

Department of Computer Science,  
University of Otago

UNIVERSITY  
of  
OTAGO



*Te Whare Wānanga o Ōtāgo*

---

Technical Report OUCS-2006-08

**View-Oriented Parallel Programming**

Author:

**Zhiyi Huang**

Department of Computer Science, University of Otago



Department of Computer Science,  
University of Otago, PO Box 56, Dunedin, Otago, New Zealand

<http://www.cs.otago.ac.nz/research/techreports.html>

# View-Oriented Parallel Programming

Z. Huang

Department of Computer Science  
University of Otago Dunedin, New Zealand  
Email:hzy@cs.otago.ac.nz

## Abstract

*Traditional parallel programming styles have many problems which hinder the development of parallel applications. The message passing style can be too complex for many programmers. While shared memory based parallel programming is relatively easy, it requires programmers to guarantee there is no data race in programs. Data race conditions are generally difficult to debug and difficult to prevent as well. The View-Oriented Parallel Programming (VOPP) is a novel shared-memory-based programming style. It removes the burden of checking data race conditions from the programmers. With the VOPP approach, shared data objects are divided into views according to the memory access pattern of the parallel algorithm. One advantage of VOPP is that programmers don't need to worry about mutual exclusion, such as locks, since mutual exclusion is automatically done by the underlying system when a view is accessed. By removing data races from programs, VOPP makes it easier to code and less difficult to debug programs. It also provides potential performance advantages on multi-core systems as well as cluster computers.*

**Key Words:** View-Oriented Parallel Programming, Cluster computer, Multi-core system, Message Passing Interface, Parallel Computing, Data Race, Deadlock

## 1 Introduction

With the advent of Sun's UltraSPARC T1, we are moving towards a new age of parallel processing. UltraSPARC T1 (also known as Niagara) is a multi-core system which consists of 8 cores and can support up to 32 processes. That means, inside the single chip, there could be 32 processes running in parallel. With conservative estimation, we expect hundreds of cores (processors) built into one chip. Therefore, with multi-core technology, we can have a super-computer built inside one chip.

What will happen in the near future is that every desktop computer will be a super-computer with huge computing power. The problem for us is how to utilize the power. The task eventually falls on the shoulders of application programmers. The programmers should be able to make their applications to run in parallel on multi-core systems. Parallel programming on multi-core systems is important for utilizing the available computing power.

However, parallel programming is regarded as difficult and error prone. Sound parallel programming methodology is needed to ease parallel programming. Traditionally, there are two camps in parallel programming methodologies. One is based on message passing such as Message Passing Interface, the other is based on shared memory which is used for communications between computing entities such as processes. Parallel programming with message passing is notoriously difficult and complex, especially when there are hundreds of processes communicating with messages. Using shared memory for communications between processes is natural and straightforward for parallel programmers, but the following problems hinder parallel programming using shared memory.

First, data race condition is difficult to debug. Data race means there are concurrent accesses to the same memory location and at least one of them is write access. If it happens in an execution of a parallel application, we say the application has a data race condition. Since a parallel execution is normally not repeatable, to find the cause of a data race is difficult. Second, deadlock makes debugging more complex. Deadlock is a situation where multiple processes/threads wait for each other due to competing for locks. To prevent data race conditions, mutually exclusive primitives such as locks are used. Improper use of locks can result in deadlocks. Also mutual exclusion has complicated the mental model of parallel programming, since the programmer not only needs to consider what data to access, but also to consider how to access them (exclusively or non-exclusively). Third, parallel applications are normally not portable. There is no popular standard for shared memory based programming. Many languages and systems are proposed, such as OpenMP, HPF, Linda, Pthread, and etc. To have

a standard API for shared memory based programming similar to MPI for message passing is very important for portability of parallel applications using shared memory. The API should be both efficient and convenient for parallel programming on different architectures and platforms such as SMPs, multi-core systems, and cluster computers.

This paper will address the above issues based on our novel View-Oriented Parallel Programming (VOPP). The rest of this paper is organised as follows. Section 2 describes the key concepts and primitives for VOPP. Section 3 illustrates VOPP with a few typical examples. Section 4 compares our work with other related work. Section 5 presents the performance results of several applications on cluster computers. Finally, our future work on VOPP is suggested in Section 6.

