

TR-0901

Parallel Computation of Gröbner Bases on
Distributed Memory Machines

by

H. Sawada (MEL), S. Terasaki (Matsushita)
& A. Aiba

December, 1994

© Copyright 1994-12-6 ICOT, JAPAN ALL RIGHTS RESERVED

ICOT

Mita Kokusai Bldg. 21F
4-28 Mita 1-Chome
Minato-ku Tokyo 108 Japan

(03)3456-3191~5

Institute for New Generation Computer Technology

Parallel Computation of Gröbner Bases on Distributed Memory Machines

HIROYUKI SAWADA[†], SATOSHI TERASAKI[‡] AND AKIRA AIBA^{*}

[†]Mechanical Engineering Laboratory,
1-2 Namiki, Tsukuba, Ibaraki 305, Japan

[‡]Matsushita Electric Industrial Co., Ltd.,
4-5-15 Higashi-shinagawa, Shinagawa-ku, Tokyo 140, Japan

^{*}Institute for New Generation Computer Technology,
1-4-28 Mita, Minato Ku, Tokyo 108, Japan

November 29, 1994

Abstract

This paper reports our work on parallelizing an algorithm computing Gröbner bases on a distributed memory parallel machine. When computing Gröbner bases, the efficiency of computation is dominated by the total number of S-polynomials. To decrease the total number of S polynomials it is necessary to apply a selection strategy that selects the minimum polynomial as a new element of an intermediate base.

On a distributed memory parallel machine, as opposed to a shared memory parallel machine, we have to take into account non-trivial communication costs between processors. To reduce such communication costs, it is better to employ coarse grained parallelism rather than fine grained parallelism.

We adopt a manager-worker model. S-polynomials are reduced in worker processes in parallel, and the minimum polynomial is selected in the manager process. To implement the selection strategy in this parallel model, synchronization between worker processes is required for every selection of a new element of the intermediate base. However, in spite of synchronization, introducing the selection strategy produces not only a better absolute computation speed but also better speedup with multiprocessors. We achieved about 8 times speedup with 64 processors for large problems, T-6 and Ex-17.

^{†‡}This research was conducted at ICOT.

1 Introduction

Constraint logic programming (*CLP*) is a programming paradigm proposed by Jaffar & Lassez (1987) and Colmerauer (1987) and is an extension of logic programming. At ICOT, we have been researching CLP since 1987, and we have been developing two CLP languages, CAL (*Contrainte Avec Logique*), reported in (Aiba *et al.*, 1988), and GDCC (*Guarded Definite Clauses with Constraints*), reported in Hawley (1991) and (Terasaki *et al.*, 1992). CAL is a sequential CLP and GDCC is a parallel CLP.

CLP improves the descriptive power of a language by introducing a facility to handle relationships in certain domains other than syntactic equivalence of terms. To implement a CLP language, one has to implement a subsystem called a *constraint solver* to handle the extra relationships.

CAL and GDCC are both CLP languages with constraint solvers that have the ability to handle non-linear algebraic equations by employing the Buchberger algorithm (Buchberger, 1965, 1983) to compute the Gröbner base of given equations.

In our research and development of the GDCC parallel constraint solver, our major concern is the absolute speed of computing Gröbner bases by parallel processing. To parallelize the Buchberger algorithm, the absolute computation speed with a single processor must be first improved. Then, the speedup has to be improved with multi-processors. If we have a slow computation speed with a single processor, it is easy to provide deceptive speedup with multi-processors. However, the efficiency of the constraint solver is determined by the absolute speed, not by the speedup. Therefore, the speedup must be evaluated along with the absolute computation speed, even though some works regard the speedup as more important than the absolute computation speed.

This paper reports our work on developing the parallel algebraic constraint solvers with a manager-worker model for GDCC. They are implemented on a parallel inference machine *PIM/m* (Taki, 1992) developed at ICOT, which is a distributed memory machine consisting of 256 processors, and by using the kernel language *KLI* (Chikayama, 1992, Ueda & Chikayama, 1990) for the parallel inference machine. One of our parallel constraint solvers with the manager-worker model could calculate the Gröbner base for the T-6 problem in about 240 minutes with a single processor and in about 30 minutes with 64 processors on PIM/m, giving a speedup factor of about 8. On a Sun Sparc Server 490, it took about 90 minutes for the same problem by Backelin & Fröberg (1991).

