Optical Interconnects: Application Analysis and Architectural Requirements Study*

Albert Reuther

Janice Onanian McMahon

Robert Bond

MIT Lincoln Laboratory

DARPA/MTO Kickoff Meeting

June 23, 2003

*This work is sponsored by DARPA, under Air Force Contract F19628-00-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the author and are not necessarily endorsed by the United States Government.

Outline

- Introduction
- Some Candidate Applications
- Some Architectural Implications
- Program Plan
- Summary

- Motivation
- System Example
- Study Objective and Methodology

DoD High Performance Military Applications

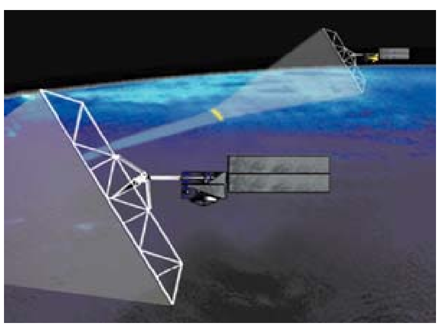

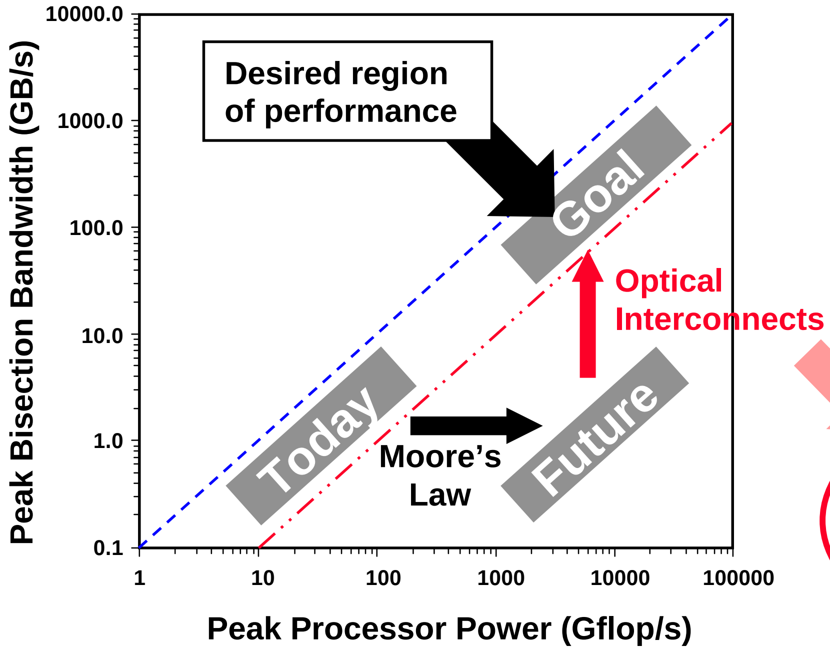

× 1000 speed/(power·volume) over electrical based interconnects

- Architectural balance is key to achieving delivered performance
- Optoelectronics could help maintain this balance for target applications

System Balance Example

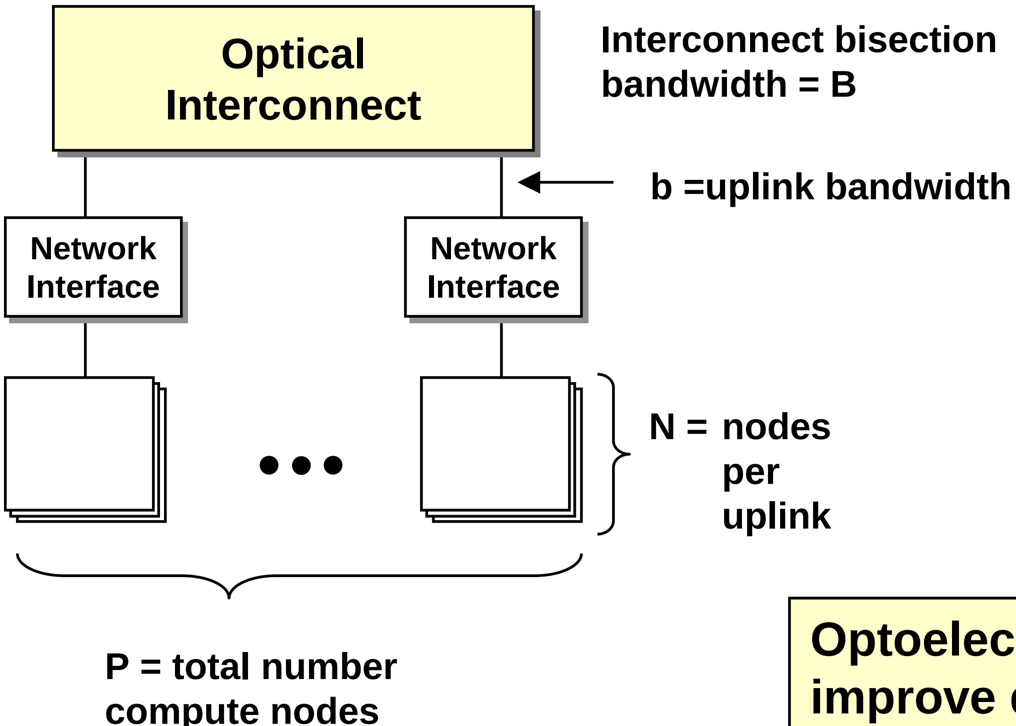

System bisection bandwidth

$$\min\left(\frac{P*b}{N}, B\right)$$

Limited by slowest link in communication path

Design issues that affect performance:

- Network interface
- Processor integration
- Communication protocol
- …

Optoelectronic interconnects can improve delivered performance only if they address system bottlenecks
⇒ Requires technology improvements motivated by *architectural* analysis
⇒ Motivates the invention of new architectures

MIT Lincoln Laboratory Summer Study

Objective

To motivate the development of advanced optoelectronic interconnect technology within the context of new computing architectures to address next generation DoD computing challenges

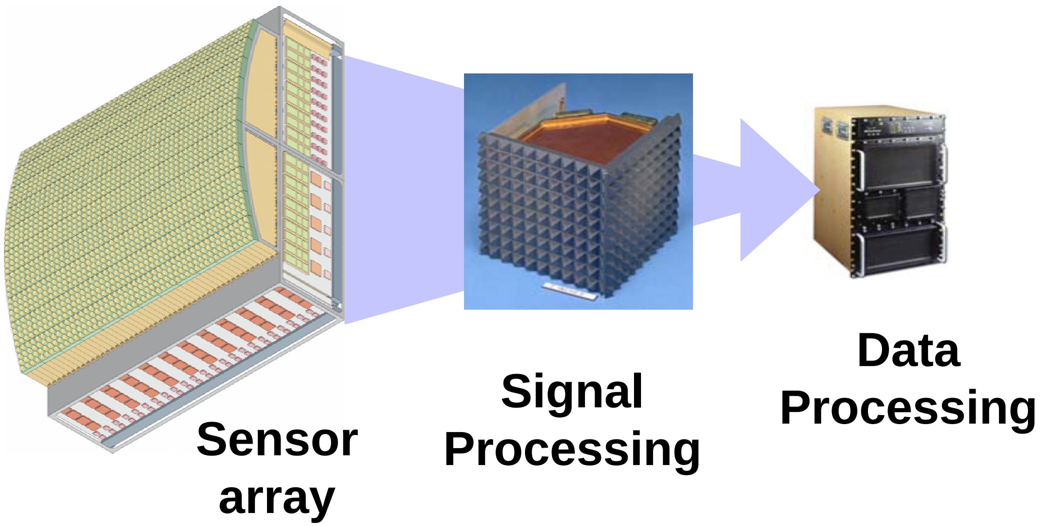

Methodology

1. Define DoD application requirements that drive new architectures
2. Define architectural approaches that drive new optoelectronic technology development
3. Identify research directions for future optoelectronic architectures

