Eliminating Bottlenecks on Parallel Systems using Adaptive Objects

Masahiro YASUGI
Department of Communications and Computer Engineering
Graduate School of Informatics, Kyoto University
Sakyo Kyoto, JAPAN 606-8501
yasugi@kuis.kyoto-u.ac.jp

Shigeyuki EGUCHI
Graduate School of Science and Technology
Kobe University

Kazuo TAKI
Department of Computer and Systems Engineering
Kobe University

Abstract

Dynamic method replacement is a new technique to eliminate bottlenecks (e.g., around a root of tree structure) using adaptive objects for concurrent accesses. The technique eliminates the frequency of mutual exclusion and remote message passing by dynamically increasing the number of read only methods and the immutable part of objects. The results of performance measurements on both shared memory and distributed memory parallel architectures indicate the effectiveness of our approach to bottleneck elimination.

1 Introduction

The efficiency of parallel execution is often limited by the presence of bottlenecks in a small region of the entire system. In the object-oriented parallel computing, a major source of such bottlenecks is the mutual exclusion (serialization) of concurrent accesses to a single object (such as a root object of a tree structure).

Some bottlenecks can be eliminated (1) by relaxing the rule of mutual exclusion (e.g., by providing relaxed mutual exclusion which permits simultaneous read-only accesses to an object) or (2) by introducing *immutable* objects and permitting local access to their copies. The relaxation of mutual exclusion can be seen in [9]. The introduction of immutable objects can be seen in [8].

In the conventional object-oriented parallel programming languages, however, these techniques are applied only at the compile time, which makes lots of bottlenecks in the practical parallel applications still remain. For instance, when the application employs a network of objects (e.g., a binary search tree and a neural network) which is constructed by updating the states of the objects during the execution, the compile time approach does not work since the bottleneck can be eliminated only after the construction.

In this paper, we propose *dynamic method replacement* which provides adaptive objects for concurrent accesses, where we can increase the number of read-only methods and the immutable part of the objects dynamically. The proposed dynamic method replacement is realized by encapsulated message renaming, where the mechanism of class inheritance and polymorphism is naturally extended. We also present a free method detection algorithm, which is used to statically analyze the sufficient conditions of the runtime increase of *free methods* which only access to the immutable part of an object.

The rest of the paper is organized as follows. In Section 2, the brief overview of our object-oriented parallel language OPA is described. We illustrate the design of dynamic method replacement in Section 3. A free method detection algorithm is presented in Section 4. The implementation on a shared memory parallel computer is presented in Section 5, and that on a distributed memory parallel computer is presented in Section 6. The results of the performance measurements on both parallel architectures are presented in Section 7. In Section 8, 9, we discuss the related work and the effectiveness of our approach.

2 Object-oriented parallel language OPA

OPA (Object-oriented language for PArallel processing) is a parallel extension of Java language[2]; we remove specifications on threads and monitors from Java then add new constructs for structured synchronization and relaxed mutual exclusion. Its design is intended to realize both ease-of-use and high performance of parallel processing.

OPA supports irregular parallelism with dynamically forked threads; however, the fork and join of threads are performed in a structured manner. Objects are mainly utilized to represent complex data structure, where they can hold a reference to another object. The mutual exclusion is necessary for concurrent accesses (message passings) to an object to read/write the object's state consistently.

2.1 Free methods

The use of immutable objects can be easily extended to that of partly immutable objects: if a method only reads an immutable part of an object, the method may be invoked without mutual exclusion. Let us call such method *free*. On distributed memory architectures, exploiting free methods allows the language system to make a closure for a free method and distribute that copies to the other nodes to permit local invocation on each copy. The effect of using free methods is significant, since using local copies almost perfectly eliminates bottlenecks.

When the dynamic method replacement is employed, however, it is difficult for the compiler to confirm that a method is free, since every possible combination of methods must be examined. Thus, we must assume that the compiler may fail to detect free methods within a limited time and space complexity. We should not change the semantics of language due to the limitation of the compiler; we can not employ free methods *directly* in the language.