The structure of this paper is as follows. In Section 2, we describe works related to parallelizing the Buchberger algorithm. In Section 3, our approach to parallelization of the constraint solver and the parallel model of the manager-worker are described, and the results of the experiments are listed.

2 Related works

Several attempts have been made to parallelize the Buchberger algorithm on shared memory machines by Melenk & Neun (1988), Vidal (1990), Buchberger (1987), and Clarke *et al.* (1990), and on distributed memory machines by Ponder (1988), Siegl (1990), and Senechaud (1989).

For shared memory machines, two parallelisms were implemented: coarse grained parallelism and fine grained parallelism. Coarse grained parallelism parallelizes reduction by reducing several S-polynomials simultaneously, while fine grained parallelism parallelizes reduction of a polynomial by dividing it based on the fact that only access to the leading monomial is necessary to control the Buchberger algorithm and that an S-polynomial is a linear combination of two polynomials.

In 1988, Melenk & Neun (1988) proposed fine grained parallelism and achieved about 2 times speedup on a two-processor CRAY X-MP. In 1990, Vidal (1990) implemented coarse grained parallelism based on the idea, mentioned by Buchberger (1987), of reducing several S-polynomials simultaneously. His parallel algorithm was implemented on a 16 processor *Encore* machine, and achieved 14 times speedup with 12 processors. In (Clarke *et al.*, 1990), these two techniques were combined and it was found that fine grained parallelism only worked for sufficiently large problems, such as Rose in (Boege *et al.*, 1986).

On the other hand, for distributed memory machines, fine grained parallelism, which parallelizes reduction of a polynomial by dividing it, was not implemented because of the non-trivial communication costs between processors. Ponder (1988) described three parallel algorithms: an algorithm to reduce several S-polynomials simultaneously (*Parallel S-polys*), an algorithm to parallelize the interreduction between polynomials in the intermediate base (*Parallel Reduction*), and an algorithm to solve a problem under different orderings among variables to see which ordering is fastest. He achieved 1 to 2 times speedup with 4 processors using the Parallel S-polys and Parallel Reduction algorithms. By using alternative orderings, he found that the execution time of the Buchberger algorithm is highly sensitive to ordering, and obtained at best that the fastest ordering was 66 times faster than the slowest. However, the ordering of the fastest computed Gröbner base and the ordering of a user's request will often be different. Therefore, the fastest computed Gröbner base is not always the base which the user really wants.

For Boolean Gröbner bases, Senechaud (1989) parallelized generation and reduction of S-polynomials by distributing polynomials of the intermediate base to processors which form a ring structure and by making subsets of the intermediate base circulate around the ring. She achieved 8 times speedup with 16 processors for a problem with 32 polynomials and 5 variables. However, it is not clear that her method is also effective for algebraic polynomials because of their complicated coefficients. In 1990, Siegl (1990) used a medium grain pipeline principle that parallelized reduction by making a pipeline of polynomials of

the intermediate base. That was implemented in STRAND88 on a transputer machine. He achieved 6 times speedup with 16 processors for a small problem.

Because our machine is a distributed memory machine and the communication costs between processors are not negligible, we also employ coarse grained parallelism to reduce several S-polynomials simultaneously and aim to solve relatively large problems efficiently. In the following section, we describe of our approach to parallel implementation of the Buchberger algorithm.

3 Parallelization

3.1 notation

The following notations are used in the following sections (Terasaki *et al.*, 1992, Hollman, 1992, Buchberger, 1983).

$Lm(f)$: Leading monomial of polynomial f .

$f \Rightarrow_r h$: Polynomial f is reduced to polynomial h by applying polynomial r to polynomial f once.

$f \downarrow_R$: Irreducible form of polynomial f w.r.t. polynomial set R .

$Spoly(f, g)$: S-polynomial of polynomials f and g .

$pp(f)$: Primitive part of polynomial f .

3.2 approach to parallelizing the constraint solver

Since reduction is the most time consuming part of computing Gröbner bases as described in Hollman & Langemyr (1991), we try to parallelize this.