Architectural approaches will be evaluated in the context of application requirements

Outline

- Introduction
- Some Candidate Applications
- Some Architectural Implications
- Program Plan
- Summary

- Radar, GMTI STAP
- Image Processing
- Data Fusion

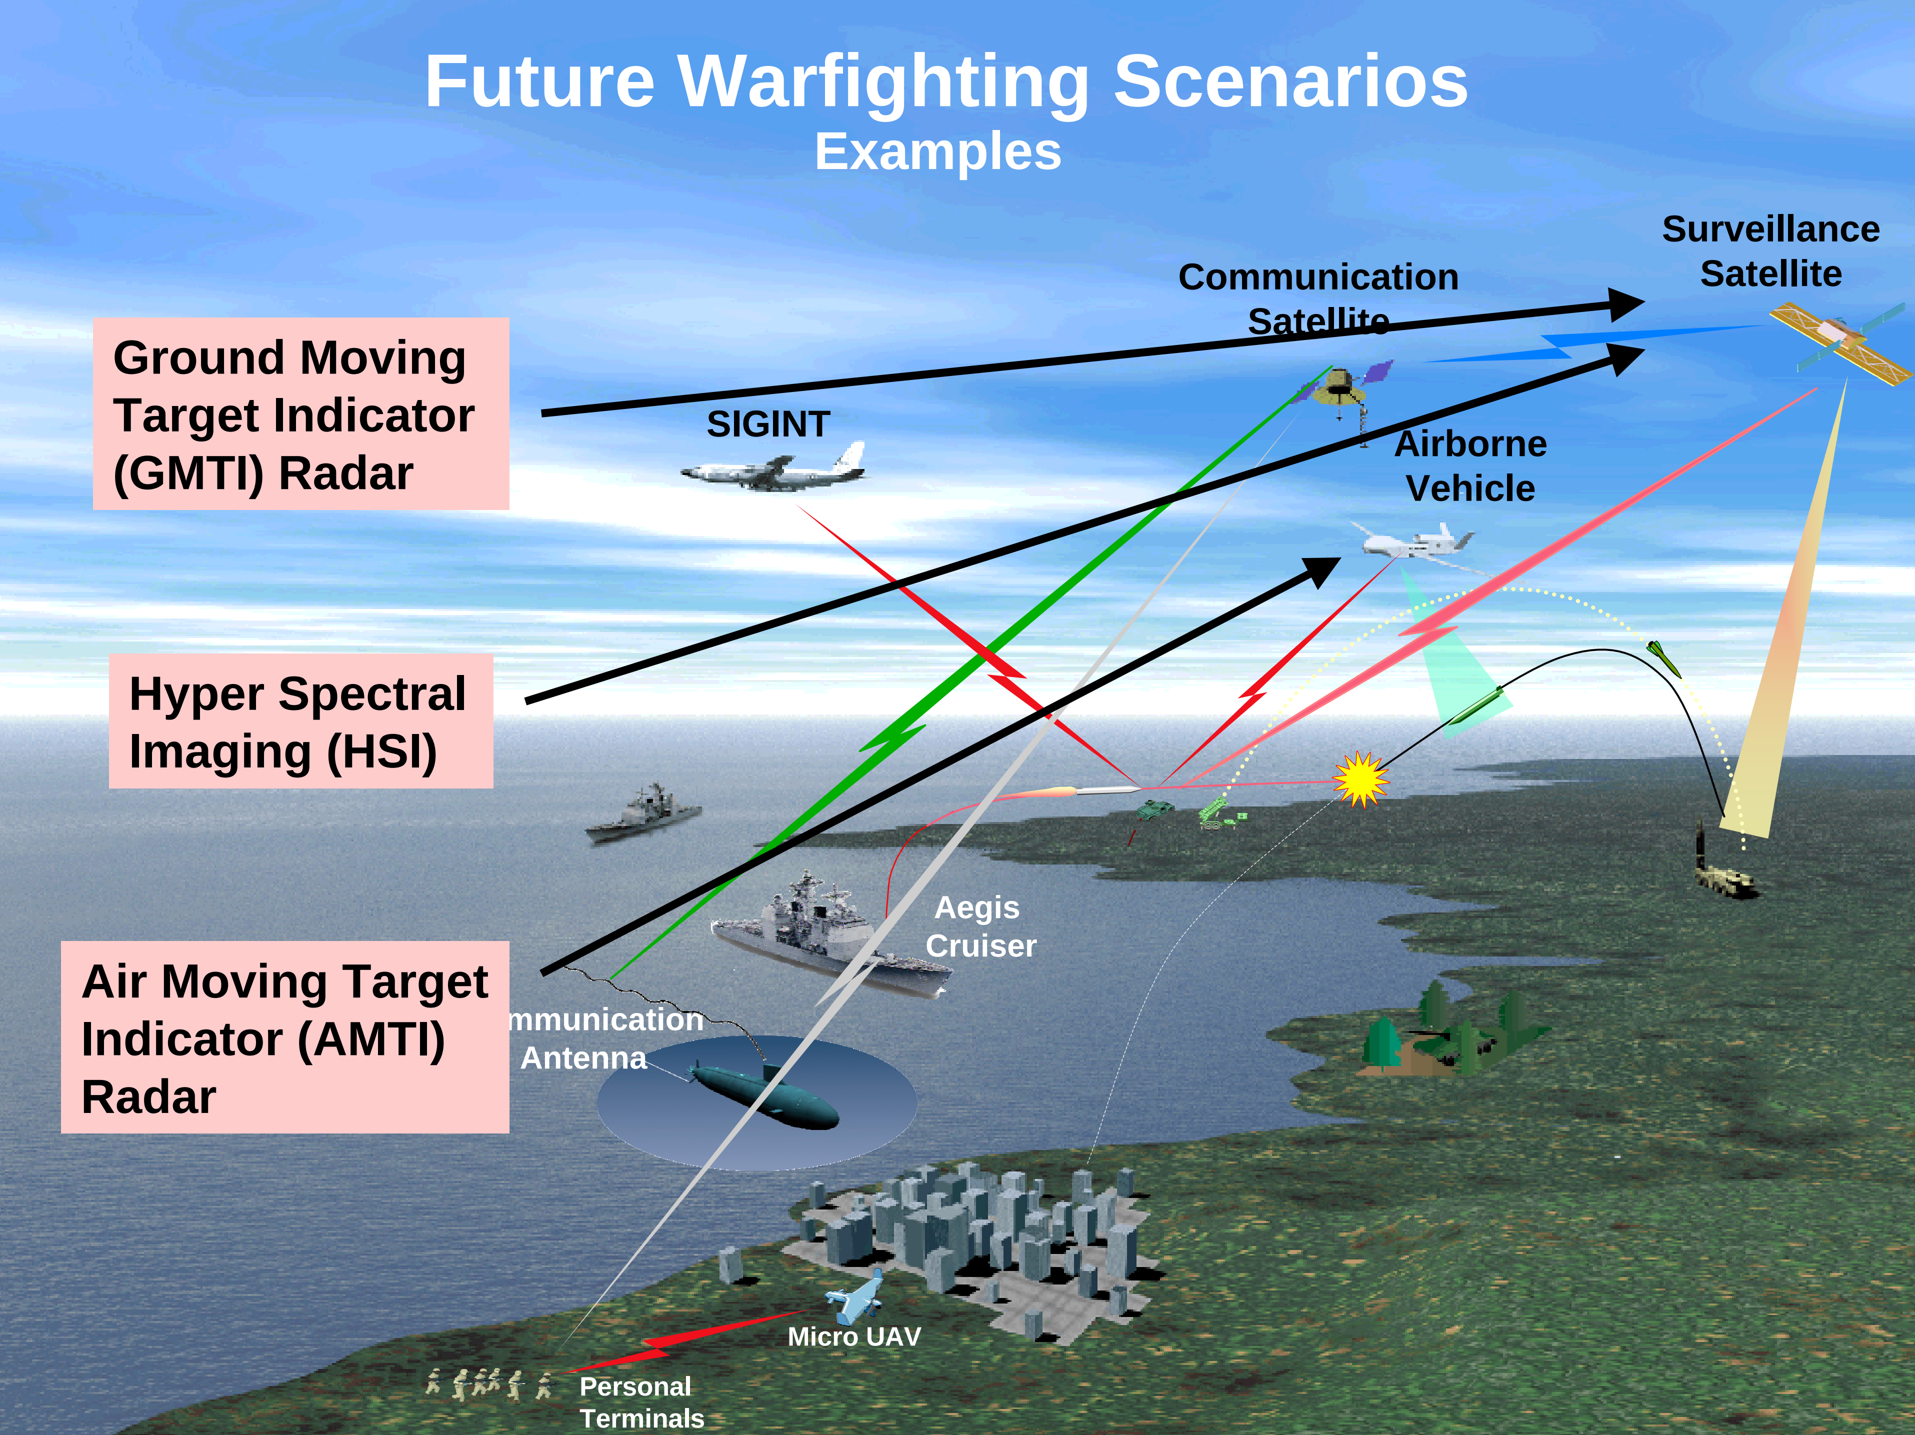
Future Warfighting Scenarios
Examples
Surveillance Satellite
Communication Satellite
Ground Moving Target Indicator (GMTI) Radar
SIGINT
Airborne Vehicle
Hyper Spectral Imaging (HSI)
Aegis Cruiser
Air Moving Target Indicator (AMTI) Radar
mmunication Antenna
Micro UAV
Personal Terminals

