

**RBC Requirements
&
Specification Review**

July 17, 1990

NASA-JPL
C. Buzzell
S. Bellenot
R. Fujimoto
M. Robb

Agenda

- ⊗ Meeting Objectives
- ⊗ Specification Outline
- ⊗ Design Changes From Phase I Prototype
- ⊗ Specification Review
- ⊗ Commercial Multiprocessor Hardware Interfaces
- ⊗ GP1000 Availability
- ⊗ Time Warp Release Version

Objectives

- ⊗ Review and finalize RBC specification sections 1-4 (functional requirements).
- ⊗ Review section 5 of RBC specification and proposed software interface methodology.
- ⊗ Present results of multiprocessor vendor survey and proposed standard hardware interface.
- ⊗ Present status on GP1000 interface task and discuss available options.
- ⊗ Agree on version, and delivery of Time Warp version to be used for modification and incorporation into Phase II program.

Specification Outline

- 1.0 Scope**
- 2.0 Applicable Documents**
- 3.0 Functional Overview**
- 4.0 Functional Requirements**
 - 4.1 Maximum Total CMF State per Processor or Node
 - 4.2 Maximum Amount of State Allowed per Process
 - 4.3 Maximum Total State Memory per Processor or Node
 - 4.4 RBC Memory Allocation
 - 4.5 Maximum Number of Processes
 - 4.6 Minimum State Memory Size per Segment
 - 4.7 Dynamic Process Creation
 - 4.8 Memory Fragmentation
 - 4.9 Dynamic Process Destruction
 - 4.10 Dynamic State Size Growth/Shrinkage
 - 4.11 Number of Previous States Retained
 - 4.12 State Memory Expansion Increments
 - 4.13 Background and Foreground Tasks
 - 4.14 RBC Response Time
 - 4.15 VME Compatibility
 - 4.16 Supported VME Transfer Types

5.0 Interface Requirements

5.1 Hardware Interface Requirements

5.1.1 VME Interface

5.1.2 Non-VME Interface

5.1.3 Compatibility With Commercial Multiprocessors

5.2 Software Interface Requirements

5.2.1 RBC Memory Usage

5.2.2 Initialization Commands

5.2.3 Run - Time Commands

5.2.4 Example Command Sequences

6.0 Physical Requirements

6.1 Environmental

6.2 Power Requirements

6.3 Mechanical

Appendix A - Abbreviations Acronyms and Definitions

Changes From Phase I Prototype(20)

Spec. ¶	Description	Prototype Module	Current Design
4.3	Total Memory per processor Node	33 Mbytes	129 Mbytes
4.4	RBC Memory allocation	2 non-contiguous Blocks	Segments
4.5	Maximum Number of Processes	1024	256
4.11	Number of Previous States (frames) Retained	31	127
4.14.2	Rollback Method	Background	Foreground + WB read clearing
4.7	Dynamic Process Creation	Not Spec'd	Yes (segments)
4.9	Dynamic Process Destruction	Not Spec'd	Yes (segments)
4.10	Dynamic State Growth/Shrinkage	Not Spec'd	Yes (segments)

Number of Previous States(4.11)

WAS
31 Frames

Phase II
127 frames

- Assume granularity of 10 - 18 msec per event.
- Assume Mark (state save) occurs at the completion of each event.
- Assume GVT updates occur 1 per second (archiving and release of archived frames).
- Assume zero time associated with any processor operation except event processing (i.e. worst case for frame usage).

Then:

RBC needs to maintain sufficient frame reserve to allow the simulation to progress without exhausting all available frames.

$$\begin{aligned} \text{Number frames required} &= \frac{\text{Time between GVT update}}{\text{Mark Frequency (Granularity)}} \\ &= \frac{1}{10 \text{ msec}} \quad \text{to} \quad \frac{1}{18 \text{ msec}} \\ &= 55 \quad \text{to} \quad 100 \quad \text{frames} \end{aligned}$$

Total Memory per Processor or Node (4.3)

WAS

33 Mbytes

Phase II

129 Mbytes

- Follows directly from change in spec. section 4.11
- Maximum state per Node is 1 Mbyte
- Number of previous frames tracked is 127 instead of 31.

Then:

Total state memory required to be available is:

- | | |
|-------------------------------------|----------------|
| a) Current mark frame | 1 Mbyte Max |
| b) Previous frames | 127 Mbytes Max |
| 127 history frames * 1 Mbyte/ frame | |
| c) Pre-GVT state memory archiving | 1 Mbyte Max |

Total RBC Memory per node 129 Mbytes Max

RBC Memory Allocation (4.4)

WAS

2 Non-contiguous
state blocks

Phase II

Allocation based
on segments.

Required:

- Both JPL's and Jade's implementation of Time Warp utilize multiple, non-contiguous memory segments to define (allocate) a state vector. Also both utilize dynamic growth and shrinkage of the state vectors.
- JPL has a programmed limit of 100 such segments and current applications utilize them in the low 10's.
- Jade's implementation allows "any number" of state block allocations and dynamic memory segments.

