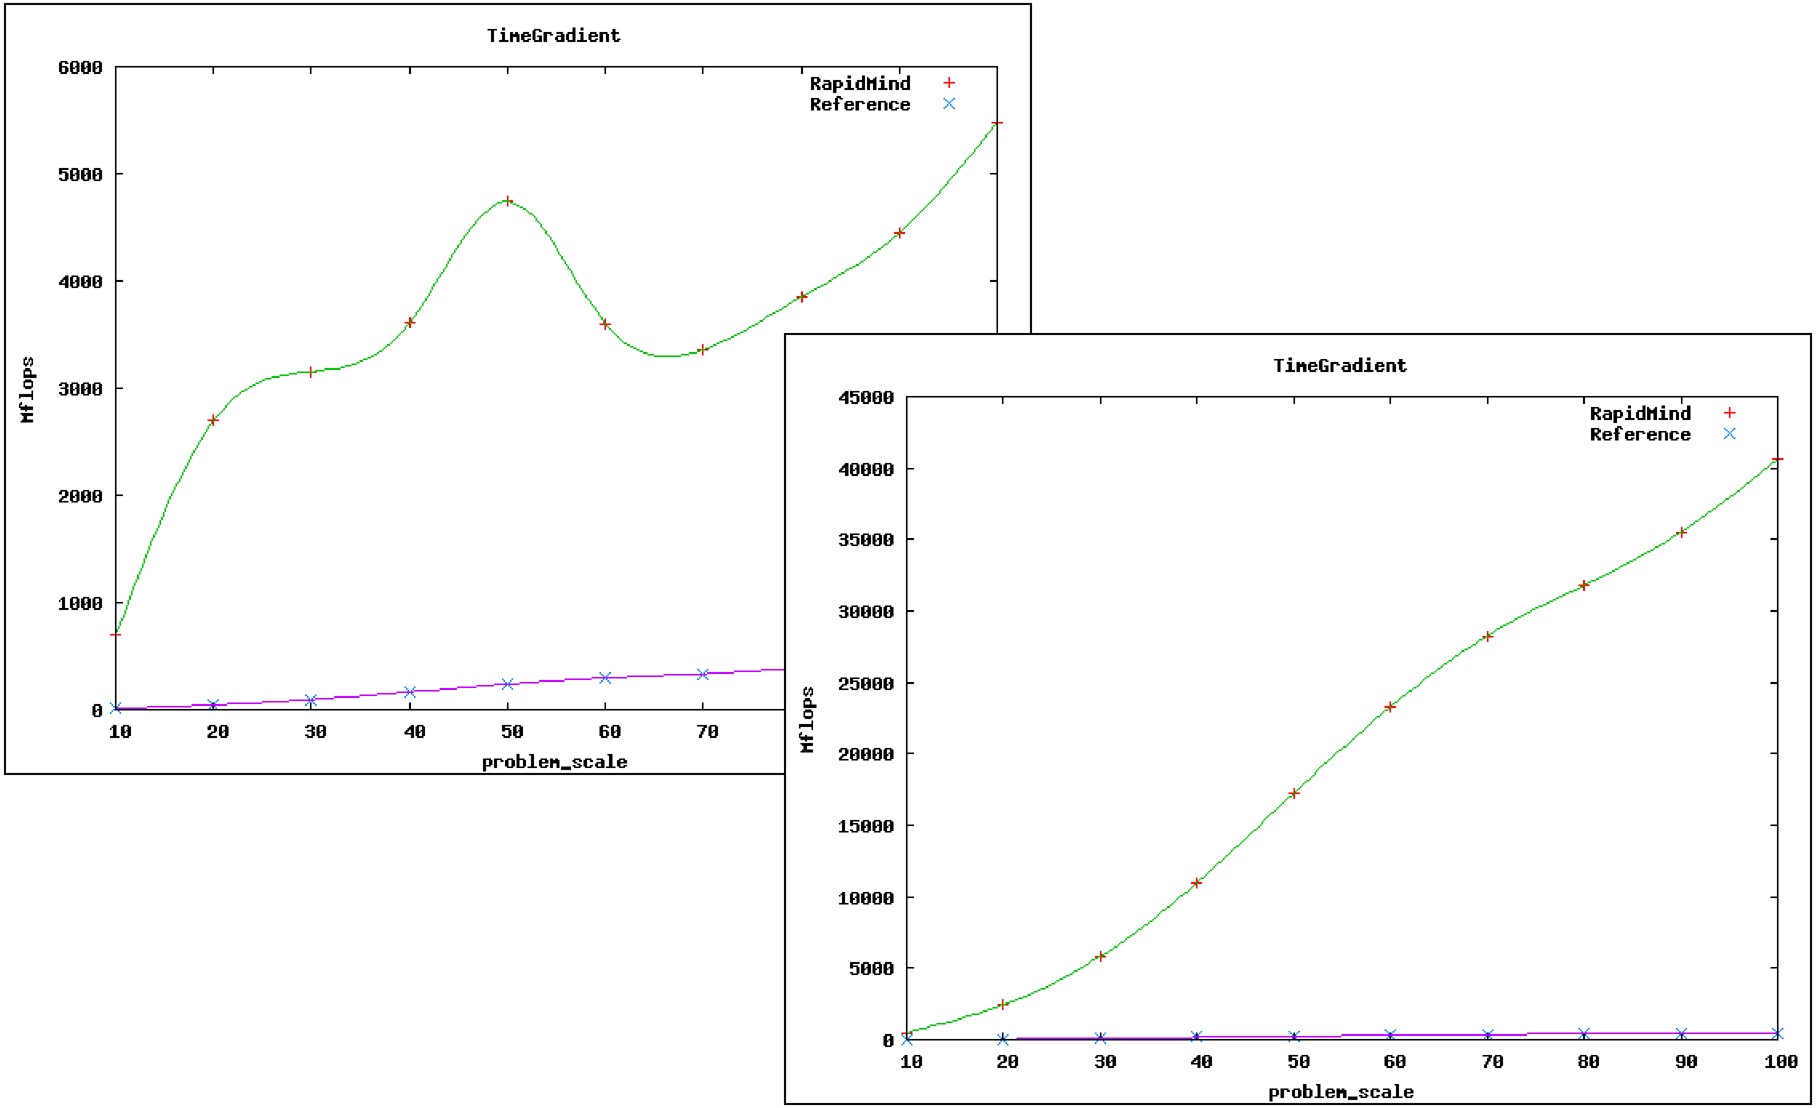



Copyright © 2009 RapidMind Inc.

RapidMind

# Imaging Components

## Components

- Libraries that solve some of the most challenging enabling functionality in the volume processing pipeline

- Accelerate performance by leveraging multiple CPU cores and GPU hardware acceleration