2.2 Mutual exclusion

Instead of free methods, non-blocking read-only methods (proposed in [9] for deadlock avoidance) can be employed in the language. The non-blocking method may read a mutable part of the object but it can be executed without infinite blocking. This is realized by updating the object's state atomically. Now, we can exploit free methods to optimize non-blocking methods without change of the language semantics.

To incorporate non-blocking methods, the OPA compiler automatically divides *instant methods* (defined with keyword `instant`) into read-write (RW) type and read-only (RO) type according to their literal access to the object. For instance, the `insertF` method and the `search` method in Figure 1 are of RO type and the `insert` method is of RW type. Keyword `internal` (Figure 1) is appended to the OPA language to indicate that the field (instant variable) or method is accessible only via (implicit) `this` reference.

Both RO and RW methods read the necessary variables of the object into implicit local variables atomically at the beginning of the method, while an RW method writes the values of local variables into the object atomically at the point where the rest of the method execution no longer updates

```
class BinTree{
 internal int key, value;
 internal BinTree left = null, right = null;
 public BinTree(int k, int val){
   key = k;
   value = val;
 }
 public instant int search(int k, int v){
   if(k == key)return value;
   else if(k < key){
     if(left != null)return left.search(k, v);
     else return v;
   }else{
     if(right != null)return right.search(k, v);
     else return v;
   }
 }
 public replaceable instant void insert(int k, int val){
   if(k < key){
     if(left != null)left.insert(k, val);
     else{
       left = new BinTree(k, val);
       if(right != null)setmethod(insert, insertF);
     }
   }else{
     if(right != null)right.insert(k, val);
     else{
       right = new BinTree(k, val);
       if(left != null)setmethod(insert, insertF);
     }
   }
 }
 internal instant void insertF(int k, int val){
   if(k < key)left.insert(k, val);
   else right.insert(k, val);
 }
}
```

Figure 1. Example of dynamic method replacement (binary tree dictionary)

the local variables. The compiler automatically determines the update point with flow analysis.

The RO method can read the state of an object, even if an RW method is running on the object. The RO method is prohibited only from reading the state which is partly updated. For an RO method and an RW method, respectively, the message passing proceeds as follows:

RO method: (RO1) lookup the method. (RO2) invoke the method. (RO3) atomically read the object's state into local variables. (RO4) execute the rest of the method.

RW method: (RW1) lookup the method. (RW2) invoke the method. (RW3) atomically read the object's state into local variables. (RW4) execute the method until the atomic update. (RW5) atomically write the updated local data into the object. (RW6) execute the rest of the method after the atomic update.

Let *RW section* be the section from RW1 to RW5, *RO section* be the section from RO1 to RO3, and *W section* be the

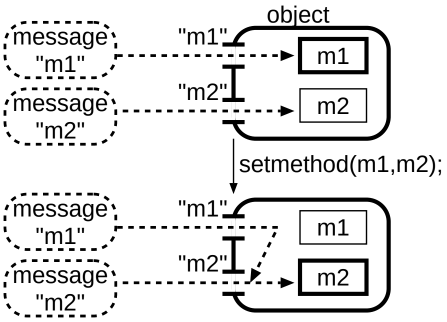

Figure 2. The message renaming by the dynamic method replacement

section of RW5. The mutual exclusion is necessary between an execution of RW section and another execution of RW section (Let us call this mutual exclusion *RW-RW mutex*). The mutual exclusion is also necessary between an execution of RO section and another execution of W section (Let us call this mutual exclusion *RO-W mutex*). Note that the method lookup is included in RO section and RW section due to the dynamic method replacement. By "mutual exclusion," we mean that the two or more threads may not be executed *in parallel*; that is, executions of high-level operations for different threads may not be interleaved. This means that the implementation-level operations may be executed in parallel as long as the semantics of higher-level mutual exclusion is preserved.

3 Dynamic method replacement

To confirm that a method is free, it must be confirmed that the method does read-only access to the object and that the read part will be never written by the other methods.

In practical applications, the object often has some RW methods which do write access to the object; however, after they are invoked finite times (usually for setting up the state of the object), it is often the case that (a) they are no longer invoked or (b) they no longer do write access to the object. If this property is confirmed by the compiler, there will be more occasions for exploiting free methods. Employing the dynamic method replacement makes it possible for the compiler to detect this property, since we can replace the RW method with an appropriate RO method if the method no longer does write access.