The total number of S-polynomials generated during computation has a great influence on the total efficiency of the Gröbner base computation. Thus, it is very important for efficient parallel implementation to decrease the number of S-polynomials. The number of S-polynomials is determined by the series of leading monomials of elements of the intermediate base R generated during computation. Thus, a new element of R should be selected so as to decrease the number of S-polynomials generated during computation. As described in Buchberger (1983), it is well known that a critical pair (f, g) is not necessary if the greatest common divisor of $Lm(f)$ and $Lm(g)$ is 1. To decrease the number of generated S-polynomials, the minimum polynomial which has the minimum leading monomial should be selected, because a smaller monomial is apt to be the prime to other monomials rather than to a larger one. We should therefore choose the global minimum polynomial as a new element of R .

Furthermore, we have to take into account the fact that our parallel inference machine PIM/m is a distributed memory machine. Unlike a shared memory

parallel machine, communication between processors is not negligible compared to the computation on a processor. For this reason, on PIM/m, relatively coarse grained parallelism is more suitable than fine grained parallelism to reduce the amount of communication between processors.

We implement an algebraic constraint solver on a manager-worker model.

3.3 manager-worker model

In the manager-worker model, reduction is parallelized by partitioning a polynomial set to be reduced. To decrease the communication cost between the manager and worker processes, all worker processes have the same intermediate bases on their own memories. They reduce polynomials in parallel and report the local minimum polynomials to the manager process. The manager process then chooses a global minimum polynomial from among these.

Figures 1 and 2 show our manager-worker model. The set of polynomials is partitioned and each worker process has a different subset W_i . The initial Gröbner base R_{init} is copied to all worker processes. New input polynomials are distributed to the worker processes by the manager process so that all worker processes have the same number of polynomials.

In order to choose a global minimum polynomial, it is necessary to reduce only the leading monomials of all polynomials by the intermediate base. To add the selected polynomial to the intermediate base, however, the selected polynomial must be reduced completely, while other polynomials need not be reduced completely. Thus, each worker process reduces polynomials by R_i according to the following reduction stages.

1. A polynomial is not reduced.
2. The leading monomial of a polynomial is reduced.
3. A polynomial is completely reduced to an irreducible form.
4. A polynomial is reduced to its primitive part.

All input polynomials and S-polynomials are input to the first stage. These polynomials are sent to the second stage after reducing the leading monomial completely by the function $ReduceLm(W_i, R_i)$. At the second stage, each worker process selects the local minimum polynomial, sends a copy of the local minimum polynomial to the manager process, and reduces the local minimum polynomial according to stages 3 and 4. After going through the reduction stages, the local minimum polynomial becomes the local candidate for the new element of R .

The manager process receives local minimum polynomials from all worker processes, selects the global minimum polynomial from among these, and sends a *global-minimum-message* or *not-global-minimum-message* to each worker process. When a worker process receives a *global-minimum-message*, it should send

(1)	input $F := F_{init}$	(24)	$F := F \cup \{f_i\}$
	% F_{init} is a set of input polynomials.	(25)	$f'_i := f_i \downarrow R_i$
(2)	input $R := R_{init}$	(26)	$W_i := W_i \cup \{f'_i\} \setminus \{f_i\}$
	% R_{init} is \emptyset or an initial Gröbner base.	(27)	$W := W \cup \{W_i\}$
(3)	input $n := n_{inst}$	(28)	else $f_i = 0$
	% n_{inst} is the number of worker	(29)	endif
	% processes.	(30)	endfor
(4)	$W := \emptyset$	(31)	$f := Choose(F)$
	% W is a set of polynomial sets.	(32)	for $i = 0$ to $n - 1$
(5)	for $i = 0$ to $n - 1$	(33)	if $f_i = f$ then
(6)	$R_i := R$	(34)	$W := W \setminus \{W_i\}$
	% R_i is i th worker's intermediate	(35)	$W_i := W_i \setminus \{f'_i\}$
	% base.	(36)	$W := W \cup \{W_i\}$
(7)	$W_i := \emptyset$	(37)	$h := pp(f'_i)$
	% W_i is i th worker's polynomial set.	(38)	endif
(8)	$W := W \cup \{W_i\}$	(39)	endfor
(9)	endfor	(40)	if $h = 1$ then
(10)	$F := ReduceLm(F, R)$	(41)	return $R = fail$
	% Reduce leading monomials of	(42)	endif
	% all polynomials in F by R .	(43)	$W := LoadBalance(R, W, h, n)$
(11)	$m := Count(F)$		% Generate S-polynomials
	% Count polynomials in F .	(44)	$R = R \cup \{h\}$
(12)	for $j = 0$ to $m - 1$	(45)	$m := 0$
(13)	$f := First(F)$	(46)	for $i = 0$ to $n - 1$
	% f is the first polynomial in F .	(47)	$R_i := R_i \cup \{h\}$
(14)	$F := F \setminus \{f\}$	(48)	$W_i := ReduceLm(W_i, R_i)$
(15)	$i := j \bmod n$	(49)	$m_i := Count(W_i)$
(16)	$W_i := W_i \cup \{f\}$	(50)	$m := m + m_i$
(17)	endfor	(51)	endfor
(18)	while $m \neq 0$	(52)	endwhile
(19)	$F := \emptyset$	(53)	$R := Interreduce(R)$
(20)	for $i = 0$ to $n - 1$		% Interreduce between elements of R .
(21)	if $W_i \neq \emptyset$ then	(54)	return R
(22)	$W := W \setminus \{W_i\}$		
(23)	$f_i := Choose(W_i)$		

