

Session 5: Current & Emerging Standards

Craig Lund, Chief Technology Officer
Mercury Computer Systems, Inc.

**High Performance Embedded Computing (HPEC) Conference
September 2003**

The Ultimate Performance Machine

Agenda

30 min	Survey Introduction	Craig Lund
15 min	Real-Time Java	Martin Rinard
30 min	CORBA for HPEC	Bill Beckwith
15 min	Break	
15 min	UML 2 for HPEC	Jeffrey Smith
15 min	VXS	Henry Wong
15 min	Data Reorganization Interface (DRI)	Ken Cain

- ✍ **There are many HPEC papers on VS IPL this year – yesterday VS IPL had its own room and parallel track dedicated to those papers.**
- ✍ **At 4:00 p.m. today we will receive an update on the DARPA PCA program’s “Morphware” activity – an activity that may motivate new standards in adaptive and reconfigurable computing.**

Long, Boring List of Lists ...

- ✍ **Fabric standards**
- ✍ **Board standards**
- ✍ **Security standards**
- ✍ **Component models**
- ✍ **Parallel programming standards**
- ✍ **Library standards**
- ✍ **Application-specific standards**
- ✍ **Next horizon ...**

Fabric Standards

- ✍ **InfiniBand** is a highly-functional fabric optimized for box-to-box communications. It is popular in the supercomputer market where it competes with Myrinet.
- ✍ **10G Ethernet** is InfiniBand's competitor within the commercial domain. Some believe Ethernet will displace every other fabric on this slide.
- ✍ **RapidIO** is a chip-to-chip fabric designed to replace a microprocessor bus. Strong support from the people who make DSP chips.
- ✍ **PCI Express** is a serial version of PCI. It is an I/O model, not a peer-to-peer model. Strong support within the PC graphics industry.
- ✍ **Advanced Switching (AS)** is an Intel product proposal often confused with PCI Express. AS targets the same applications that RapidIO does.
- ✍ **HyperTransport** is the only fabric on this list that is shipping today in high volume. However, it is also the only fabric on this list that industry pundits appear to have written off as a long term contender in the embedded market.

VME Board Standards

- ✍ **VME has hit the performance wall. Successors are jockeying for military and aerospace mindshare.**
 - ✍ **VITA 41 (called VXS) avoids split backplanes**
 - ✍ **VITA 46 represents greater change**
- ✍ **Few VME suppliers have declared specific allegiance to 41 or 46.**
- ✍ **Another option is VITA 42 (called xMC) daughtercards.**

PCI Board Standards

- ✍ **The telecommunications downturn has impacted the PICMG world. The CompactPCI market is downsized and fragmenting.**
- ✍ **Advanced TCA moves technically far away from cPCI – it does not even have PCI in it. Historically Intel driven.**
- ✍ **Several efforts exist to patch the legacy cPCI infrastructure in various ways.**

Security

- ✍ **The government is supporting a major effort to define a new solution to multi-level security and safety-critical (D0-178B) challenges. The proposed solution impacts many parts of embedded systems.**
- ✍ **The initiative is called MILS. The open work appears in the OMG High Assurance Task Group. OMG is the standards authority for things like CORBA, UML, C4I and software radio.**
- ✍ **Bill Beckwith of OIS will soon provide us with a broad overview of OMG standards that impact HPEC. Jeff Smith will follow Bill with an update on the OMG's UML 2 revision's HPEC features.**

Component Models

- ✍ **Two major component models have emerged:**
 - ✍ **Microsoft's .Net**
 - ✍ **Java, which is complemented by OMG's CORBA components for languages other than Java.**
 - **An embedded personality for CORBA component "deployment and configuration" was just adopted. Bill Beckwith will likely mention it in his talk.**
- ✍ **The supercomputer community is doing its own thing called Common Component Architecture (CCA). We had a poster on CCA listed on early HPEC 2003 agendas, but the author needed to withdraw.**

Parallel Programming

- ✍ **POSIX Threads on SMP hardware**
- ✍ **Linda™ (Tuples) and JavaSpaces in Jini**
- ✍ **Message Passing Interface (MPI)**

DMA for Parallel Systems

- ✍ **Linked lists of DMA commands are closely associated with hardware and thus inherently proprietary. Busy waiting is a common synchronization technique associated with DMA.**
- ✍ **Everything in this domain has been proprietary.**

Remote DMA


- ✍ **The mainstream has recognized that DMA is a better paradigm for low-level network interfaces than standard sockets.**
- ✍ **Standard transports and APIs are emerging that enable the DMA paradigm over network fabrics.**
- ✍ **Arkady Kanevsky was invited to HPEC 2002 to provide a full 30 minutes on RDMA standards. The next two slides summarize his talk.**

Remote DMA Transports

- ✍ **Virtual Interface Architecture (VI)**
 - ✍ Just an “architecture” - no API or protocol
 - ✍ FC/VI and VI/TCP protocols emerged
- ✍ **InfiniBand**
 - ✍ Evolution of VI
 - ✍ Complete protocol with no API (just “verbs”)
 - ✍ Also licensed Windows Socket Direct protocol
- ✍ **IETF Remote Direct Data Placement (RDDP or iWARP)**
 - ✍ Evolution of InfiniBand RDMA architecture, applied to TCP
 - ✍ RDMA Consortium created a very complete protocol proposal and defined a superset of InfiniBand “verbs”

Remote DMA APIs

VIDF

-  Offers an API for VI based upon an example VI implementation from Intel
-  Standard not actively maintained, but products exist in the marketplace today.

DAT Collaborative

-  An evolution of VIDF with an open implementation
-  Targets InfiniBand and iWARP

The Open Group's Interconnect Transport API

-  Creating RDMA extensions to UNIX sockets
-  UNIX standardization path for DAT Collaborative work

Microsoft Windows Socket Direct

Data Reorganization

- ✍ **These APIs are focused specifically on parallelizing signal and image processing applications. They make data layout in memory a major focus of the API.**
- ✍ **Ken Cain will soon talk about standards work in this area.**
- ✍ **CORBA adopted these concepts in its optional “Data Parallel” API. However, so far, no commercial CORBA offers that API. The concepts are moving forward in OMG as part of “streams” for the CORBA component model.**
- ✍ **Ken Cain’s talk on DRI will summarize implementation options.**

Library Standards

- ✍ **A defunct array processor supplier named Floating Point Systems offered a library with its product. Their interface became the de facto standard in the HPEC industry.**
- ✍ **VSIPPL is a recent option. Programs are starting to use it. VSIPPL++ and Parallel VSIPPL are under development.**

- ✍ **Some government buyers create standards. Examples include:**
 - ✍ **Sonar middleware (MTM)**
 - ✍ **Software Communications Architecture (SCA)**

Next Horizon

- ✍ **FPGA computing is the next horizon for standards.**
- ✍ **Today FPGA investments are not very portable.**