  - Minimum of 10x the performance of existing CPU-based implementations
  - Flexibility to customize through configuration, source license, or professional services

- Order of magnitude performance improvement can enable you to deliver new workflows and new applications

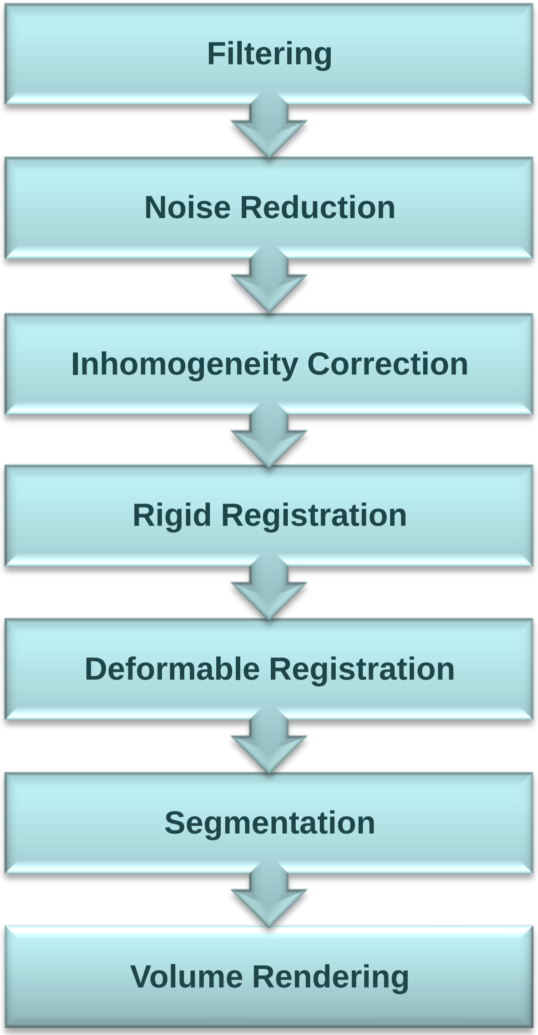



Copyright © 2009 RapidMind Inc.