## 2 View-Oriented Parallel Programming (VOPP)

VOPP is a novel parallel programming style [6] based on the concept of view.

### Definition 1 *Definition of View*

- A view is a set of memory units (bytes or pages) in shared memory. Suppose  $M$  is the set of total units in shared memory and  $V_i$  is a view, then  $\forall V_i, V_i \subseteq M$ .
- Views do not overlap with each other. Suppose there are two different views  $V_i$  and  $V_j, i \neq j$ , then  $\forall V_i \forall V_j, V_i \cap V_j = \phi$

### Definition 2 *Properties of View*

- Views are created and destroyed dynamically.
- Each view has a unique view identifier.
- Before a view is accessed (read or written), it must be acquired (with view primitives); after the access of a view, it must be released (with view primitives).
- Multiple views can be merged together.

So far we have identified the following classes of views.

### Definition 3 *Classes of Views*

- Single-Writer view (SWV): only one process is allowed to acquire the view for write purpose at any particular time; multiple processes can acquire the view for read-only purpose at the same time. This class contains the following two special subclasses:
  - Consumable View (CV): This class of views is similar to a pipe. A writer (producer) produces the view while multiple readers (consumers) consumes the view. The producer and the consumers are synchronized with each other. Before a view is produced, the consumers should wait to consume the view; before the view is consumed by all consumers, the producer should wait to produce another version of the view.
  - Atomic View (AV): This class of views is accessed with atomic operations such as *read\_view* and *write\_view*.
- Multiple-Writer view (MWV): multiple processes can acquire the view for write purpose at the same time but they must write on different locations of the view.

### Definition 4 *Consistency of View*

- When a view is acquired by a process  $P_i$ , all previous *write* accesses to the view must be performed with respect to  $P_i$  according to their causal order.
- After a view is acquired by a process  $P_i$ ,  $P_i$ 's local copy of the view will not be affected by other processes until it acquires the view next time.

A write access to a unit of a view is said to *be performed with respect to* process  $P_i$  at a time point when a subsequent read access to that unit by  $P_i$  returns the value set by the write access. The above consistency is in conformity with the View-based Consistency (VC) proposed in [4].

There are a number of requirements for VOPP programmers.

- The programmer should use a number of views to store data objects according to the data sharing pattern of the parallel algorithm.
- Each view should consist of data objects that are always processed as an atomic set in the program.
- When any data object of a view is accessed, view primitives must be used (see below).

The following view primitives are needed to manipulate and access views:

- *int alloc\_view(int \*view\_id, int size, int flag)*: create the specified view by allocating a memory space with *size* and specifying its class with *flag*. The return value indicates if there is any error or not. If *view\_id* is null, the function will allocate an identifier to the view and return the identifier as its return value.
- *void \*view\_brk(int size)*: extend the view with size *size*, then return the pointer to the start of the extended area.
- *int free\_view(int view\_id)*: destroy the specified view and free the memory space of the view.
- *void \*acquire\_view(int view\_id)*: acquire exclusive write access to the specified view; the calling process is blocked if the view is held by another process. The address of the view is returned if the acquire is successful.
- *void release\_view(int view\_id, int nr)*: release the specified view. If the view is consumable, the number of consumers the producer waits is *nr*.
- *void \*acquire\_Rview(int view\_id)*: acquire read-only access to the specified view; the calling process gets an up-to-date version of the specified view. The address of the view is returned if the acquire is successful.
- *void release\_Rview(int view\_id)*: finish read-only access to the specified view.
- *int read\_view(int view\_id, void \*buf, int \*size, int offset)*: atomic read operation on views.
- *int write\_view(int view\_id, void \*buf, int size, int offset)*: atomic write operation on views.
- *void enqueue\_view(int view\_id)*: enqueue a view identifier into the system queue.
- *int dequeue\_view()*: dequeue a view identifier from the system queue.

VOPP allows programmers to participate in performance optimization through wise partitioning of shared data into views. The rule of thumb for VOPP overhead is that, the more view acquisitions, the more messages incurred in the network; and the larger a view is, the more amount of data transmitted in the view acquisition. The programmers will be able to finely tune the program by reducing both the number of view acquisitions and the size of views.