Prototype Capability:

- Allowed up to 2 non-contiguous state blocks be defined for each process. State blocks could reside anywhere in the 32 bit address space of the node processor.
- RBC hardware performed the necessary address translation to map the state blocks into the WB memory.
- Each process required that its state block boundaries be defined.

2 words for each block * 2 blocks per process * up to 1024 processes = 4 K words

RBC Memory Allocation (Cont'd)

- Entire 32 bit address space covered on 1 Kbyte increments, 22 bits required for each definition. To perform the address comparison and translation required over 50 IC's to implement.
- Each process directly linked to its defined state blocks and the definitions fixed after initialization.
- Node required to inform the RBC of which process was executing (so the correct values for address translation could be loaded) through the use of the NEW PID command.
- Utilizing the prototype technique will not support the requirements for multiple segment definitions without excessive chip counts.

Phase II Capability:

- RBC memory allocation based on segment definitions which are not directly linked to the processes they serve.
- Processes may consist of multiple segments and the segments need not be contiguous in memory. A single segment cannot contain multiple processes.
- Segments must be at least 1 Kbyte in size and multiples of 1 K bytes.
- Up to 256 segments may be defined.
- Requires that the segment be identified for RBC commands (Mark, Rollback, Advance....) but not for reads and writes.

Maximum Number of Processes (4.5)

WAS

1024

Phase II

256

- Follows directly from paragraph 4.4 RBC memory allocation
- 256 segments allowed and each segment can be a separate process. Thus 256 processes per processor are permitted.
- Change to 256 based on:
 - a. Comments at the completion of Phase I indicated that 1024 processes utilized in Phase I was a high estimate and that T.W. would utilize substantially fewer processes.
 - b. Printed Circuit board area utilization and form factor: To accommodate existing processors which support VME I/F's the form factor for the RBC module has been switched from 9U X 400mm (14.5 X 15.7 inches) to 6U X 160 mm (10.3 X 7.2 inches). This represents a 65% reduction in board area.

Reducing segment ID width from 10 to 8 ($2^8 = 256$), allows single chip implementation of segment ID storage and configuration storage data path buffers instead of double thus reducing chip count.
 - c. Fujimoto has estimated that 64 processes per processor is adequate for most applications.
 - d. The need for adding the capability to have multiple segments and dynamic segments is perceived to have higher priority than number of processes

Rollback Method (4.14.2)

WAS

Background

Phase II

Foreground +
Read WB

Prototype:

- Relied on clearing all rolled back frames for ALL WB addresses in the current process.
- Rollbacks could not be queued and the current rollback must finish prior to either a Mark or additional Rollback.
- Implementation Advantages:

a. High speedup when rollbacks are infrequent :

Work load (clearing of the WB) only occurs when the rollback occurs (i.e. infrequently). If rollbacks are infrequent with respect to Marks then speedup versus Phase II method is maximized.

b. High speedup when processor " dead" time is available during the rollback or when long periods elapse after the RB command and Mark or additional RB:

Rollback is a background task and operates when other tasks are not actively using the RBC but are still processing.

c. Conceptually simple

Rollback Method (cont'd)

- Implementation Disadvantages:
 - a. Design complexity vs conceptual simplicity
 - b. High chip count to implement
 - c. If Mark or additional RB commands follow closely, implementation becomes an expensive foreground task.

Rollback Method (cont'd)

Phase II:

- Decrement CMF counter in foreground based on RB distance.
- Delay clearing of the WB memory bits until the addresses are read from the WB during normal foreground operations
- Implementation Advantages:
 - a. Rollback used data. Only addresses which are utilized (i.e. read) are updated. Time is not wasted rolling back unused WB addresses.
 - b. Combining of rollback commands. Addresses are updated when read, rollbacks are combined to give a total rollback distance for addresses which are accessed infrequently.
 - c. Sharing of WB accesses. Combining the rollback with the WB read, reduces the time to perform each operation independently.
 - d. Requires lower chip count than for background clearing
- Implementation Disadvantages:
 - a. Conceptually more difficult than background clearing
 - b. Requires insertion of an operation in the read/write path of the processor. This must be carefully designed to prevent a "decision hit" penalty on every read/write cycle.

Specification Review

Commercial Multiprocessor Hardware Interfaces

- Twelve (12) vendors surveyed

BBN Advanced Computers

Intel Scientific Computers

NCUBE

Cogent Research

Kendall Square Research

Paracom Inc.

Concurrent Computer Corp.

Meiko Scientific

Sun Microsystems

Encore Computer

Myrias Computer Corp

Topologix

- Two vendors failed to respond

Topologix

Kendall Square Research

- Very few support a standard bus configuration at the node level.

- Of those that do or are planned VME is the accepted standard.

