# LOAD BALANCING OBJECTS FOR TIME WARP ON A HYPERCUBE

Steven F. Bellenot

The Florida State University
Math and Computer Science
Tallahassee, Florida 32306

Originally, I was to develop a "benchmark" for the study of speed up of the time warp operating system running on a hypercube and study the load balancing question. Load balancing here means where (which processor) do you put your objects to obtain best speed up. However, the last half of my summer was spend debugging time warp code. (Perhaps as all projects, it was a bit behind schedule and help was needed.)

My benchmark program was about the dining philosophers. The dining philosophers is as close as any problem in computer science to being a classic. It is regular enough so that load balancing is easy. Thus parameters like the cost of message distance, communication vs computation and message length could be easily measured. However these tests can't be made until the time warp operating system is up and running. The code for this benchmark has been tested on the simulator.

As for load balancing in general there are two results. One is a general idea on how to measure success. The other is an object to node assignment for the COMMO* simulation, the main test case for the sponsors. In general, the question of load balancing is accepted to being a hard one. Thus any results are at least somewhat rewarding.

Debuging an operating system isn't just a summers project. It isn't done nor is it likely to be completed for several months. However, I think that I have made significant progress in this area. (Perhaps the opinions of co-workers would be of value here.) In anycase, I got time warp up to a point where its major functions of rollback, message annihilation and GVT calculations could be tested. In doing so I brought up the low level units of time warp on memory management, lists and queues.

# ON LOAD BALANCING OBJECTS FOR TIME WARP ON A HYPERCUBE

Steven F. Bellenot

The Florida State University
Department of Mathematics
Department of Computer Science
Tallahassee, Florida 32306

This summer's activities were preliminaries needed to make Time Warp availible for load balancing studies. A considerable amount of work was done on memory management, both from the design and from the implementation viewpoint. Some statistical work was also done, a bit on design, some preliminary runs and some timing results. Also the Dining Philosophers benchmark, which I designed and implemented last summer, was heavy used for debuging, for the first statistics tests, for the first speedup results and it was the first simulation to complete on the full half-hypercube.

Memory management is an important part of any operating system but is especially important on the Mark II Hypercube. From a design viewpoint, I wrote an interoffice memo suggesting new ways to garbage collect memory, heuristics for deciding when to use which garbage collection method, and ways of detecting deadlock caused by running simulations that require more memory than available on the Mark II. From the implementation side, I help to debug the "current version" of Time Warp. At the start of my debugging, Time Warp was like the old joke: "The memory goes first, and I forget what goes next." This memory death problem, as I called it, seems to be completely fixed now. I also helped debug other parts of Time Warp code. In particular, the message subsystem and timers (both in the so called MI) were fixed.

The dining philosophers simulation turned out to be very useful as a debugging tool. This benchmark was updated from last summer so that it could test a number of new things. (For examples: queries to self and a timekeeper object was added.) Preliminary statistics on the amount of VAX time needed to run the simulation were obtained. More statistics were taken on the number of rollbacks and the number of anti-messages sent on multi-node tests (also on the VAX). These Time Warp statistics indicated that our statistics gathering needed to be more carefully done. (For example there is more than one kind of rollback.) Finally, our first speedup results were obtained this summer using the philosophers.