### 3 Examples of VOPP

In this section, we will use some examples to illustrate VOPP. These examples use a C interface provided by VODCA [8], a system supporting VOPP. VODCA provides the above view primitives with a prefix *Vdc\_*, though current version of VODCA only implements a subset of those view primitives. In addition, VODCA also provides the following C interface.

- *VDC\_NPROCS*: the maximum number of parallel processes supported by VODCA.
- *VDC\_NVIEWS*: the number of view identifiers available for use.
- *VDC\_NPAGES*: the number of pages in the shared memory.
- *Vdc\_nprocs*: the actual number of parallel processes in an execution.

- *Vdc\_proc\_id*: the process id, an integer ranging from 0 to *Vdc\_nprocs-1*.
- *void Vdc\_startup(int argc, char \*\*argv)*: initialise VODCA and start remote processes.
- *void Vdc\_exit(int status)*: terminate the calling process.
- *void Vdc\_barrier(unsigned id)*: block the calling process until every other process arrives at the barrier.

The following VOPP examples are based on typical problems in parallel programming.

### 3.1 Sum problem

In this problem, every process has its local array and needs to sum up all local arrays. We create a view for the subtotal of each local array in every process. Finally the master process (process 0) will sum up all subtotals.

```
int *local_array, a_size;
long *subtotal;

main(int argc, char **argv)
{
    int i, j;
    long sum=0;

    initialise a_size;

    Vdc_startup(argc, argv);

    local_array=malloc(a_size*sizeof(int));
    initialise local_array;

    Vdc_alloc_view(&Vdc_proc_id, sizeof(long), SWV);

    subtotal=Vdc_acquire_view(Vdc_proc_id);
    *subtotal = 0;
    for (i=0;i<a_size;i++)
        *subtotal+=local_array[i];
    Vdc_release_view(Vdc_proc_id, 0);

    Vdc_barrier(0);

    if(Vdc_proc_id==0){
        for(j=0;j<Vdc_nprocs;j++){
            subtotal = Vdc_acquire_Rview(j);
            sum += *subtotal;
            Vdc_release_Rview(j);
        }
        printf("The total sum is %l\n", sum);
    }
}
```

We may also use a multiple-writer view to write the following VOPP program.

```
if(Vdc_proc_id==0)
    Vdc_alloc_view(&Vdc_proc_id, Vdc_nprocs*sizeof(long), MWV);
```

```

Vdc_barrier(0);

subtotal=Vdc_acquire_view(0);
subtotal[Vdc_proc_id] = 0;
for (i=0;i<a_size;i++)
    subtotal[Vdc_proc_id]+=local_array[i];
Vdc_release_view(0, 0);

Vdc_barrier(0);

if(Vdc_proc_id==0){
    subtotal = Vdc_acquire_Rview(0);
    for(j=0;j<Vdc_nprocs;j++)
        sum += subtotal[j];
    Vdc_release_Rview(0);
    printf("The total sum is %l\n", sum);
}

```

### 3.2 Producer/consumer problem

Suppose there is a piece of shared data produced by process 0 and consumed by *Vdc nprocs* processes. We can use a consumable view to contain the shared data as below.

```

if(Vdc_proc_id==0)
    Vdc_alloc_view(&Vdc_proc_id, size, SWV|CV);

Vdc_barrier(0);

while(condition){
    if(Vdc_proc_id==0){
        vp=Vdc_acquire_view(0);
        produce(vp);
        Vdc_release_view(0, Vdc_nprocs);
    }

    vp=Vdc_acquire_Rview(0);
    consume(vp);
    Vdc_release_Rview(0);
}

```

### 3.3 Task queue problem

Using task queue for parallel computing is common. The shared data in each task is regarded as a different view and a unique identifier is allocated to the view when the task is created. The VOPP code pieces of the program are as below.

```

struct task {
    char state;
    char *task_data;
}

/* task producer */
struct task *t;

```

```

int vid;

vid = Vdc_alloc_view(NULL, sizeof(struct task), SWV);
t=Vdc_acquire_view(vid);
t->task_data=Vdc_view_brk(data_size);
create_task(t);
Vdc_release_view(vid);
Vdc_enqueue_view(vid);

/* task consumer */
struct task *t;
unsigned vid;

vid=Vdc_dequeue_view();
t=Vdc_acquire_view(vid);
consume_task(t);
Vdc_release_view(vid);

```

From the above examples, we see that VOPP does not place extra burden on programmers since the partitioning of shared data is an implicit task in parallel programming. VOPP just makes the task explicit by adding view primitives, which renders parallel programming less error-prone in handling shared data.

