



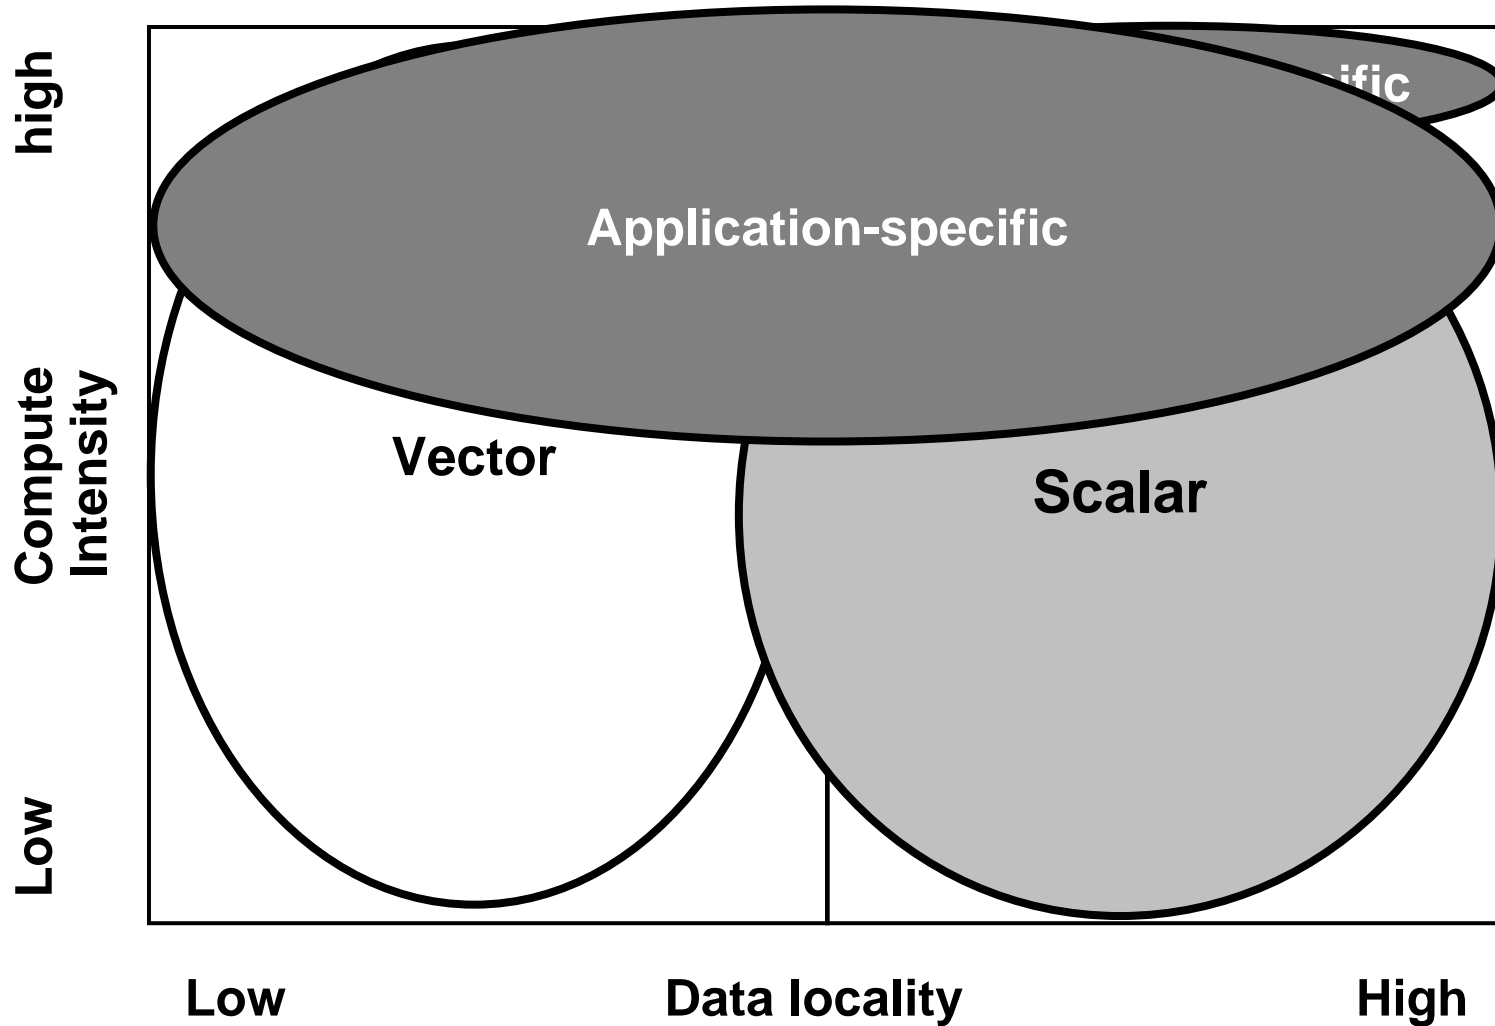
*Requirements for
Scalable Application
Specific Processing
in Commercial HPEC*