Space-Time Adaptive Processing
- Air Moving Target Indicator -

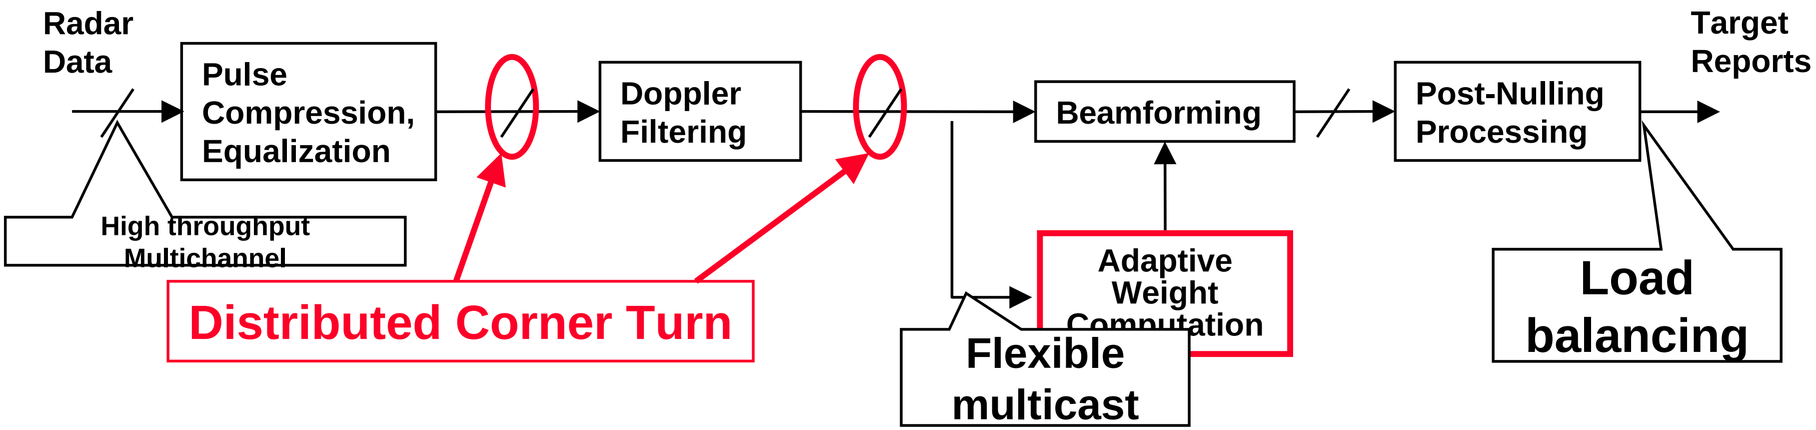

	AMTI Radar 1		AMTI Radar 2	
	Current	Future	Current	Future
Number of Channels	50	100	50	100
Data Cube Size (MB)	105	210	105	210
Processing Throughput (Gflops/s)	63	2,739	386	12,522
I/O Bandwidth (MB/s)	210	4,200	1,700	34,900

- Corner turn bandwidth ≈ factor of 10 over aggregate I/O bandwidth
- Distributed QR decomposition (fine grain communication) required

Corner Turn for Signal Processing

Corner turn ≡ change in matrix distribution to exploit parallelism in successive pipeline stages

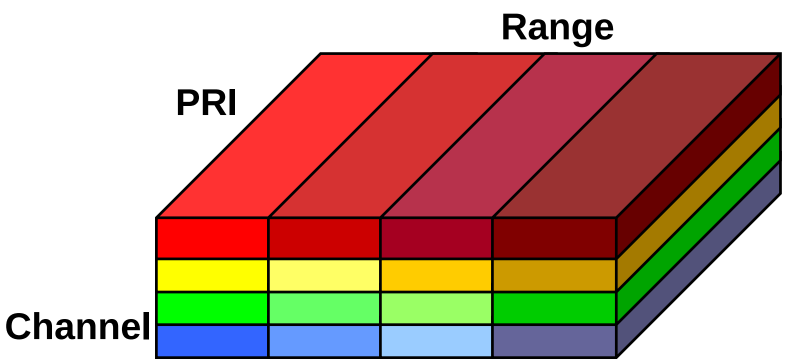

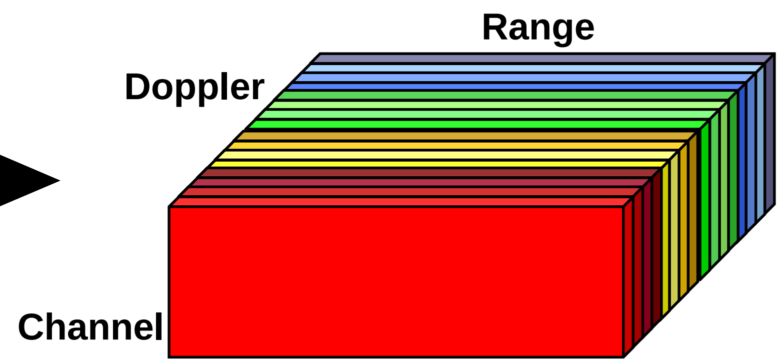

FFT of PRI data for each channel and range gate

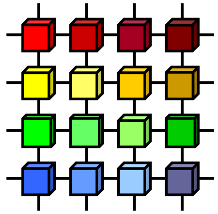

P_1 Processors

All-to-all communication where each of P_1 processors sends a message of size M to each of P_2 processors. Total data cube size is P_1P_2M.