The dynamic method replacement can increase the number of RO methods. However, the dynamic method replacement changes the set of methods which can be invoked (or *method set* for short) dynamically. Thus, to confirm that the read part is immutable, it must be confirmed that the read part will be never written by any method the object may have in the future.

Of course, by the dynamic method replacement, we can simply replace an RW method with an RO method. The simple RO method does not eliminate bottlenecks as perfectly as a free method; but it eliminates some bottlenecks since it can be executed with less strict mutual exclusion than the RW method. For this reason, the dynamic method replacement is still effective even when it does not bring additional free methods.

For example, in a binary tree dictionary class (Figure 1), the `insert` method no longer updates the variables of the object once it creates both `left` and `right` child nodes. The `insert` RW method is replaced with the `insertF` RO method by the replacement statement:

```
setmethod(insert, insertF);
```

The replaced method, `insert`, is defined with keyword `replaceable` to indicate that the method is replaceable. After the dynamic method replacement of this example, since both `insert` and `search` methods are free and the dynamic method replacement is no longer performed, we should exploit the free methods. This property should be detected by the compiler, since there are only two possible method sets.

The effect of the dynamic method replacement can be explained more precisely in terms of *message renaming*. For example, let us consider the following replacement statement:

```
setmethod(m1, m2);
```

where method `m1` must be replaceable. The effect of the replacement is that message `m1` sent to the object will be renamed to `m2`. The renamed message `m2` is resent to the object for further method lookup (see Figure 2). In other words, when the object `obj`'s method is replaced, the message passing "`obj.m1(...)`" is executed as if it were "`obj.m2(...)`." If the method `m2` is also replaceable, the further renaming will take place.

When the class with replaceable methods is inherited, both the replaceable method and the substitutive method can be overridden as usual. Actually, the substitutive method may be simply a usual method. The late binding mechanism still works but it is extended: the method lookup is based on not only the actual class of the object but also the current method set (or renaming rules) of the object, and the result of method lookup can be not only a method but also a message which needs further method lookup. Since the inheritance and the late binding mechanism brings unknown uses of variables and messages by subclasses, keyword `internal`, `replaceable` and `instant` may not be omitted.

The update for the dynamic method replacement is performed together with the atomic update of the instance variables of the object. The mutual exclusion for the method lookup is required together with the mutual exclusion for the method invocation as was described in Section 2.2. An implementation issue is that the mutual exclusion for the RO method and that for the RW method differ but which type of method will be obtained is unknown at the moment of starting method lookup.

4 Free method detection algorithm

This section outlines an algorithm for the detection and exploitation of free methods.

The task of the method lookup can be regarded as a function which maps messages (ranged over by a) to methods (ranged over by m). As is described in Section 5, basically, a method table (T) is prepared for each class. If free methods are exploited, however, several method tables will be prepared for that class. Every T maps messages to table entries (ranged over by e):

$$T = \{a_1 \mapsto e_1, \ldots, a_n \mapsto e_n\}$$

The storage for the current state of each replaceable method is prepared in every object with that replaceable method. The location x of the storage is common to the subclasses. The table entry e is either of an actual method m with its type (RO, RW, or Free) or a location x:

$$e ::= \text{RO}\ m \mid \text{RW}\ m \mid \text{Free}\ m \mid \text{Replaceable}\ x$$

where 'Free$m$' is used only when a free method is exploited.

A method state is represented by $M = (T, S, H, R)$, where R is the renaming rule which maps the locations to the corresponding messages:

$$R = \{x_1 \mapsto a_1, \ldots, x_n \mapsto a_n\}$$

and $S(m)$ gives a set of possible method replacements (also denoted by R) by the execution of method m:

$$S = \{m_1 \mapsto \{R_{11}, \ldots, R_{1l_1}\}, \ldots, m_n \mapsto \{R_{n1}, \ldots, R_{nl_n}\}\}$$

and the set $H = \{a_1, \ldots, a_n\}$ consists of hidden messages which are introduced behind replaceable methods for the implementation. Method replacement R updates the renaming rule R_1 to R_2 (denoted by R_1R) by overwriting the relevant locations as follows:

$$\text{Dom}(R_2) = \text{Dom}(R_1) \cup \text{Dom}(R)$$
and
$$R_2(x) = \begin{cases} R_1(x) & \text{if } x \in \text{Dom}(R_1) - \text{Dom}(R) \\ R(x) & \text{if } x \in \text{Dom}(R) \end{cases}$$