Figure 1: Manager-worker model

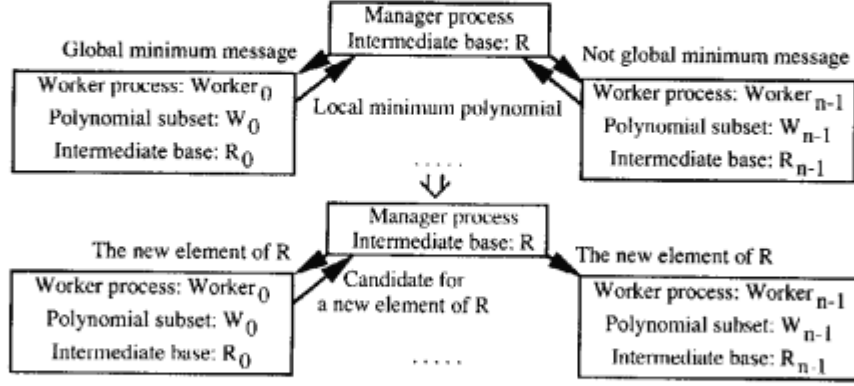


Figure 2: Architecture of the manager-worker model

its local candidate to the manager process. The manager process receives the local new-element-candidate, adds it to R , and sends it to all worker processes. Each worker process receives the new element of R , adds it to R_i , then generates critical pairs and S-polynomials.

In the manager-worker model with such synchronization, idling time exists in worker processes. In our implementation, to reduce such idling time, each worker process executes a subtask after computing its local candidate for a new element of R until it receives a global-minimum-message or the new element of R from the manager process. The main task is computation of the local candidate, and the subtask is the reduction of other polynomials, that is, the reduction of monomials other than the leading monomials. Since such subtasks are not always executed, these subtasks may cause *nondeterminacy* in our parallel algorithm, that is, the sequence of generating elements of R and the total number of generated S-polynomials are not the same for every execution. In our parallel algorithm, if several local candidates for the new element of R have the same leading monomials, then the manager process compares other monomials in those local candidates. Since the result of comparison by the manager process depends on whether such monomials of these local candidates have been reduced previously, the manager process does not always choose the same candidate at every execution, and nondeterminacy may occur. However, these subtasks may decrease the subsequent load.