Matrix inverse for each Doppler

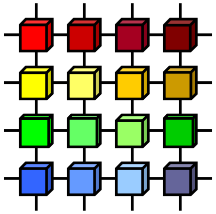

P_2 Processors

Corner turn operations require large amounts of connectivity and bandwidth

Ground Moving Target Indicator

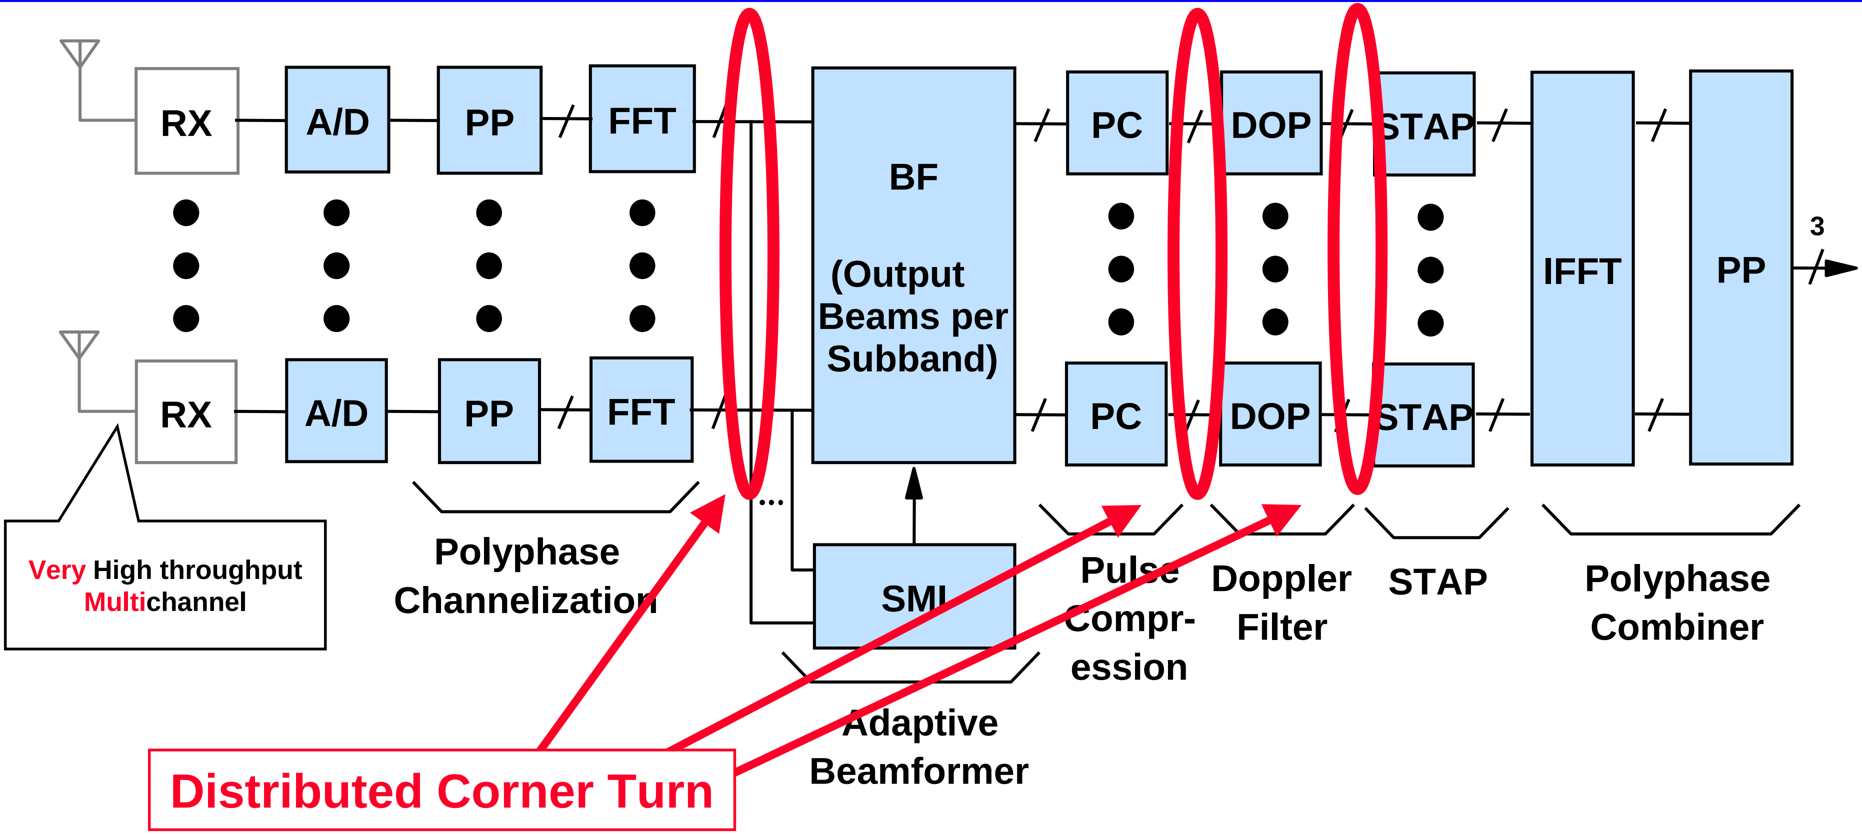

	Current	Future
Number of SubBands	50	200
Processing Throughput (Tflops/sec)	2.4	37
I/O Bandwidth (GB/sec)	80	1,000

Corner turn bandwidth ≈ factor of 10 over aggregate I/O bandwidth

Hyper-Spectral Imaging

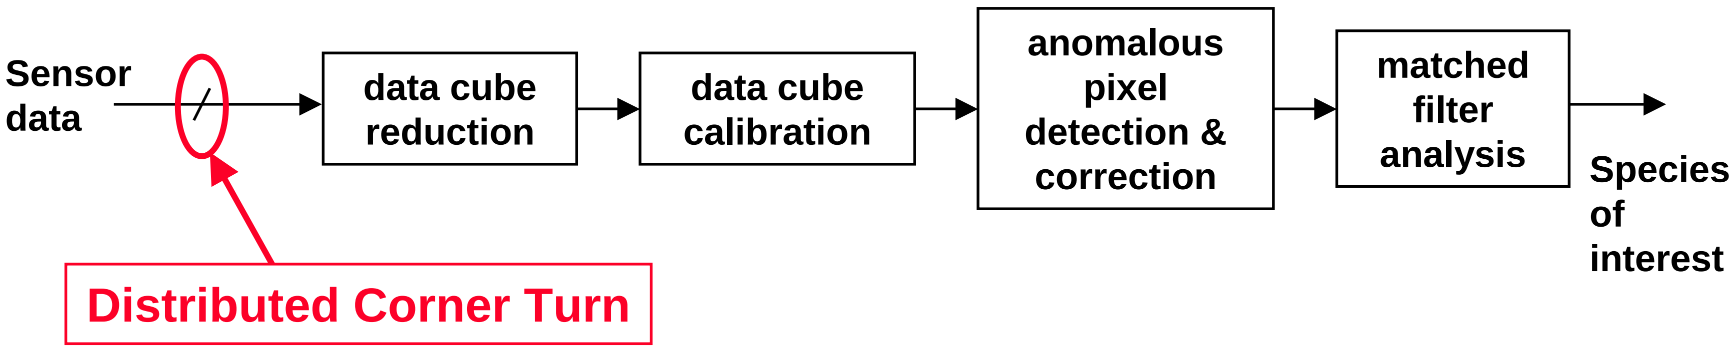

Hyperspectral InfraRed Imaging Spectrometer (HIRIS)

Number of Pixels	128^2	512^2	$(4K)^2$
Focal Plane Size (Gbits)	1.46	23.44	1500
Processing Throughput (Gflops/sec)	42	646	41200
I/O Bandwidth (GBytes/sec)	.18	2.9	188