- Processor Mix

Transputer	68020	80386	Other
Cogent Topologix Meiko Scientific Paracom	Myrias BBN GP1000 Sun 3	Intel IPSC/2	NCUBE Intel IPSC/860 Encore NS32332 Concurrent (prop) BBN TC2000 (88100) Sun4 (prop)

- Interface capabilities

Direct VME link available:	VME link available easily through adapter:	VME link difficult or performance Prob.
BBN TC2000 Paracom <u>Some Sun's</u>	Intel IPSC/2 BBN GP1000 Cogent	Intel IPSC /860 } DMA channel } long setup. Myrias Meiko Scientific } serial link NCUBE Encore Concurrent } shared memory


GP1000 Availability and H/W Interface

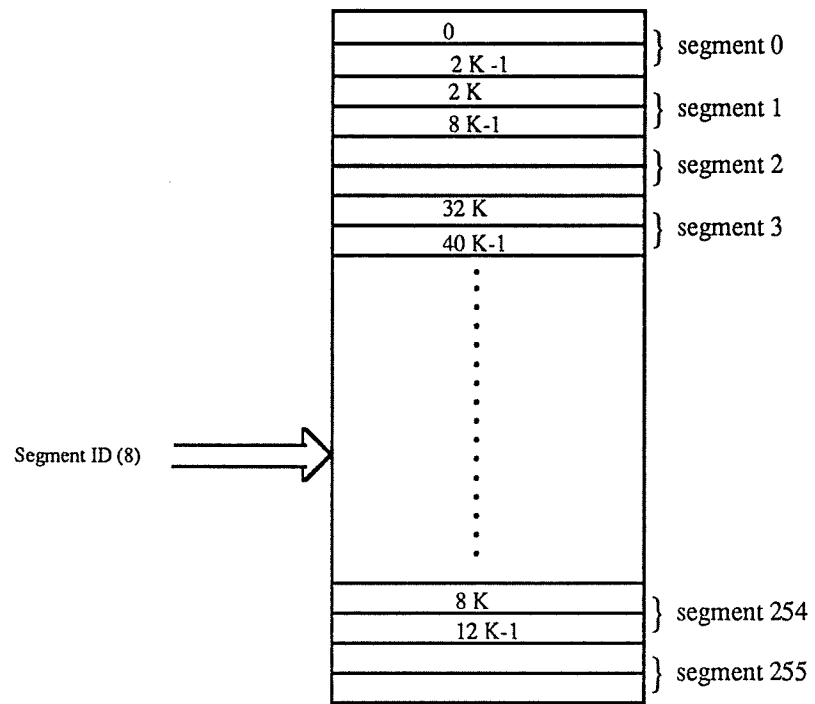
- GP1000 availability in Mountain View Facility agreed to by BBN is no longer possible due to closure of the field office.
- IPT currently waiting on BBN for designation of GP1000 and integration location to fulfil the agreement.
- IPT has engaged a local consultant to assist in resolution of switch and VME latency concerns on the GP. Consultant was a design engineer on the GP1000 node processor board (R. Quiros).

Time Warp Release Version

- Need to provide to IPT a copy of Time Warp which is to be the version designated as deliverable for the Phase II development. IPT will modify the TW code to interface it with the RBC hardware.
- Source modifications performed by IPT consultants and engineering staff.
- Need source listing and any available documentation
- Need test Time Warp application for the purpose of testing RBC and TW modifications.
- Date for delivery: open
- Media for delivery: open
- Delivery is direct to IPT for configuration control.

68020 'A' & 'B' DEFINITION





Segment Definition Table Example,

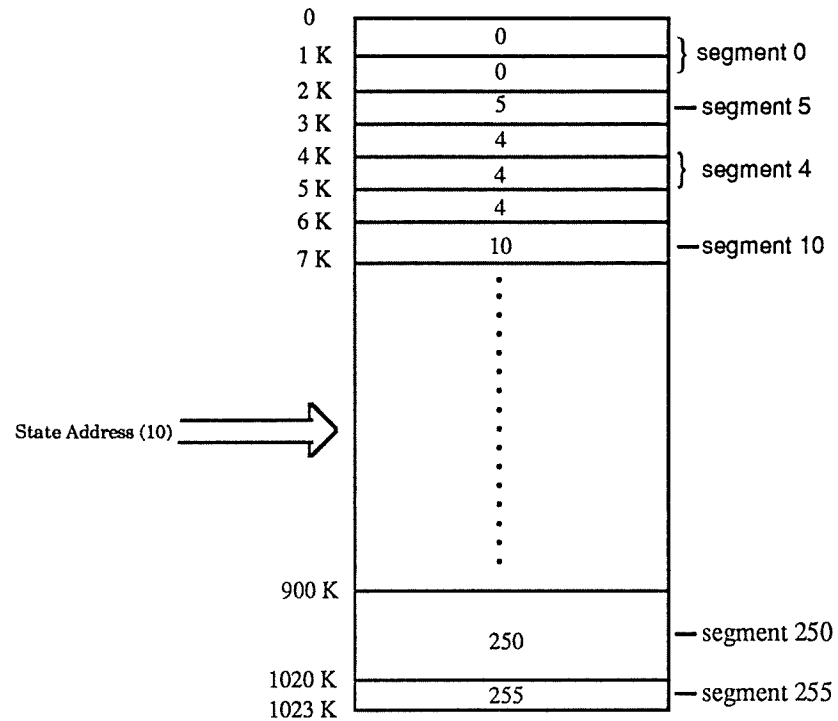


Figure 4, Segment Encoder Table Example