Steve Miller
Chief Engineer

The 3 Single-Paradigm Architectures

<u>Scalar</u>	<u>Vector</u>	<u>App-Specific</u>
Intel Itanium	Cray X1	Graphics - GPU
SGI MIPS	NEC SX	Signals - DSP
IBM Power		Prog'ble - FPGA
Sun SPARC		Other ASICs
HP PA		

Paradigms to Applications



Architectural Challenges

- Hardware

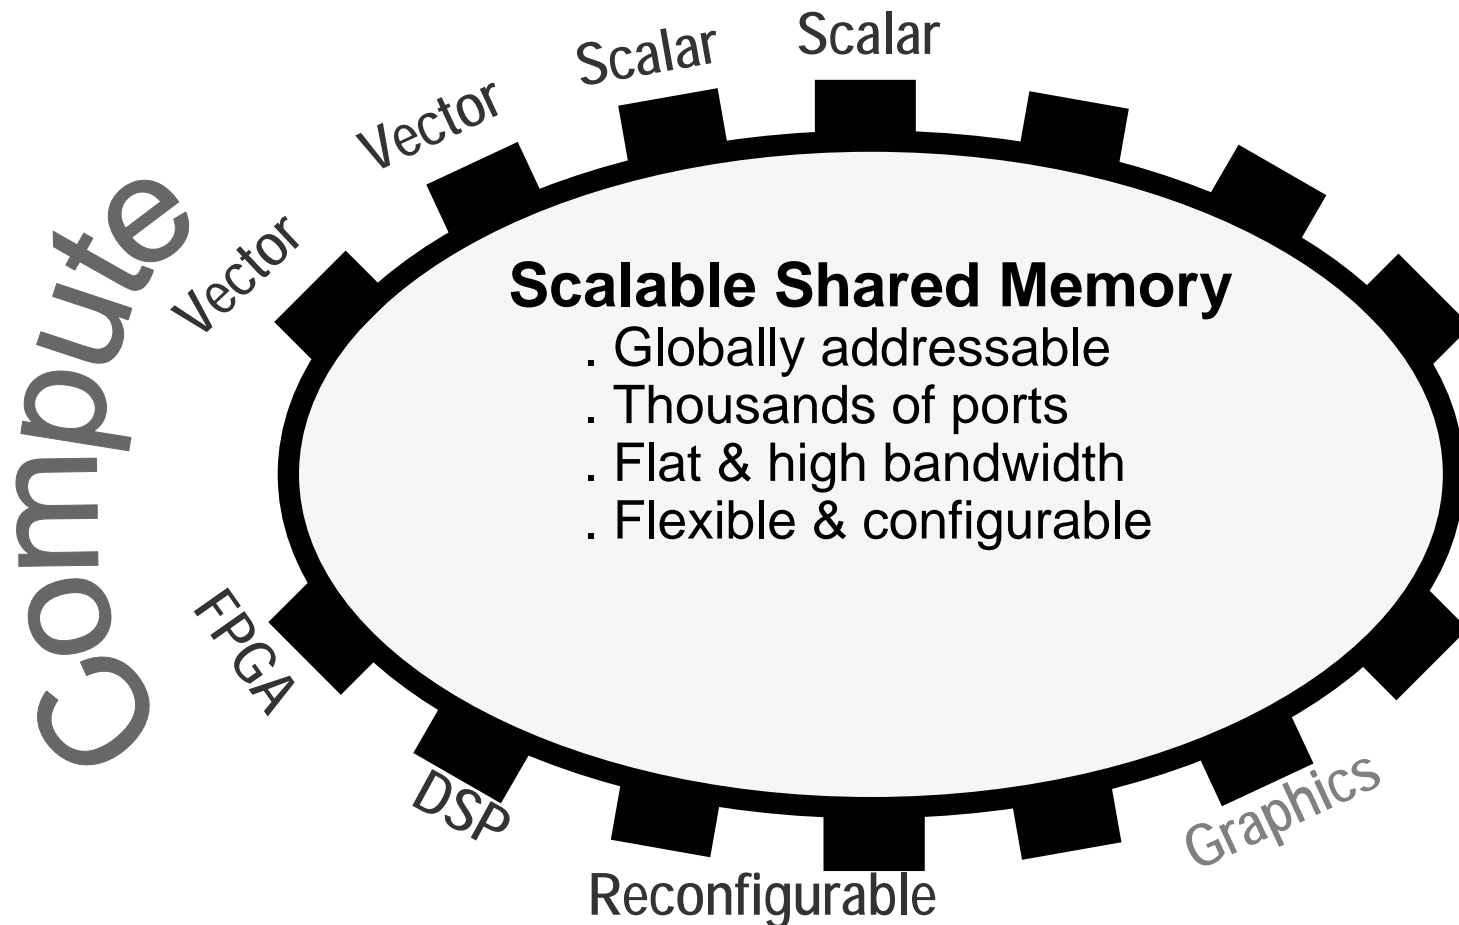
- Bandwidth to/from System
- Scalability