- HIRIS Testbed Data (01)
- Current Technology Limit (03)
- Available circa 05-07

- **Large volumes of sensor data are exchanged prior to processing**
- **Fine grain kernels similar to STAP (QR decomposition, etc.)**

Honeywell VIVACE

Application Balance

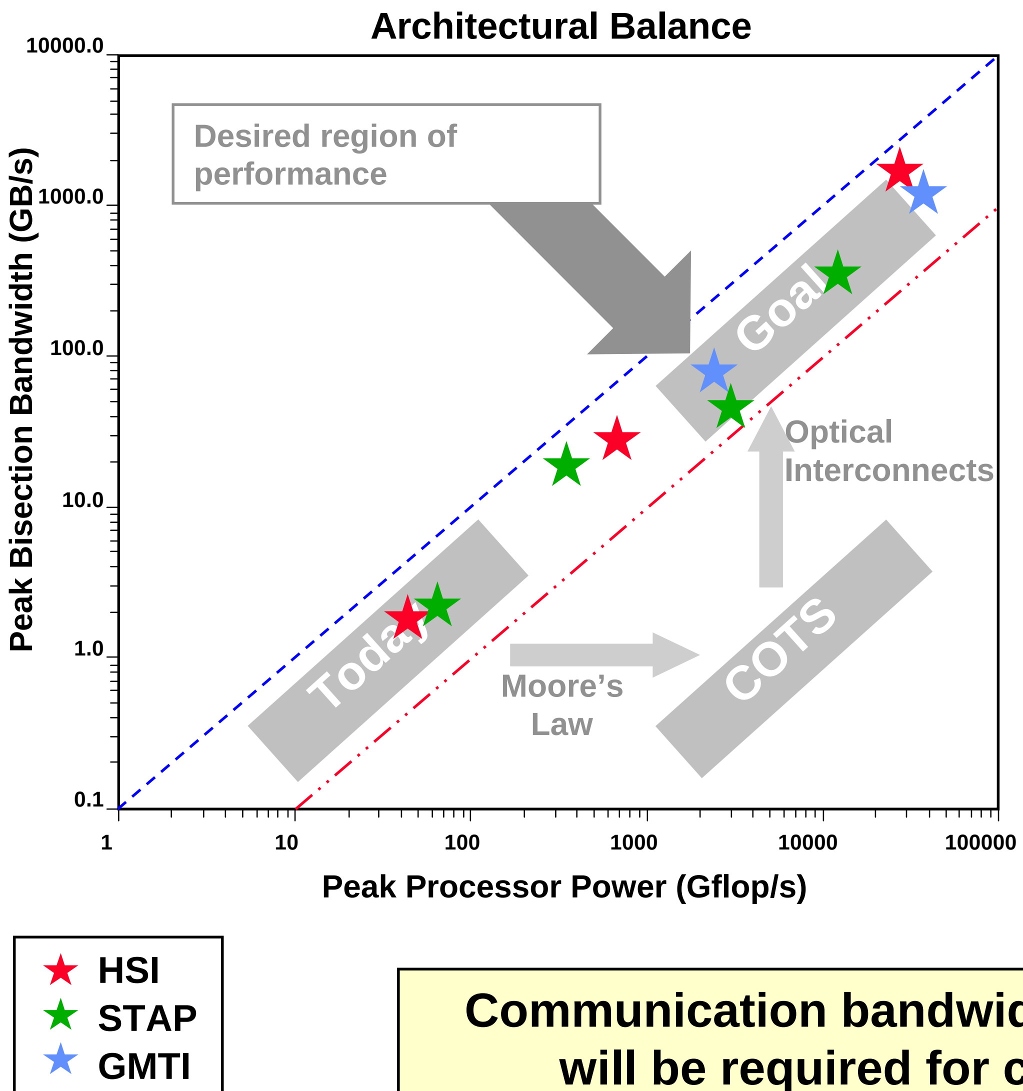

- Processing and I/O scale together
- Capabilities of optoelectronic interconnects are required to maintain architectural balance

Communication bandwidth of approaching TB/sec will be required for corner turn operations