The method lookup function *lookup* can be defined as follows:

$$\textit{lookup}(a, T, R) = \textit{rename}(T(a), T, R)$$

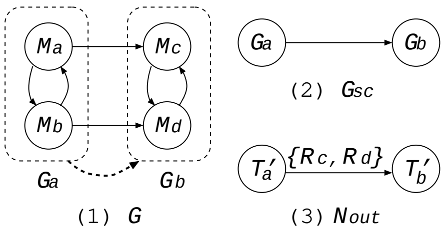

Figure 3. Graphs for free method detection algorithm

where,

$$\textit{rename}(\text{RO}\ m, T, R) = m$$
$$\textit{rename}(\text{RW}\ m, T, R) = m$$
$$\textit{rename}(\text{Free}\ m, T, R) = m$$
$$\textit{rename}(\text{Replaceable}\ x, T, R) = \textit{lookup}(R(x), T, R)$$

The intuitive meaning of the above definitions is pictured in Figure 4 and Figure 5.

Now, we can find a possible transition of method states $M_1 \longrightarrow M_2$ by the following definition:

$$\begin{aligned} & M_1 \longrightarrow M_2 \\ & \quad \text{iff} \\ & (\exists m \in \textit{MS}(T_1, H_1, R_1) \,.\, \exists R \in S(m) \,.\, R_2 = R_1R) \wedge \\ & T_1 = T_2 \ \wedge \ S_1 = S_2 \ \wedge \ H_1 = H_2 \\ & \text{where } M_1 = (T_1, S_1, H_1, R_1) \\ & \text{and } M_2 = (T_2, S_2, H_2, R_2) \end{aligned}$$

where $\textit{MS}(T, H, R)$ gives the invokable method set and is defined as follows:

$$\textit{MS}(T, H, R) = \{m \mid \exists a \in (\text{Dom}(T) - H) \,.\, m = \textit{lookup}(a, T, R)\}$$

With the above definitions, we can outline the free method detection algorithm as follows (see Figure 3):

[**Free method detection algorithm:**]

input: an initial method state $M_0 = (T_0, S_0, H_0, R_0)$

output: a directed acyclic graph (DAG) with transition conditions $N_{out} = (G_{out}, c_{out})$

where

G_{out}: DAG of the adaptive method tables

c_{out}: function which maps edges of G_{out} to transition conditions $\{R_1, \ldots, R_n\}$

$\{R_1, \ldots, R_n\}$: transition condition that the renaming rule is either of $R_1, \ldots, R_n$.

The meaning of the output: when the transition condition is satisfied, the method table can be switched alone with the edge of G_{out}.

1. Let directed graph $G_0 = (V_0, E_0) = (\{M_0\}, \{\})$.

2. To generate G (the graph of all possible method states), let $G_{n+1} = (V_{n+1}, E_{n+1}) = (V_n \cup \{M' \mid \exists M \in V_n.\ M \longrightarrow M'\}, E_n \cup \{(M, M') \mid \exists M \in V_n.\ M \longrightarrow M'\})$ repeatedly until $G = G_{n+1} = G_n$ (Figure 3–(1)).

3. Find all strongly connected components on G; let G_{sc} be the DAG of those components (Figure 3–(2)).

4. For each strongly connected component, find free methods. To save the amount of computation, the detection is performed from the child vertexes to parent vertexes. (a, m) is a free method if (1) m is of RO type, (2) for all present and future components, the result of $lookup$ by a is always m and (3) for all present and future components, there is no method (in method sets calculated by *MS*) which writes the instance variables read by m.

5. To prepare multiple method tables, for each strongly connected component, T is updated to $T' = T\{a_1 \mapsto$ Free $m_1, \ldots\}$ for all found free methods.