<pre> (1) $B := \emptyset$ % B is a load balance information number set. (2) $M := \emptyset$ % M is a set of number of polynomials. (3) for $i = 0$ to $n - 1$ (4) $B := B \cup \{i\}$ (5) $m_i := \text{Count}(W_i)$ (6) $M := M \cup \{m_i\}$ (7) endfor (8) for $i = 0$ to $n - 1$ (9) $B' := B$ (10) $B := \emptyset$ (11) $M' := M$ (12) $M := \emptyset$ (13) $b_i := n$ (14) while $b_i = n$ (15) $b := \text{Minimum}(B')$ % b is the minimum number in B'. (16) $B' := B' \setminus \{b\}$ (17) $m := \text{Minimum}(M')$ (18) $M' := M' \setminus \{m\}$ (19) if $m_i = m$ then (20) $b_i := b$ (21) $B := B \cup B'$ (22) $M := M \cup M'$ </pre>	<pre> (23) else (24) $B := B \cup \{b\}$ (25) $M := M \cup \{m\}$ (26) endif (27) endwhile (28) endfor (29) for $i = 0$ to $n - 1$ (30) $W := W \setminus \{W_i\}$ (31) $C_i := \text{CriticalPairs}(R_i, h)$ % Make critical pairs. (32) $m := \text{Count}(C_i)$ % Count critical pairs in C_i. (33) for $c = 0$ to $m - 1$ (34) $(f, h) := \text{First}(C_i)$ (35) $C_i := C_i \setminus \{(f, h)\}$ (36) $j := c \bmod n$ (37) if $j = b_i$ then (38) $s := \text{Spoly}(f, h)$ (39) $W_i := W_i \cup \{s\}$ (40) endif (41) endfor (42) $W := W \cup \{W_i\}$ (43) endfor (44) return W </pre>
---	---

Figure 3: Load balancing

3.4 load balancing

For parallel processing, it is necessary to balance the load on each processor. Since each worker process reduces its own polynomials, we should balance loads according to the reduction cost. It is quite difficult, however, to predict the reduction cost for every S-polynomial beforehand. Thus, we balance loads so that each worker process has the same number of polynomials.

The most naive method for this load balancing is for the manager process to generate all S-polynomials and distribute them to worker processes. But this method increases the communication cost between the manager and worker processes, and results in a communication bottleneck.

Since each worker process has the same intermediate base R_i , it can generate the same critical pairs in the same order. Based on this fact, we have implemented a method where each worker process generates critical pairs in parallel, as follows. Figure 3 shows the load balancing executed by the function $\text{LoadBalance}(R, W, h, n)$ in Figure 1. Each worker process reports the number




Figure 4: Example of load balancing

of polynomials it has to the manager process. The manager process decides on a load balance information number b_i for each worker process, and sends that number to them. Each worker process assigns a critical pair number, c , to each generated critical pair, where c is common to all worker processes. Then, each worker process generates S-polynomials from such critical pairs such that $c \bmod n = b_i$.

Figure 4 shows an example of load balancing. There are three worker processes, $Worker_0$, $Worker_1$, and $Worker_2$.

The total number of worker processes, n , is 3. $Worker_0$ has 3 polynomials, $Worker_1$ has 2 polynomials, and $Worker_2$ has 4 polynomials. The processes report the number of polynomials they have to the manager process. The manager process decides on a load balance information number b_i . In this case, $b_0 = 1$, $b_1 = 0$, and $b_2 = 2$. The manager process then sends these load balance information numbers to all worker processes. Based on these load balance information numbers, $Worker_0$ generates S-polynomials i from $Critical\ pair_c$ so that $c \bmod 3 = 1$. $Worker_1$ generates S-polynomials from $Critical\ pair_c$ so that $c \bmod 3 = 0$, and $Worker_2$ generates S-polynomials i from $Critical\ pair_c$ so that $c \bmod 3 = 2$.

3.5 various implementations

action may take a long time if the polynomials are complicated, that is, they have many monomials and their monomials have large coefficients. To improve the efficiency, we should avoid creating complicated polynomials by reduction.

In our parallel implementation, there are four issues that may have a great influence on complexity of the polynomials created by reduction: admissible ordering (Buchberger, 1983), order of using polynomials of the intermediate base for reduction, interreduction, and using redundant elements of the intermediate base to reduce polynomials.

As for the admissible ordering and the order of using polynomials, we have implementations that are regarded as the best for computing Gröbner bases efficiently: the total degree reverse lexicographic ordering and the order that *older* polynomials are used earlier than newer ones. The total degree reverse lexicographic ordering is said to give the best theoretical and practical complexity to the computation of Gröbner bases as described in Hollman (1992).