# MSD Rigid Registration

## Key features and functionality

- Rigid registration with Euclidean transformations
- Mean squared differences comparison metric
- Linear interpolation of moving volume
- Bounding-box ROI on fixed volume
- Two volumes of up to 512x512x256 voxels
- Support for 16-bit precision data

## Use cases

- Comparison of rigid volumes taken at different times with the same modality


Copyright © 2009 RapidMind Inc.

# Performance

*Test case using **entire** 512x512x64 volume and a GeForce 280 GPU*

## Inner loop performance:

- Approximately 25 SSD comparison metric evaluations
- Metric comparisons require 95% of runtime
- Speedups of 75x to 95x observed over serial ITK baseline

## Overall performance:

- Serial ITK baseline: 254s
- RapidMind implementation: 16s (approx 16x faster overall)

## Observations:

- After acceleration, overall performance dominated by non-accelerated serial preprocessing. A portion of that could *also* be accelerated.
- ITK is still useful for preprocessing, as long as these operations do not dominate performance.


Copyright © 2009 RapidMind Inc.

# RapidMind Rigid Registration
# Performance vs. ROI Size

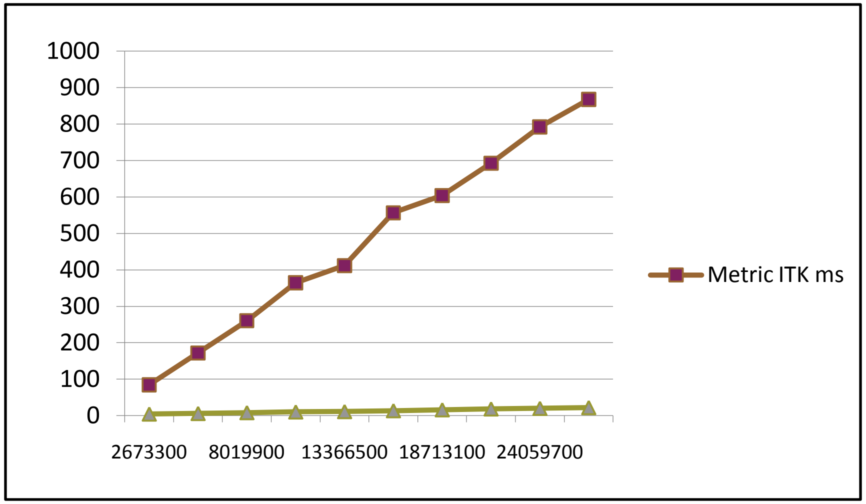


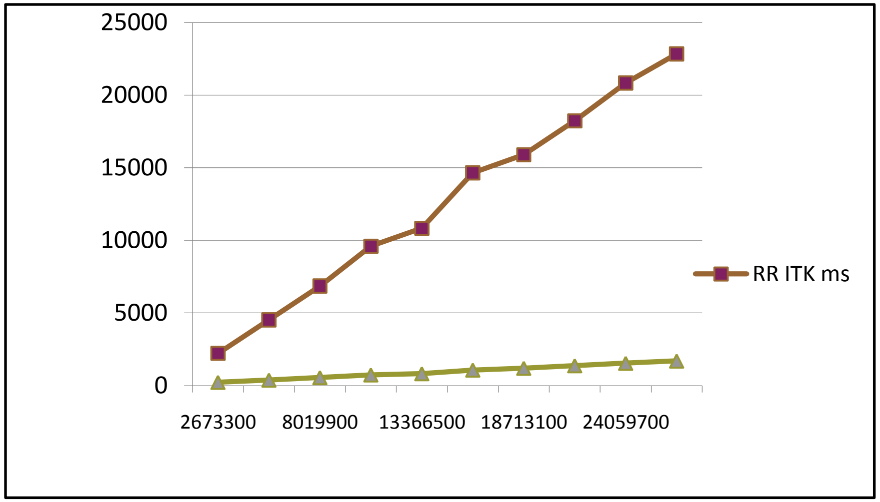


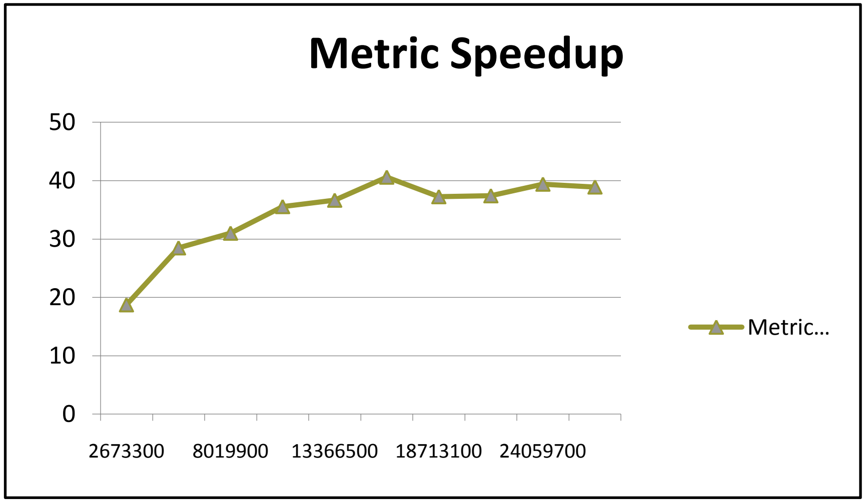


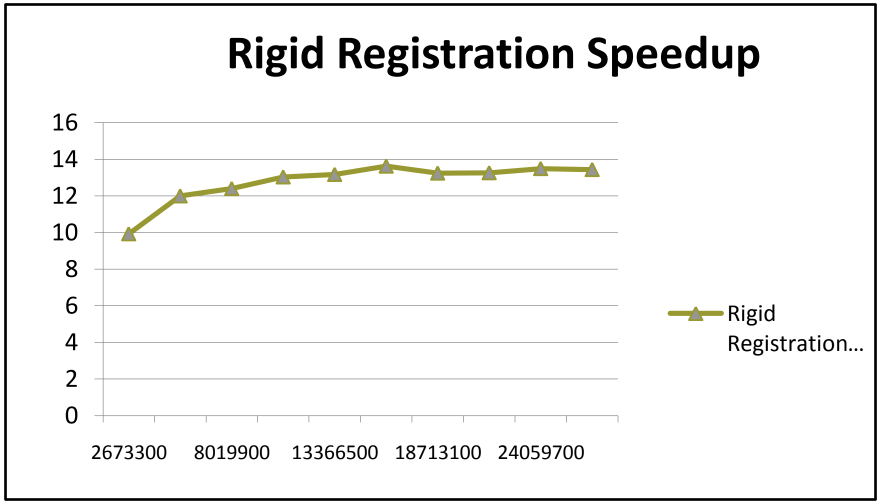

# Matching Goals and Capabilities

| Goal | Capabilities Needed |
|---|---|
| Reduce development costs | • Application developers can be responsible for optimization without having to become GPU programming experts<br>• Architects can leverage a set of fundamental components that let them concentrate on specific value-add functionality |