6. To know the conditions of table switches in advance, let G_{out} be the DAG (isomorphic to G_{sc}) of the updated T's and let $c_{out}(E) = \{R \mid (T\ , S\ , H\ , R)$ is either of terminal vertexes of the G edges which correspond to $E\}$ for edge E of G_{out}. (Figure 3–(3))

For the binary tree dictionary class (Figure 1), G in the above algorithm is as follows (ignoring `search` method for simplicity):

$$G = (\{M_0, M_1\}, \{(M_0, M_1)\})$$

where,

$$\begin{array}{l} M_0 = (T_0,\, S_0,\, H_0,\, R_0) \qquad M_1 = (T_1,\, S_1,\, H_1,\, R_1) \\ T_0 = T_1 = \ \{\texttt{insert} \mapsto \text{Replaceable}\ x_{\text{insert}}, \\ \qquad\qquad \texttt{insertF} \mapsto \text{RO}\ m_{\text{insertF}}, \\ \qquad\qquad a_{\text{insert}} \mapsto \text{RW}\ m_{\text{insert}}\} \\ S_0 = S_1 = \ \{m_{\text{insert}} \mapsto \{\{x_{\text{insert}} \mapsto \texttt{insertF}\}\}, \\ \qquad\qquad m_{\text{insertF}} \mapsto \{\}\} \\ H_0 = H_1 = \{a_{\text{insert}}\} \\ R_0 = \{x_{\text{insert}} \mapsto a_{\text{insert}}\} \qquad R_1 = \{x_{\text{insert}} \mapsto \texttt{insertF}\} \end{array}$$

In this example, each of M_0 and M_1 is a strongly connected component and finally we will obtain:

$$\begin{array}{l} N_{out} = (G_{out}, c_{out}),\ G_{out} = (\{T_0, T_1'\}, \{(T_0, T_1')\}), \\ c_{out} = \{(T_0, T_1') \mapsto \{R_1\}\} \end{array}$$

where,

$$\begin{array}{l} T_1' = \ \{\texttt{insert} \mapsto \text{Free}\ m_{\text{insertF}}, \\ \qquad \texttt{insertF} \mapsto \text{Free}\ m_{\text{insertF}},\ a_{\text{insert}} \mapsto \text{RW}\ m_{\text{insert}}\} \end{array}$$

Note that RW m_{insert} is no longer invoked since $a_{\text{insert}} \in H_1$ holds.

Since the space complexity of the given algorithm is exponential to the number of replaceable methods, the compiler may fail to detect free methods within limited computational resources. In such cases, the compiler will abort the detection and solely use the original method table.

5 Implementation for a shared memory architecture

Our target machine is the SGI's shared-memory computer, POWER Onyx (12 MIPS R10000 processors, 195MHz, 2GB memory). The OPA compiler generates C code, then the C compiler generates the binary code. The OPA runtime initially creates an OS thread (a process sharing the memory with other processes) on each processor. The OPA runtime non-preemptively schedules a number of OPA threads on the top of the fixed number of OS threads which will run actually in parallel on the same number of processors.

The RW-RW mutex is implemented with (a) a flag which indicates whether or not RW section is being executed by a thread and (b) a thread queue to which threads waiting for their turn are enqueued. To modify the flag and the queue atomically, lock/unlock operations are employed. In the actual implementation, the lock operation is deferred until the first method lookup is finished, and not performed if the obtained method is of RO type.

The RO-W mutex (to prohibit the RO method from reading the state which is partly updated) is implemented with a version numbering method: every object maintains a version number (initially 1), which is added by 2 with every update of the object and is set to 0 during the update. The reader retries if the version number is 0 (i.e., it is during the update) or it changes after reading. This method is a variant of Lamport's version numbering method[6] that requires two version counters. The benefit of this method is that it permits concurrent reads and does not require costly lock operations.

OPA objects are implemented with C structures composed of variables for mutual exclusion, a pointer to the method table for that class, and an instance variable table. Objects of its subclass are implemented with similar C structures which may have more methods and instance variables; but they keep, as usual, the shape of its superclass to realize subtyping with little overhead (see Figure 4). The status of method replacement is stored in the instance variable table.