The focus of VOPP is shifted more towards data management (e.g. data partitioning and sharing), instead of mutual exclusion and data race as in traditional shared memory based parallel programming. Mutual exclusion is automatically achieved by VODCA system when a process acquires a view. The memory area of a view is protected by the system so that no other processes will be able to "touch" the area. In this way, the bug of "data race" is removed from VOPP programs.

To avoid deadlock, we are going to provide some wrapper functions to acquire multiple views together in the same order. The programmers can use these wrapper functions when they are not sure if there is a deadlock because of acquiring multiple views.

## 4 Comparison with other related work

The idea of combining data with mutual exclusion is not new. There were some related work such as Entry Consistency (EC) [2] and Scope Consistency (ScC) [9]. However, their programming interfaces are very different from VOPP.

VOPP is different from the programming style of Entry Consistency in terms of the association between data objects and views (or locks). Entry Consistency [2] requires the programmer to associate data objects with locks and barriers in programs, while the VOPP programmer simply creates views with *alloc view* primitive, which is just as easy as memory allocation.

VOPP is also different from the programming style of Scope Consistency (ScC) [9] Programs based on ScC are extended from the traditional DSM programs, i.e., lock primitives are normally used in those programs while scope primitives such as *open\_scope* are used only when required by memory consistency. Therefore, the programming model provided in ScC is a mixture. The programmer has to think of mutual exclusion when lock primitives are used, but has to think of memory consistency when scope primitives are used. This blended programming model simply confuses programmers. However, in contrast to the traditional DSM programs, the focus of VOPP is shifted towards shared data (views) rather than mutual exclusion. Programmers only think of shared data (views) when view primitives are used, while mutual exclusion and view consistency are left to the underlying system.

VOPP is very different from MPI. From programming point of view, VOPP is more convenient and easier for programmers than MPI, since VOPP is still based on the concept of shared memory (except the consistency of the shared memory is maintained according to views). In addition, VOPP provides experienced programmers an opportunity to finely-tune the performance of their programs by carefully dividing the shared memory into views.

Since partitioning of shared data into views becomes part of the design of a parallel algorithm in VOPP, VOPP offers the potential to make VOPP programs perform as well as MPI programs. The reason is that a VOPP program can be finely tuned so that its underlying message passing behavior can match that of its MPI counterpart. That is, if there is a finely-tuned MPI program, we can make a VOPP program whose underlying message passing behavior is similar to that of the MPI program.

The VOPP program can imitate the MPI program in a way that wherever there is data transfer between processors in the MPI program, the VOPP program allocates a shared view for the data and uses view acquisition instead of sending and receiving data. In this way, the overhead of message passing in VOPP can be almost the same as that in MPI program, since the cost of view acquisition in VODCA is almost the same as that of sending and receiving a block of data in MPI.

Though the message passing behavior of VOPP programs can be made similar to that of MPI programs, the programming interface provided in VOPP is very much different from MPI. MPI programmers have to know where a block of data is located, while location of a view is transparent to VOPP programmers. VOPP programmers only need to worry about which view to acquire, but not the location of the view.

## 5 Performance comparison

In this section, we present the performance results of several applications coded with MPI, VOPP and the traditional DSM programming style. The MPI applications are run on MPICH [3], the VOPP programs are run on VODCA, and the traditional DSM applications are run on TreadMarks [1].

The performance tests are carried out on an Itanium cluster connected by InfiniBand network. The cluster consists of 128 nodes running Linux kernel 2.4.21, 16 nodes of which are used for performance evaluation. We run two processes on each node, since each node has two 1.3 GHz Itanium 2 processors with 4 Gbytes of memory. The page size of the virtual memory is 16 KB. The C compiler used is GCC 3.2.3. All programs are compiled with "-O2" option.