The order of using older polynomials earlier is recommended by both Giovini *et al.* (1991) and Hollman & Langemyr (1991). Hollman (1992) also pointed out that the result of reduction of a polynomial tends to be more complicated than the polynomial before reduction. That is, there is a tendency for a newer polynomial to be more complicated than an older one, because a newer polynomial is reduced many more times than an older one. Therefore, to avoid creating complicated polynomials by reduction, older polynomials should be used before newer ones.

These implementations are also employed to our algebraic constraint solvers. As for interreduction and using redundant elements, however, it is necessary to examine their effect on the efficiency of the parallel computation of Gröbner bases.

3.5.1 interreduction

Interreduction is the reduction of every polynomial in an intermediate base by other polynomials. Since newer polynomials have already been reduced by older polynomials, interreduction means that older polynomials are reduced by newer ones. Thus, since newer polynomials tend to be more complicated than older ones, interreduction may lead to complicated polynomials.

Giovini *et al.* (1991) described no interreduction as better. Though Czapor (1991) said that interreduction is necessary to efficiently compute Gröbner bases in lexicographic ordering, we use the total degree reverse lexicographic ordering instead of lexicographic ordering. In our parallel implementation, we compare a no interreduction strategy to a one-step interreduction strategy that replaces every polynomial r_i in an intermediate base with r'_i , where $r_i \Rightarrow_h r'_i$ and h is a polynomial obtained as a new element of the intermediate base.

3.5.2 using redundant polynomials for reduction

If a leading monomial of a polynomial in the intermediate base can be reduced by a newer polynomial, then the reducible polynomial is *redundant*, as described in Gebauer & Möller (1988). Is it better to use redundant polynomials for reduction or not? Since there is a tendency for older polynomials to be less complicated than newer ones, a redundant polynomial is expected to be less complicated than a newer polynomial which can reduce its leading monomial. Therefore, the result of reduction of a polynomial by the redundant polynomial is expected to be less complicated than that by the newer one. In this sense,

a redundant polynomial may be effective for reduction. Giovini *et al.* (1991) described that redundant polynomials should be used. On the other hand, Gebauer & Möller (1988) said that the intermediate base should be kept as small as possible, that is, redundant polynomials should be cancelled.

In order to check whether using redundant polynomials is effective for reduction, we conducted several experiments.

3.6 experiment

In this section, we describe the results of our experiments. There are two issues to be examined. One is interreduction, and the other is whether to use redundant polynomials for reduction. Thus, we made four parallel algebraic constraint solvers as bellow.

NI-UR: No Interreduction and Using Redundant polynomials.

NI-NR: No Interreduction and Not using Redundant polynomials.

OI-UR: One step Interreduction and Using Redundant polynomials.

OI-NR: One step Interreduction and Not using Redundant polynomials.

In each parallel algebraic constraint solver, the admissible ordering is the total degree reverse lexicographic ordering, and older polynomials are used earlier than newer ones for reduction.

The parallel algebraic constraint solvers are implemented on the parallel inference machine PIM/m developed at ICOT, and in the kernel language KL1 for the parallel inference machine. To measure the speed of the machine, we use a unit called *LIPS* (logical inference per second) described in Matsumoto (1990). Nakashima (1992) reports that each PIM/m processor has a clock of 15.4 MHz and a speed of 615 KLIPS at its highest peak performance, while the performance of SICStus Prolog systems on a SUN Sparc Station 10/30 is 1053 KLIPS. PIM/m is a distributed memory machine, and Matsumoto (1990) reports that it takes 30 logical inference steps in PIM/m to read one primitive data, such as a variable in a monomial, from the other processor.

To examine the performance of our parallel algebraic constraint solvers, we use the twelve benchmarks listed below.

1. Katsura-4 (Boege *et al.*, 1986): (5 variables and 5 polynomials)
2. Katsura-5 (Boege *et al.*, 1986): (6 variables and 6 polynomials)
3. Katsura-6 (Katsura, 1986): (7 variables and 7 polynomials)
4. Butcher (Boege *et al.*, 1986): (8 variables and 8 polynomials)
5. Modified Hairer-2 (Sawada *et al.*, 1994): (13 variables and 11 polynomials)