Each entry of the method table for a class usually holds the method information that consists of a method type (`RO` or `RW`) and a function pointer for the method (see Figure 5). However, if the method is replaceable, the method information consists of the method type `Replaceable` and an offset for the replacement status in the instance variable ta-

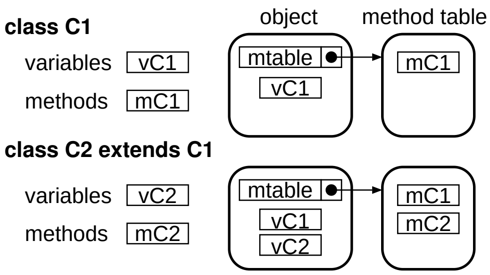

Figure 4. The representation of object

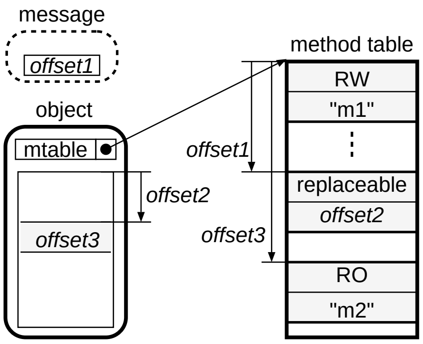

Figure 5. The method lookup for replaceable method

```
START:
do { vn=o.vn; } while(vn==0);
RETRY:
msg = message;
LOOKUP:
e = lookup(o, msg); method = e.method;
switch(e.mtype) {
case Replaceable:
  i = method;  msg = o[i]; goto LOOKUP;
case RO:
  if (vn && vn != o.vn) goto ROfail;
  status = invoke(method, vn, o, args);
  if (status == ROfail) {
ROfail:
    lock(o); vn=0; goto RETRY;
  }
  break;
case RW:
  if(vn){
    lock(o); if (vn != o.vn){vn=0; goto RETRY;}
  }
  if (o.use) { suspends(o); goto START; };
  status = invoke(method, o, args);
  break;
case Free:
  /* 0 means that lock has been acquired */
  if (vn == 0) unlock(o);
  /* 2 means that the method is free */
  status = invoke(method, 2, o, args);
  break;
}
```

Figure 6. Pseudo code of method lookup and method invocation for `o.message(args)` on shared memory architectures

ble. The replacement status consists of solely an offset in the method table for the substitutive method.

Figure 6 illustrates the pseudo code of the method lookup and the method invocation for a message passing expression of `o.message(args)`.

The thread first confirms that the object is not being updated (i.e., the version number `o.vn` is not zero) and memories the current version number in `vn`. Then the thread fetches a method table and lookups the method with an offset (*offset1* in Figure 5) given in the message. If the method is not replaceable, the thread will get the reference to the method code. However, if the method is replaceable, the entry holds an offset (*offset2*) for the replacement status in the instance variable table. This replacement status contains a new offset (*offset3*) for further method lookup. This process is repeated until a non-replaceable method is obtained.

If the obtained method is of RO type, the RO method can be invoked without acquiring the lock unless the update of the object is detected until the RO method finished the atomic read of instance variables. If an update is detected, the thread retries the same sequence from the method lookup but with the lock (`vn`=0 indicates this) to avoid further retries.

On the other hand, if the obtained method is of RW type, the thread locks the object and checks the use flag. If the object is in use by some other execution of RW section, the thread enqueues itself into the thread queue of the object and suspends its execution. If the object is not in use, the thread invokes the method which will set the use flag and release the lock.

6 Implementation for a distributed memory architecture

For a distributed memory architecture, we use a workstation cluster which consists of 4 Sun Microsystems SPARCstation5 (CPU: 110MHz microSPARC-II) workstations perfectly connected with each other by high speed serial communication interfaces called STAFF-link[7], which exhibits approximately 70$\sim$80μs latency for a small size ($\leq$ 64bytes) message.

The OPA runtime initially creates a process on each workstation. The OPA runtime non-preemptively schedules a number of OPA threads on those processes which will run actually in parallel on the same number of processors.

