

Putting Messaging Middleware in Perspective

Anthony Skjellum

MPI Software Technology, Inc.

Arkady Kanevsky

Mercury Computer Systems, Inc.

Ken Cain

The MITRE Corporation

James Lebak

MIT Lincoln Laboratory

Abstract

This talk is about comparing MPI, MPI/RT, and DataReorg, as well as potential contributions in the future from CORBA. The strengths and weaknesses of each effort from the perspective of these metrics will be offered:

- 1) achievable performance (latency, bandwidth, overhead, overlap)
- 2) portability
- 3) scalability
- 4) maturity in the SIP area
- 5) future technology path
- 6) ease of programming, training, maintaining codes
- 7) market penetration so far

Technical issues concerning the upcoming integration of DataReorg into MPI and MPI/RT extensions (the former de facto), as well as potential for MPI, MPI/RT, and DataReorg primitives in future CORBA will be given. This talk is meant to complement other talk proposals (e.g., from Arkady Kanevsky about MPI/RT for RACE, and from Ken Cain about DataReorg).

MPI/RT Standard Status and Plans

**Anthony Skjellum, MPI Software Technology
Inc.(Presenter)**

**Arkady Kanevsky (Mercury Computer Systems)
Dennis Cotel (SPAWAR Systems Center)**

20 September 2000

The MPI/RT Forum



<http://www.mpi-rt.org>

Join the mailing lists discussion!

MPI/RT 1.0 specification finished!

Goal: Final 1.1 specification by June 2001

- ◊ **Five year history, emerging from MPI 2 Standard**
- ◊ **Broad community participation includes:**
 - **FFRDCs and Government/Defense Laboratories**
 - **Defense integrators**
 - **Commercial embedded multicomputer vendors**
 - **Commercial HPC tool vendors**
- ◊ **Creating real-time parallel message passing API**

**What Problems Does
MPI/RT Try To Solve?**

MPI/RT is about QoS Models and High Performance Semantics

- ◊ **Channel based message transfer**
- ◊ **Planned Transfer Paradigm (Early Binding)**
- ◊ **Advanced models: event, priority/event, time**
- ◊ **Generation of middleout architectures for performance, QoS portable applications**
- ◊ **Addresses “Games’ Goal” of Zero Price of Portability**

Interface Maturity/Rollout

Long-term future: Integration with CORBA OOP?

Future Practice (tools and better s/w)

- Coordinates with Data Reorg services
- True QoS-based implementations
- Tools generate the MPI/RT programs
- Compute using VSIPL and BLAS

**Mechanisms
and policies
of MPI/RT
used by tools**

State of the Art (current standard APIs)

- Programmer defines channels manually
- Programmer manually redistributes data (MPI or MPI/RT)
- Coordinates with VSIPL usage [weakly]

**MPI/RT is
demanding
on the
programmer
but FAST!**

Challenges to Achieving Adoption of MPI/RT

MPI/RT has been implemented, but so what?

- ◊ **Three implementations have been done, all in “limbo”**
- ◊ **SIP customers have not had strong demand for MPI/RT**
- ◊ **MPI/RT is in demand by European/Japanese vendors wanting to compete in the system space**
- ◊ **Strong opportunity to use MPI/RT in the SAN space**
- ◊ **Alternatives: “use SMPs” or “MPI+DR is enough”**

MPI/RT Standard & Committee Status

MPI/RT Forum Plan

- ◊ **Three more official meetings - finish 1.1**
- ◊ **Several informal “working” meetings**
 - **Resolve issues with existing proposals**
 - **Resolve issues for new implementors**

- ◊ **Near-Term activities:**
 - **Clean up, publish MPI/RT 1.1**
 - **Offer further errata for 1.0**
 - **Promote MPI/RT 1.1 into new application areas**

The Forum is committed to completing MPI/RT 1.1, which is nearly complete as of now. The Forum then intends to explore other areas of interest in real-time message passing (“Charter issues”).