We have tested VODCA, MPI, and TreadMarks on the cluster with Integer Sort (IS), Gauss, Successive Over-Relaxation (SOR), and Neural network (NN). Figures 1, 2, 3, and 4 have shown the speedups of the applications on the three systems.

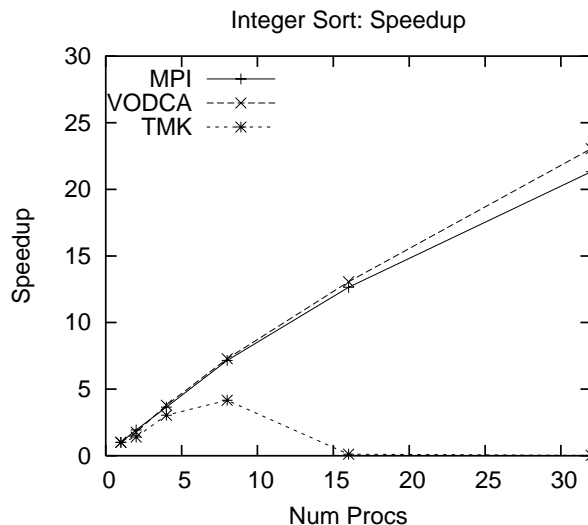


Figure 1: Speedup curves of IS

The figures show that the performance of VODCA is comparable with that of MPI. For some applications, VODCA performs even better than MPI. Since InfiniBand is a fast network with low latency, the barriers running on such a network is much faster than on Ethernet (refer to [7] for results on fast Ethernet). Faster barriers can improve the performance of both VODCA and TreadMarks, which explains why TreadMarks performs better for Gauss and SOR on InfiniBand. For applications like Gauss, there is still a small performance gap between VODCA and MPI due to thousands of barriers used in Gauss. Detailed reasons for this performance gap has been discussed in [7]. To avoid the thousands of barriers in Gauss, we can use a consumable view for sharing between the producer and the consumers in the program. We will publish the performance results in the near future.



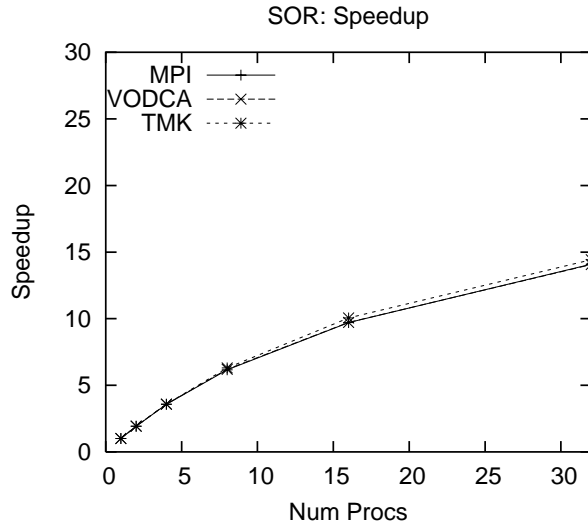


Figure 2: Speedup curves of SOR

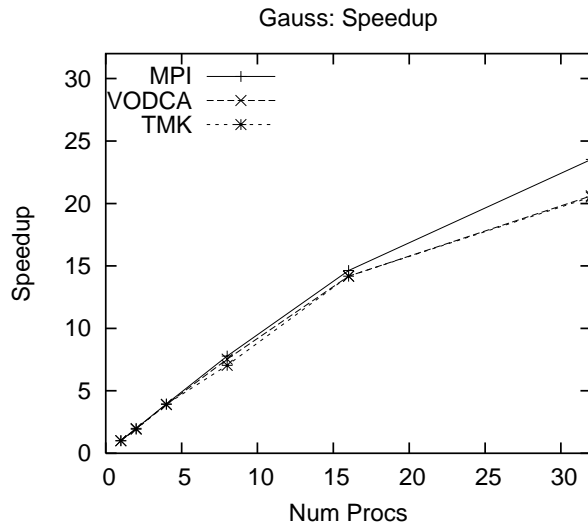


Figure 3: Speedup curves of Gauss

## 6 Future work

VOPP provides a convenient programming interface for parallel computing. It offers high performance comparable with that of MPI systems.