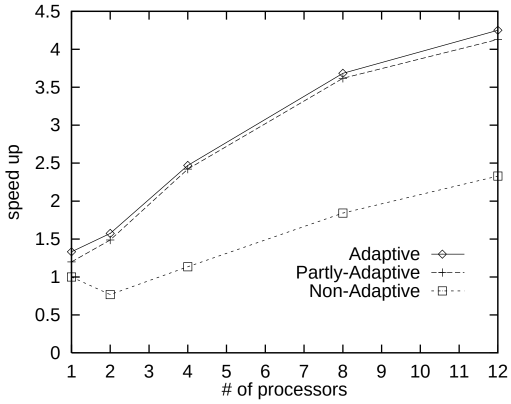

Figure 7. The speedup with dynamic method replacement (Adaptive/Partly-Adaptive) on a shared memory architecture (relative speedup to Non-Adaptive 1 PE case: the elapsed time is 0.464s for 30,000 nodes insertion to a shared binary tree.)

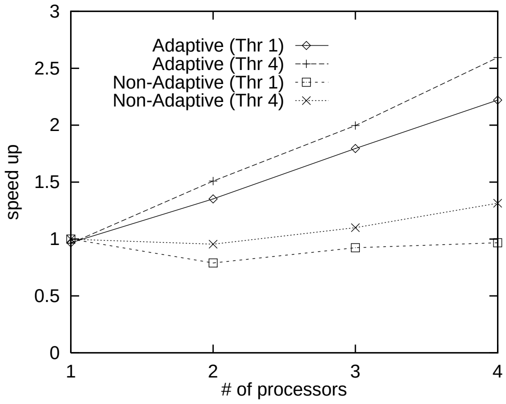

Figure 8. The speedup with dynamic method replacement (Adaptive) on a distributed memory architecture (relative speedup to Non-Adaptive 1 PE case: the elapsed time is 2.80s for 30,000 nodes insertion to a shared binary tree.)

Since the objects on a remote node is not directly accessible, either a remote method invocation to the object or a remote fetch of the object is required for the message passing to the object. We currently employ the former method.

The latency for the access to a remote object must be eliminated or hidden. The elimination of remote access can be realized by using free methods: the invocation of the free method can be performed on the local copy (proxy) of the remote object. The latency hiding is realized by using multiple threads on each processor. The OPA runtime manages the OPA (language-level) threads non-preemptively and very efficiently by performing the context switch fully by the compiled code.

The proxy is created when the reference to the corresponding remote object is received from the network for the first time. The proxy is created with the same size as the corresponding actual object. After that, the reference to the remote object, when it is received from the network, is converted into the reference to the proxy.

When the proxy's method is invoked, the proxy basically performs remote message passing to the remote actual object via a server on the remote node. When the server returns the return value to the proxy, the server also checks if the invoked method has been replaced with a free method: if it is a free method, the server attaches the data (method table and fields) of the actual object to the return value message. Once the proxy has the free method and the immutable copy of the fields, the subsequent method invocations will be performed entirely locally.

7 Performance measurements on parallel architectures

The performance measurements are performed with programs exploiting the binary tree class in Figure 1. We compared the following three programs. (Each program inserts 30,000 nodes with random keys into the shared tree structure):

Adaptive the program that exploits the dynamic method replacement and the resultant free methods (i.e., after dynamic method replacement, the `insert` method becomes free.),

Partly-Adaptive the program that exploits the dynamic method replacement but does not involve free methods (i.e., after dynamic method replacement, the `insert` method simply becomes of RO type.),

Non-Adaptive the program that has no replaceable methods.

Figure 7 shows the relative speedups to Non-Adaptive 1 PE case (the elapsed time is 0.464s and it is the 3.9 times slower than the corresponding C++ code that does not include any synchronization and context switching code.) for three programs on the shared memory architecture. The result shows the programs exploiting the dynamic method replacement is faster than the non-adaptive program for all 1 to 12 PEs. Since the speedup is achieved, we can say the bottlenecks are eliminated to some degree with the dynamic method replacement. The difference of Partly-Adaptive from Adaptive is not significant except for 1 PE case, because copying of the immutable parts of objects (which is achieved by software in the distributed memory architecture) is automatically performed by hardware cache.