- Software

- Compilers/Languages
- Debuggers
- APIs

Multi-Paradigm Computing *UltraViolet*


Terascale to Petascale Data Set :
Bring Function to Data



Software

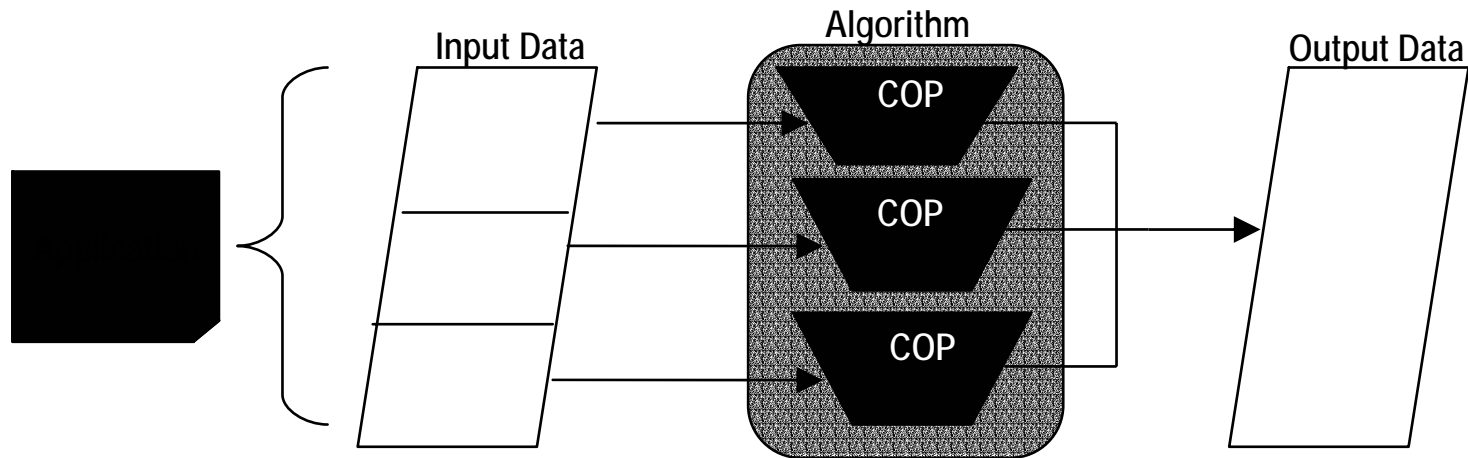
- **Provide for HDL modules**
 - Integrated environment with debugger
 - Highest performance
- **Leverage 3rd Party Std Language Tools**
 - Celoxia, Impulse Acceleration, Mitrion, Mentor Graphics
- **Developed an FPGA aware version of GDB**
 - Capable of debugging the FPGA and System Software
 - Capable of multiple CPUs and multiple FPGAs
- **Developed RASC Abstraction Layer (RASCAL)**

Software Overview

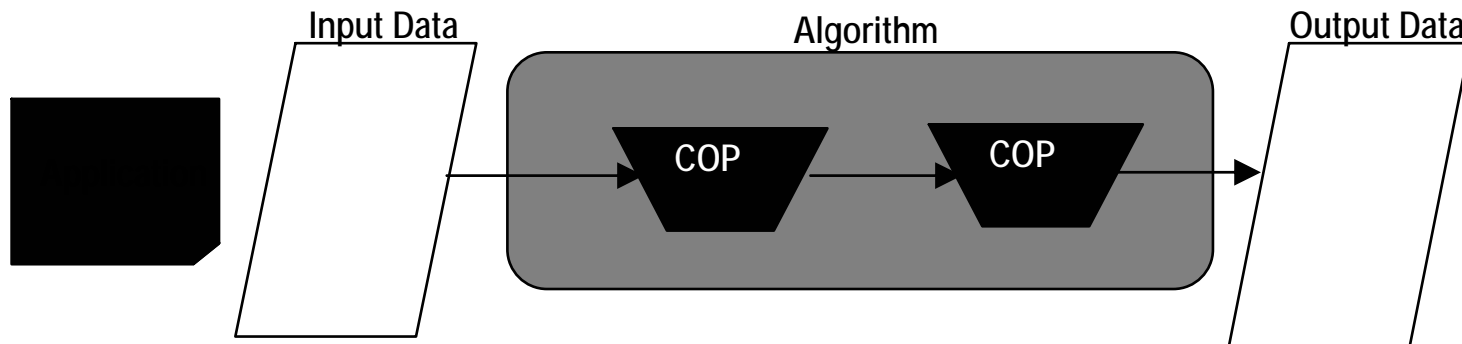


Abstraction Layer: Algorithm API

The Abstraction Layer's algorithm API mirrors the COP API with a few additions that enable wide scaling,



and deep scaling.



Hardware

- **Direct Connection to NUMALink4**
6.4GB/s/connection
- **Fast System Level Reprogramming of FPGA**
- **Atomic Memory Operations**
Same set as System CPUs
- **Hardware Barriers**
- **Configurations to 8191 NUMA/FPGA connections**

MOATB Block Diagram

NUMAlink 12.8 GB/s
SSP 6.4 GB/s
QDR SRAM 9.6GB/s
3 reads @ 1.6GB/s
3 writes @ 1.6GB/s