6. Modified Hairer-3 (Sawada *et al.*, 1994): (13 variables and 13 polynomials)
7. Modified Gerdt (Sawada *et al.*, 1994): (8 variables and 13 polynomials)
8. Cyclic 4-roots: (4 variables and 4 polynomials)
9. Cyclic 5-roots: (5 variables and 5 polynomials)
10. Cyclic 6-roots: (6 variables and 6 polynomials)
11. T-6 (Backelin & Fröberg, 1991): (7 variables and 6 polynomials)
12. Ex-17 (Gebauer, 1985): (12 variables and 12 polynomials)

3.6.1 results of experiments

Figures 5 to 7 show the results of our experiments. We found that the computation time is not decreased by using two processors instead of one. Although this may seem strange, it is in fact quite reasonable. In our parallel implementations, each worker process is allocated to a different processor, and the manager process is also allocated to a different processor. With a single processor, however, the manager process and the worker process are, obviously, allocated to the same processor. The relationship between the number of worker processes n and the number of processors pn is given by

$$n = \begin{cases} 1 & (pn = 1) \\ pn - 1 & (pn \geq 2). \end{cases}$$

Thus, there is only a single worker process in both cases of using one and two processors, that is, the computation power is not increased by using two processors instead of one.

The effect of *nondeterminacy* in our experiments should be mentioned. As described in Section 3.3, the subtask causes nondeterminacy in our parallel algorithm. The effect of nondeterminacy is significant in the problems Modified Hairer-2 (Figure 6 (a)) and Modified Hairer-3 (Figure 6 (b)). In Modified Hairer-2, a remarkable speedup by using two processors is obtained by NI-UR as bellow.

1 processor: Computation time = 413 (seconds)

2 processors: Computation time = 191 (seconds)

As described above, however, the computing power does not increase by using two processors; the obtained speedup is due to the nondeterminacy caused by the subtask. With a single processor, since the manager process and the worker processes are allocated to the same processor, the subtask is not executed because the worker process cannot work when the manager process works. With two processors, however, since the worker process can work when the manager

process works, the subtask is executed. Thus, the sequence of generated polynomials of the intermediate base R is different from that with a single processor, and the total computational cost is decreased. The details of the computation are shown below.

- (1) 1 processor:
 - Number of generated polynomials of $R = 100$
 - Number of redundant polynomials of $R = 62$
 - Number of generated S-polynomials = 284
- (2) 2 processors:
 - Number of generated polynomials of $R = 82$
 - Number of redundant polynomials of $R = 44$
 - Number of generated S-polynomials = 257

In this case, the nondeterminacy has a good effect on the efficiency of the computation.

In Modified Hairer-3, however, the nondeterminacy has an adverse effect as shown in NI-NR and OI-UR. In these cases, the total computational cost increases drastically due to the nondeterminacy. Especially in OI-UR, the Gröbner base could not be computed in 24 hours with two processors, yet it was computed in 40 minutes with a single processor. The details of the computation by NI-NR are shown below.

- (1) 1 processor:
 - Number of generated polynomials of $R = 109$
 - Number of redundant polynomials of $R = 69$
 - Number of generated S-polynomials = 399
- (2) 2 processors:
 - Number of generated polynomials of $R = 109$
 - Number of redundant polynomials of $R = 69$
 - Number of generated S-polynomials = 406

The number of polynomials and S-polynomials generated during the computation are almost the same in both cases. Thus, the higher computational cost is due to the complexity of the polynomials.

3.6.2 influence of nondeterminacy caused by the subtask

Based on the results of experiments, NI-UR is employed as our parallel algebraic constraint solver. As described in the previous section, the subtask causes nondeterminacy and may adversely increase the total computational cost. However, the subtask is necessary to improve the efficiency of computing Gröbner bases by decreasing the idling time of processors. Therefore, a parallel algebraic constraint solver in which the nondeterminacy has little effect on the efficiency

should be employed. NI-UR, which is a parallel algebraic constraint solver without interreduction and using redundant polynomials, showed a good absolute speed of computation and relatively stable speedup with multi-processors compared to others, though it does not compute all the Gröbner bases faster than others.

The reason why the nondeterminacy has little effect on the total computational cost in NI-UR is as follows. When we have two polynomials of the same leading monomial, we compare the rest of the polynomials and choose the smaller polynomial as a new element of the intermediate base. Thus, the subtask causes the nondeterminacy since the subtask is not always executed, as described in Section 3.3. Furthermore, when polynomials in the intermediate base are reduced or cancelled, the irreducible form of a polynomial is not unique but dependent on the reduction timing. Thus, in NI-UR, the influence of nondeterminacy is smaller than in any other solvers.