| Reduce time to market | • Leverage existing accelerated libraries for volume processing<br>• Access an external group of experts in developing accelerated imaging algorithms |
| Reduce risk | • Simultaneously support both NVIDIA and ATI GPUs<br>• Solution will continue to work on future processors such as Intel Larrabee<br>• Volume processing capabilities built on a customizable imaging platform that is easily extended |


Copyright © 2009 RapidMind Inc.

# RapidMind Conclusion

- Since 2004, entire company dedicated to delivering accelerated solutions on GPUs and multi-core CPUs
  - Spun off from 5 years of University of Waterloo research
  - Company is a recognized leader in accelerated computing
- RapidMind has helped software organizations accelerate their imaging applications by an order of magnitude
- RapidMind provides
  - A general-purpose software development platform that allows users to implement and deploy their own high-performance portable algorithms on accelerators
  - Customizable and portable core software components for accelerated medical imaging on both GPUs and multi-core CPUs

FROST & SULLIVAN
2009 BEST PRACTICES AWARD
North American Technology Innovation of the Year for Medical Imaging Software Solution

TR35 2008

Deloitte. Technology Fast 50 2007

HPCwire 2007 Editors' Choice Award

Best of Microsoft Tech·Ed 2007 WindowsITPro SQL Server SharePoint PRO

IEEE SPECTRUM Winners & Losers Winner 2007


Copyright © 2009 RapidMind Inc.