Dynamic Load Balancing

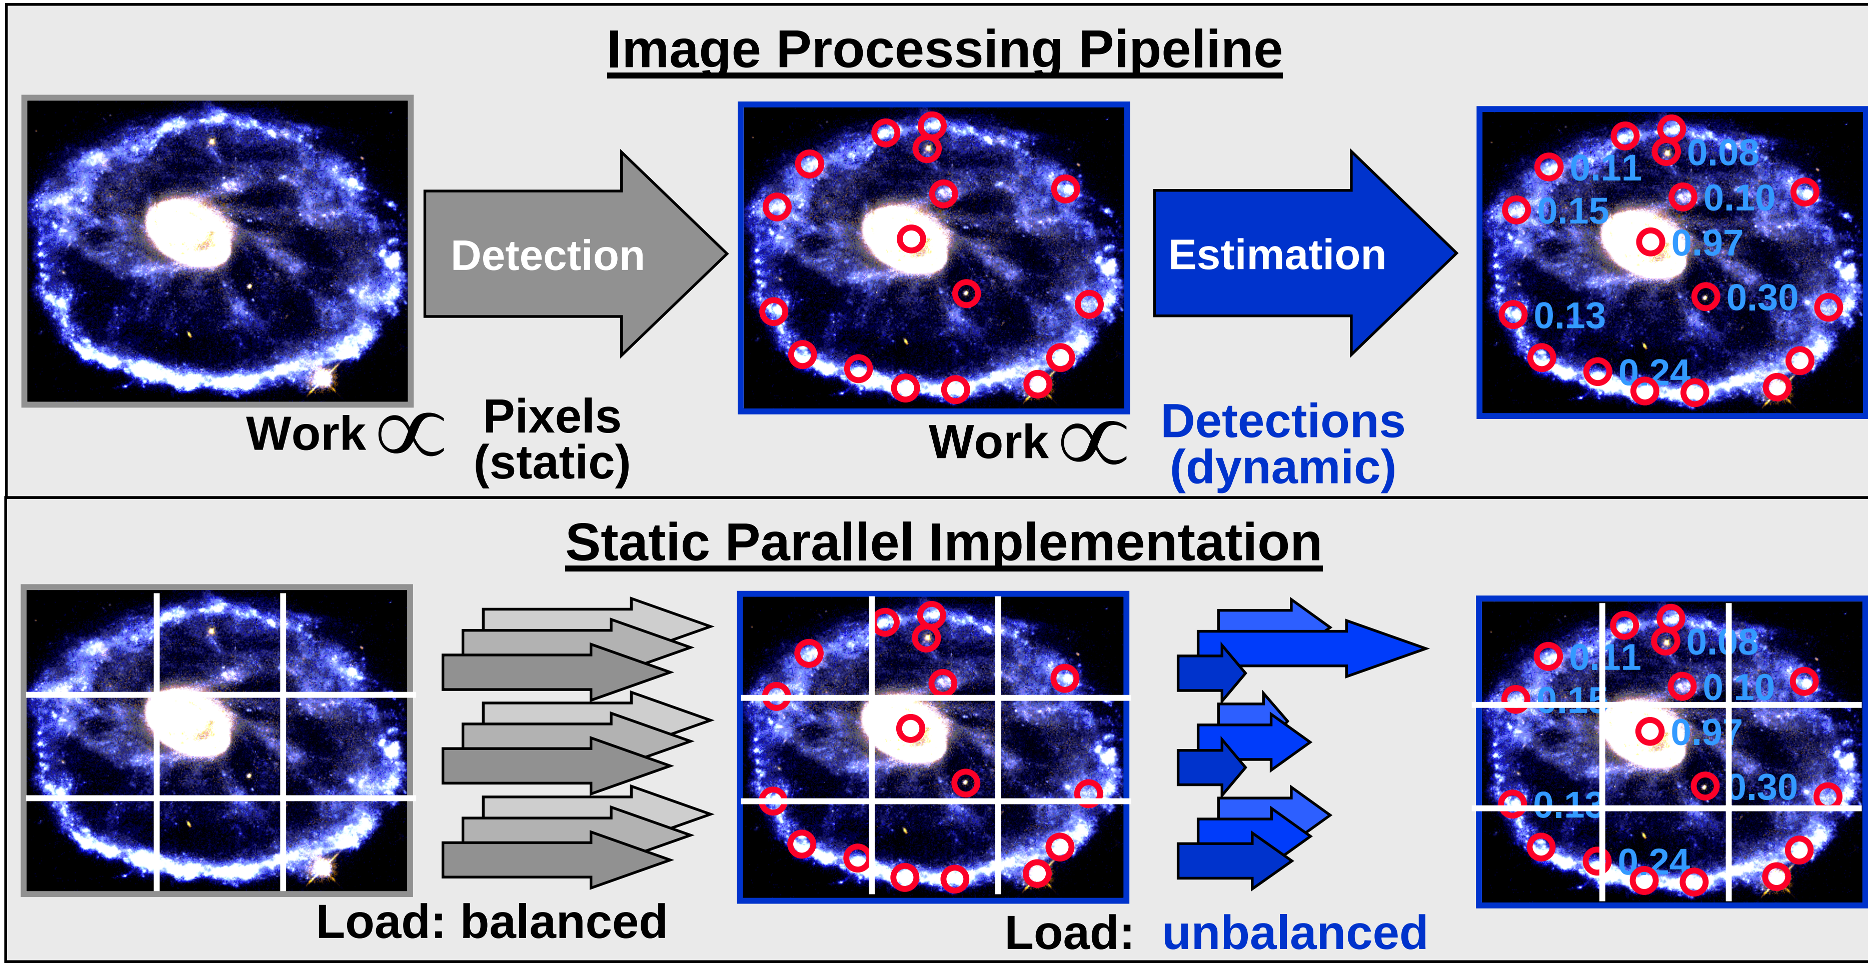

- Static parallelism implementations lead to unbalanced loads
- Dynamic load balancing stresses interconnection network

Imaging Data Overwhelms Computing Technology

New Sensors

Wide Field Imaging

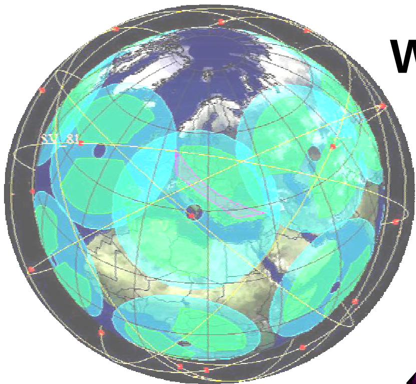

Hyper Spectral Imaging

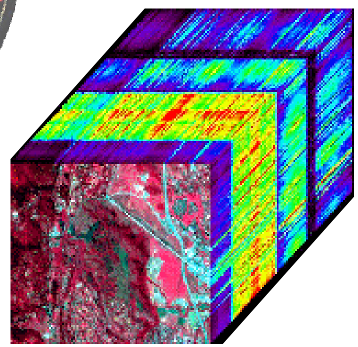

Massive Datasets

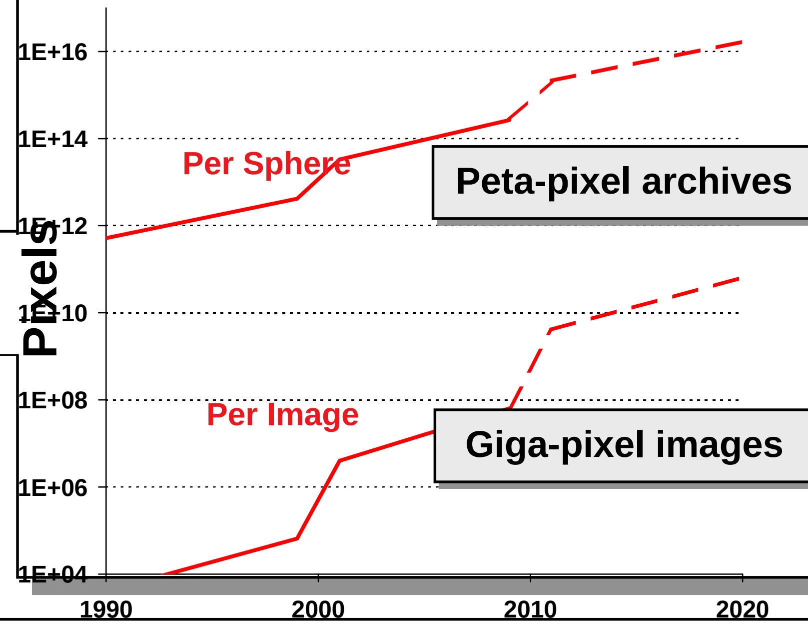

Novel Algorithms and Petascale Computing Architectures

Input Data → 3D Site Model Generation → Cross-Modality Registration → Image Conditioning →

Multi-Image Fusion ($O(N_{pixel}N_{sensor}^2)$) → Context Feature Extraction ($O(N_{pixel}N_{sensor}^2)$) → Site Mining → Detection Processing

10^{12}-10^{15} *Ops/scene* 10^{15}-10^{18} *Ops/scene*

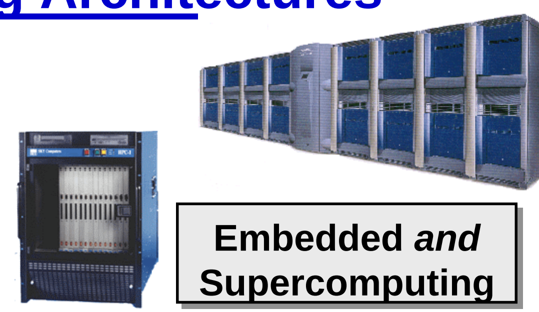

Embedded *and* Supercomputing

Outline

- Introduction
- Some Candidate Applications
- Some Architectural Implications
- Program Plan
- Summary

- Potential for optics
- Evaluation framework

Potential for Optical Interconnects

Advantages of Optics

- Higher density (Gbits/cm^2) information flow than wires
- No noise through parasitic capacitance and inductance
- Scalable interconnection bandwidths
- Low loss in free space or waveguides

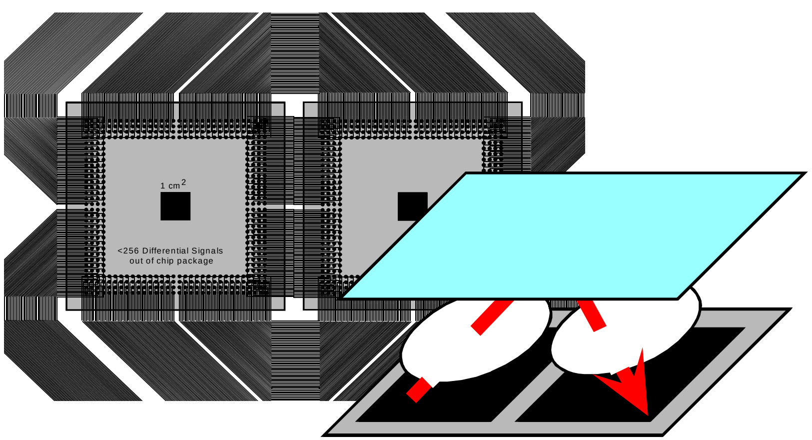

1 cm^2 1024 Signals

DoD High Performance Application Needs

- Balanced architecture
- Large scale (many processors)
- High data bandwidth, connectivity and flexibility, low latency
- Non-uniform memory access, dynamic data movement
- Low power and size
- High software productivity
- Reasonable cost