We will use more applications to investigate the convenience and performance of VOPP. Techniques for fast barriers need to be investigated when the number of processors is very large, e.g., the order of hundreds or thousands. In order to make VOPP convenient for parallel computing, a view-based debugger for VOPP programmers is needed in the near future.

## References

- [1] Amza, C., Cox, A.L., Dwarkadas, S., Keleher, P., Lu, H., Rajamony, R., Yu, W., Zwaenepoel, W.: TreadMarks: Shared memory computing on networks of workstations. *IEEE Computer* 29 (1996) 18–28
- [2] Bershad, B.N., Zekauskas, M.J.: Midway: Shared memory parallel programming with Entry Consistency for distributed memory multiprocessors. *CMU Technical Report (CMU-CS-91-170)* Carnegie-Mellon University (1991)

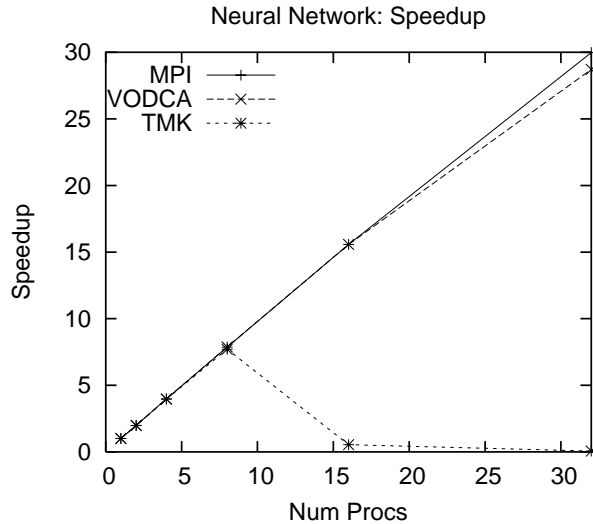


Figure 4: Speedup curves of NN

- [3] Gropp, W., Lusk, E., Skjellum, A.: A high-performance, portable implementation of the MPI message passing interface standard. *Parallel Computing* 22 (1996) 789–828
- [4] Huang, Z., Purvis M., and Werstein P.: View-Oriented Parallel Programming and View-based Consistency. In: *Proc. of the Fifth International Conference on Parallel and Distributed Computing, Applications and Technologies (LNCS 3320)* (2004) 505-518, Singapore.
- [5] Huang Z., Purvis M., and Werstein P.: View Oriented Update Protocol with Integrated Diff for View-based Consistency. *DSM Workshop 2005*, In: *Proc. of the IEEE/ACM Symposium on Cluster Computing and Grid 2005 (CCGrid05)*, IEEE Computer Society (2005)
- [6] Huang Z., Purvis M., and Werstein P.: Performance Evaluation of View-Oriented Parallel Programming. In: *Proc. of the 2005 International Conference on Parallel Processing (ICPP05)*, pp.251-258, IEEE Computer Society (2005)
- [7] Huang Z., Purvis M., and Werstein P.: Performance Comparison between VOPP and MPI. In: *Proc. of the Sixth International Conference on Parallel and Distributed Computing, Applications and Technologies*, pp.343-347, IEEE Computer Society (2005)
- [8] Huang Z., Chen W., Purvis M., and Zheng W.: VODCA: View-Oriented, Distributed, Cluster-based Approach to parallel computing, *DSM Workshop 2006*, In: *Proc. of the IEEE/ACM Symposium on Cluster Computing and Grid 2006 (CCGrid06)*, IEEE Computer Society (2006)
- [9] Iftode, L., Singh, J.P., Li, K.: Scope Consistency: A bridge between Release Consistency and Entry Consistency. In: *Proc. of the 8th Annual ACM Symposium on Parallel Algorithms and Architectures* (1996)
- [10] Keleher, P.: Lazy Release Consistency for distributed shared memory. Ph.D. Thesis (Rice Univ) (1995)