For the distributed memory architecture, each program creates 1 or 4 thread(s) on each processor (for the latency hiding). Figure 8 shows the relative speedups to Non-Adaptive 1 PE case (the elapsed time is 2.80s.) for Adaptive and Non-Adaptive programs. The result shows the programs exploiting the dynamic method replacement is approximately 2 times faster than the non-adaptive program for 2 to 4 PEs. The result also indicates the latency hiding is effective on distributed memory architectures.

8 Related work

Actors[1] has a language construct `become`, which dynamically change the behavior of objects. This mechanism corresponds to a mechanism which switches the entire classes dynamically. Although OPA compiler exploits the switch of the entire method tables for free methods, it is too low level to be programmed directly. Furthermore, our technique allows to switch not only the invoked code of the method but also the invocation sequence including mutual exclusion and remote communication.

Concurrent Aggregates[4] introduces the notion of aggregate, a set of the same kind of objects, called representatives. When a message is sent to an aggregate, the message is sent to an arbitrarily chosen representative. This mechanism also eliminates bottlenecks around frequently accessed objects. OPA provides the same kind of object (such as root node of `BinTree` class) in a systematically confirmed way. Again, we think it is a little bit low level to program the aggregate directly.

Diniz and Rinard[5] use *dynamic feedback*, which is an adaptive technique that enables a compiler to automatically generate code that dynamically adapts to changes in the computing environment, in the context of a parallelizing compiler for object-based languages and to switch synchronization optimization policies. They have also proposed a set of synchronization transformations to reduce the synchronization overhead.

9 Discussion

Since the dynamic method replacement can be applied only to the replaceable methods, there is (virtually) no overhead for the other non-replaceable methods. The programmer should employ replaceable methods only when it is worth while doing so. Otherwise, two more memory references is incurred for the method lookup. This additional cost can be reduced by switching method tables prepared even for not-free methods with the additional switching costs.

The proposed technique is effective especially for data structures in which the immutable parts can be dynamically increased. The examples include a tree structure for the N-body simulation[3].

Since our target is bottlenecks whose parallelization is relatively difficult, the effect of the bottleneck elimination on the total speedup is strengthened by Amdahl's law.

10 Conclusion

In this paper, we proposed *dynamic method replacement* which provides adaptive objects for concurrent accesses, where we can increase the number of read-only methods and the immutable part of the objects dynamically to eliminate or relax the mutual exclusion and remote communication. This paper also presented its implementation and the results of the performance measurements on parallel architectures, which indicate the effectiveness of our approach.

The proposed feature was implemented and examined in our object-oriented parallel programming language OPA. Future work includes performance measurements using practical applications; we have already tested the N-body problem application using tree structure[3] and observed a considerable bottleneck elimination.

References

[1] G. Agha. *Actors : A Model of Concurrent Computation in Distributed Systems*. The MIT Press, 1987.

[2] K. Arnold and J. Gosling. *The Java Programming Language*. Addison-Wesley Publishing Company, 1996.

[3] J. Barnes and P. Hut. A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature*, 324:446–449, 1986.

[4] A. A. Chien. *Concurrent Aggregates*. The MIT Press, 1993.

[5] P. C. Diniz and M. C. Rinard. Dynamic feedback: An effective technique for adaptive computing. In *Proceedings of the SIGPLAN '97 Conference on Program Language Design and Implementation*, June 1997.

[6] L. Lamport. Concurrent reading and writing. *Communication of the ACM*, 20(11):806–811, 1977.

[7] H. Nakajo, A. Ichikawa, and Y. Kaneda. An Implementation and Evaluation of a Distributed Shared-Memory System on Workstation Clusters Using Fast Serial Links. In *Proc. of Int.Symp. on High Performance Computing(ISHPC)*, volume 1336 of *LNCS*, pages 143–158. Springer-Verlag, 1997.

[8] J. Sargeant. United functions and objects: an overview. Technical report, Department of Computer Science, University of Manchester, 1993.

[9] K. Taura and A. Yonezawa. Schematic: A concurrent object-oriented extension to Scheme. Technical report, Department of Information Science, University of Tokyo, 1995.