Optical interconnects may provide scale, bandwidth, flexibility and connectivity required for supercomputing applications

Example: Optical Interconnects and Cache Architectures

Current Cache Architectures

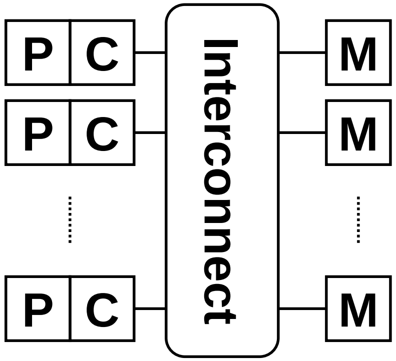

Optical Memory Interconnect

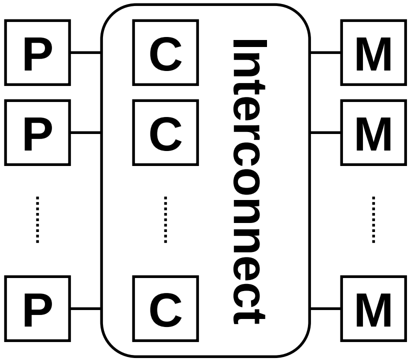

- **Exploit data locality to hide memory latency via on-chip caches**
- **<u>Problem:</u> cache performs badly on random, non-uniform data access**

- **Interconnect keeps a cache of connection topologies for memory access patterns**
- **Improve memory access time for random patterns via flexibility of interconnect**

Capabilities of optics could dramatically change the design of high performance computing architectures
⇒ Issues: feasibility, cost, power, size

Example Framework for Development of Metrics and Approaches

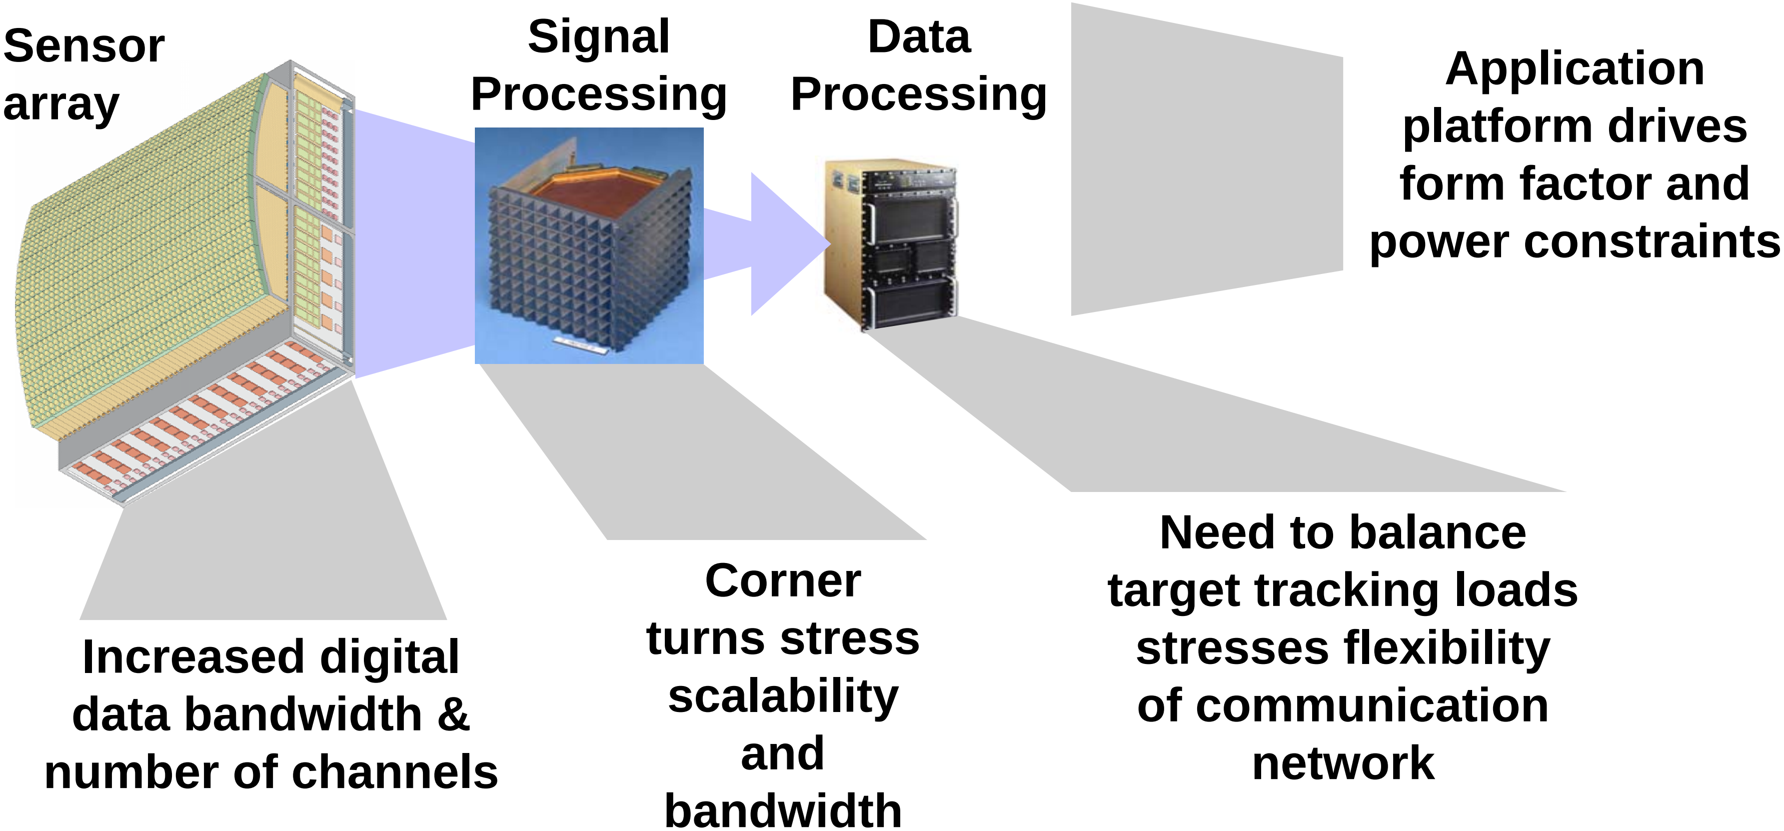

Framework will be used to derive benchmarks and metrics that drive development

Outline

- Introduction
- Some Candidate Applications
- Some Architectural Implications
- Program Plan
  - Team Members
  - Schedule/Milestones
  - Deliverables
- Summary

Plan, Schedule, Milestones, Deliverables FY03 Summer Study

2003

Month	Milestone
May/June	Study formulation
July	Preliminary Results
August	Final Results
September	Final Report

Team Members
- Larry Rudolph - MIT
- Bob Bond - MIT LL
- Jeremy Kepner - MIT LL
- Janice McMahon - MIT LL
- Albert Reuther - MIT LL

Deliverables
- Framework definition
- Metrics
- Requirements

$\Rightarrow$ Architectural approaches

Summary

- MIT Lincoln Laboratory is performing a summer study (FY03) with the objective of motivating the development of advanced optoelectronic interconnect technology within the context of new computing architectures to address next generation DoD computing challenges
- Basis for a new program
- Deliverables include an evaluation framework synthesized from military DoD high performance applications
  – Framework will showcase potential benefits of joint development of optical interconnect technology and new system architectures.
- The study will identify metrics to be applied at multiple scales
  – Component
  – Architecture
  – Application

