PROGRAMMING TEMPLATES FOR ACTIVE STORAGES

Ekaterina Tyutlyaeva[1] and Alexander Moskovsky[2]

[1]Program Systems Institute of Russian Academy of Sciences,

Pereslavl-Zalessky, Russian Federation

Storing and processing data in a most efficient way is a major concern in every aspect of human activity. In this paper, we consider the Active Storage concept, implemented using the TSim template parallel programming library and the Lustre parallel file system.

The main idea of active storages is, in the context of parallel file systems, to use computational power of the storage node to process data located on this node. Processing data directly on the storage units can give a dramatic performance increase, by avoiding redundant transfers of data.

The TSim base scheduler provides support for different load balancing strategies and scheduling policies.

For our Active Storage system we developed the Active Storage policy class. User just needs to change the template parameter for the TSim scheduler's registration to change the strategy of the scheduler for the Active Storage policy and back.

```
//MyScheduler::Registrar<RoundRobin<TaskBase> >(NAME);
MyScheduler::Registrar<ActiveSched<TaskBase> >(NAME);
TSimRuntime rt(new MyScheduler);
```

For fault tolerance support we have implemented the modified scheduler and derived type T-variables (also known as futures) with timeout. This technique allows to complete the execution in a determined period of time even if a number of tasks has been uncompleted. Example of usage the T-variable with timeout:

```
TValT<int, N> fin;
 t = new ActiveTask(i, fin);
TSubmit(t);
```

Besides, we study testing results for our Active Storage system using reprojection remote sensing data task and compare it with the standard round robin scheduling policy. Study shows that the Active Storage system is effective even for a small grain of parallelism.

Finally, we suggest a cost model for estimate storing and processing costs for N duplicates in Active Storage environment. We analyzed a large amount of works on characterizing failures in computer systems to choose a statistical distribution that can provide the best fit.