Architectural Scalability

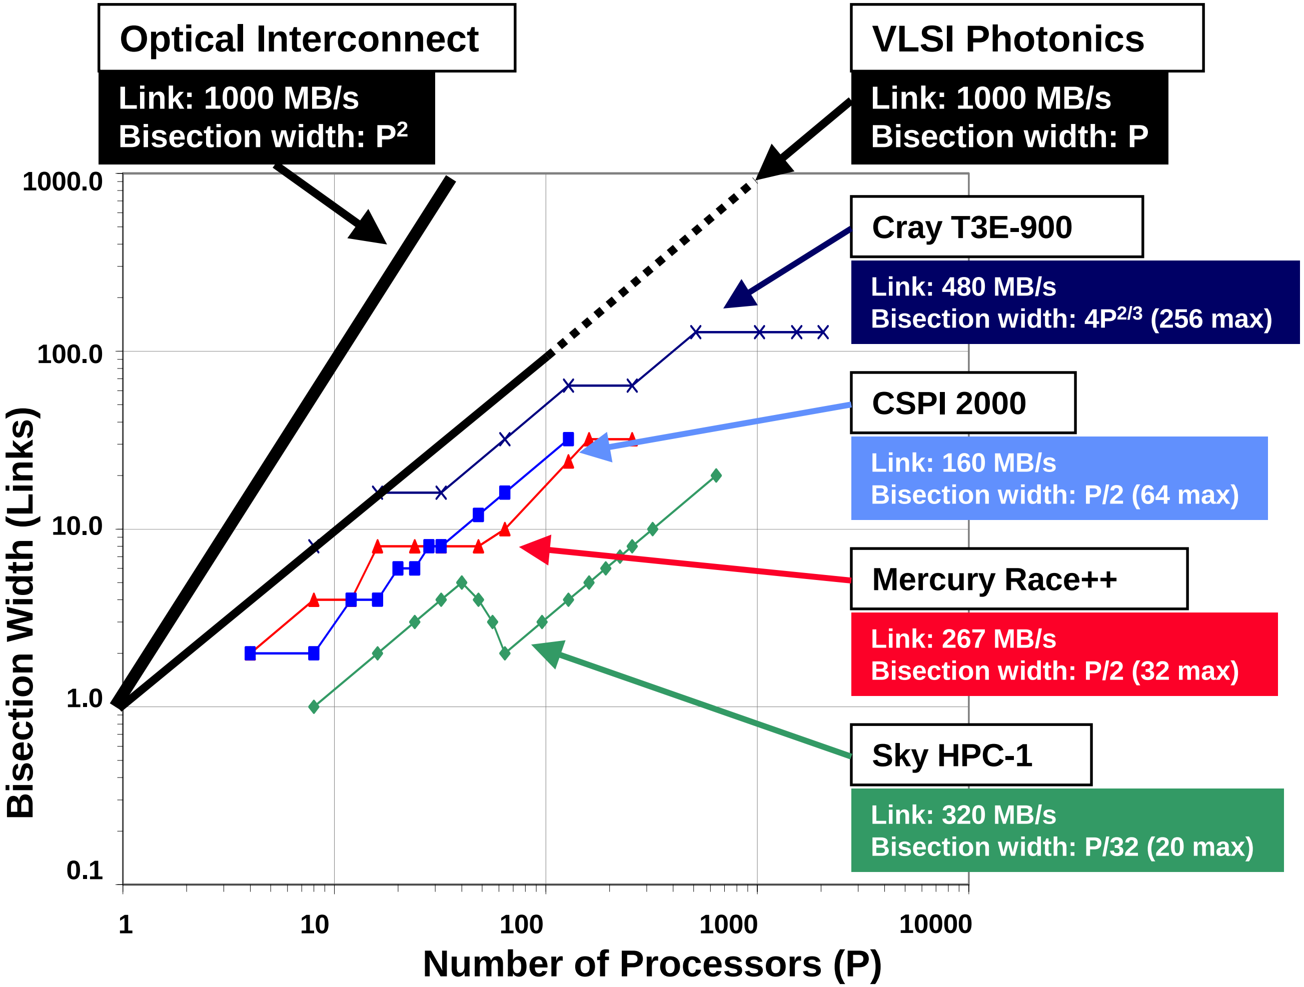

Optoelectronic interconnects mitigate scalability issues by providing high levels of connectivity per form factor

Potential benefits:

- Link bandwidth (significant)
- Interconnect (very significant)
- Latency (unknown)

Example Framework for Development of Metrics and Approaches

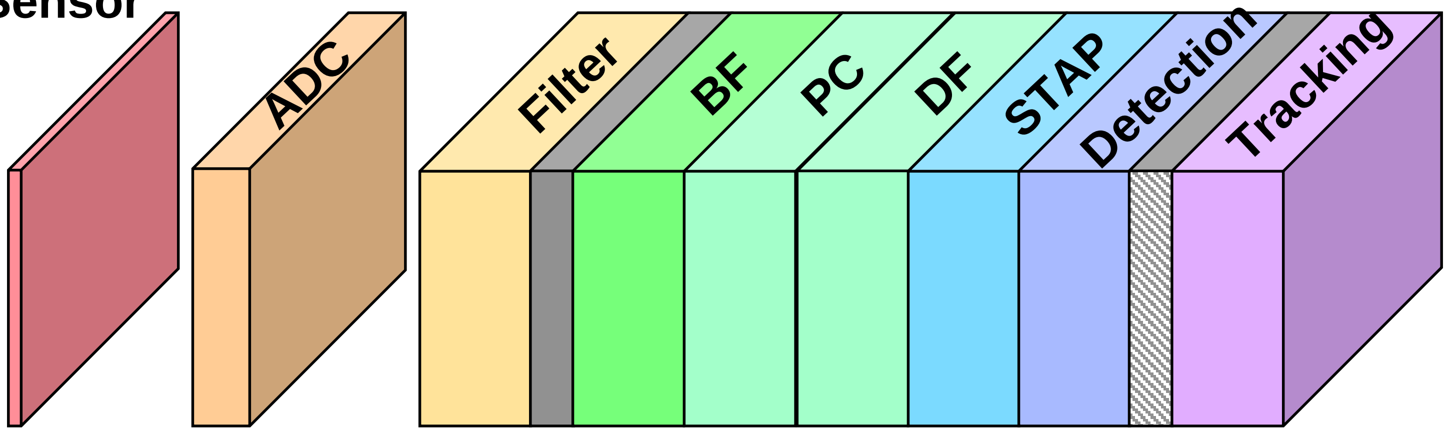

Application platform drives form factor and power constraints

Trend toward digital front end increases sensor size, data bandwidth

Corner turns stress scalability and bandwidth

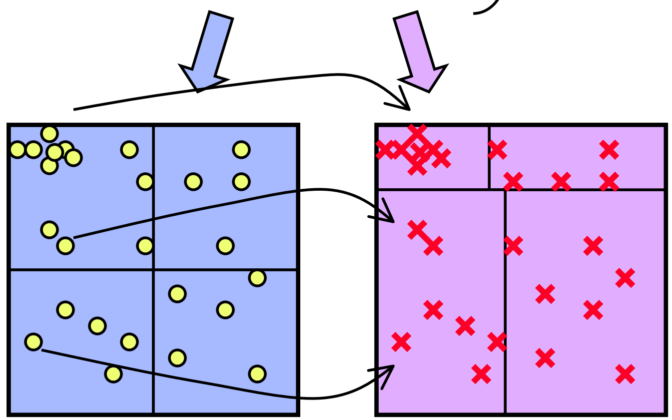

Need to balance target loads stresses flexibility of communication network

Framework will be used to derive benchmarks and metrics that drive development