4 Conclusion

We have developed several parallel algebraic constraint solvers in order to improve the absolute computation speed, because the efficiency of parallel algebraic constraint solvers are determined by the absolute computational speed and not by the speedup due to multi-processors. For example, for the Modified Gerdt problem in Figure 6 (c), though the absolute computation speed of NI-UR is superior to that of NI-NR, the speedup of NI-NR is better than that of NI-UR because the computation speed of NI-NR is slow with a single processor. From this example, we find that the speedup must be evaluated along with the absolute computation speed.

In our research and development of parallel algebraic constraint solvers on the distributed memory machine PIM/m, there are three major issues: the communication costs between processors, the nondeterminacy caused by the subtask, and the selection of a minimum polynomial as a new element of the intermediate base.

To decrease the communication costs, we have implemented our algebraic constraint solvers so that each processor has the same intermediate base on its memory and so that each worker process generates S-polynomials for itself by exchanging an integer with the manager process. This implementation requires sufficient memory for each processor because the total memory necessary for the computation increases almost linearly with the increase of worker processes. However, we used this implementation since it is essential to improve the efficiency by decreasing the communication costs.

The nondeterminacy may have either a good or bad effect on the efficiency, and it is very difficult to predict whether nondeterminacy improves the efficiency. If the subtask is not executed, then nondeterminacy does not occur. The subtask is, however, necessary to improve the efficiency by decreasing the idling time of

each processor. Thus, we have chosen NI-UR as our parallel algebraic constraint solver because the influence of nondeterminacy in NI-UR is less than that in any other.

Selecting a minimum polynomial as a new element of the intermediate base is known to decrease the total computational cost in the sequential Buchberger algorithm. To parallelize the Buchberger algorithm, it is also necessary to choose a minimum polynomial. We apply this selection strategy in our manager-worker model by synchronizing parallel worker processes. Such synchronization seems to decrease the speedup with multi-processors. By applying the selection strategy, however, it is guaranteed that the global minimum polynomial is selected as a new element, and the total computational cost does not increase with multi-processors. Furthermore, since each worker process reduces S-polynomials using the same algorithm, the total computation power increases even though synchronization may make some processors idle on every selection of a new element. As a result, the effect of the selection strategy on improving efficiency is greater than the decrease in efficiency caused by the defect in synchronization, and the speedup with multi-processors is also improved when compared with algebraic constraint solvers without the selection strategy (Hawley, 1991).

Finally, we must note the influence of the bignum operation. The calculation of coefficients, that is, the bignum operation, occupies a large part of the total computation. Therefore, the performance of the bignum operation affects the strategy used to compute the Gröbner bases. If the performance of the bignum operation is improved in the future, then the strategy might be changed for the same benchmarks employed in our experiment.

5 Acknowledgments

We would like to thank Prof. Kazuhiro Fuchi and Dr. Ryuzo Hasegawa for their encouragements and support of this research. We would also like to thank Mr. David J. Hawley, an important member of this research in the early stage, and Dr. Joachim Hollman for productive discussions and suggestions.


References

- [1] Aiba, A., Sakai, K., Sato, Y., Hawley, D. and Hasegawa, R. (1988). Constraint Logic Programming Language CAL. *Proc. International Conference on Fifth Generation Computer Systems*, 263-276.
- [2] Backelin, J. and Fröberg, R. (1991). How we proved that there are exactly 924 cyclic 7-roots. In: (Watt, S.M. editor) *Proc. ISSAC'91*, ACM, July, 103-111.


- [3] Boege, W., Gebauer, R. and Kredel, H. (1986). Some Examples for Solving Systems of Algebraic Equations by Calculating Groebner Bases *J. Symbolic Computation*, 83-98.
- [4] Buchberger, B. (1965). Ein Algorithmus zum Auffinden der Basiselemente des Restklassenringes nach einem nulldimensionalen Polynomideal. PhD Thesis, University of Innsbruck.
- [5] Buchberger, B. (1979). A Criterion for Detecting Unnecessary Reductions in the Construction of Gröbner Bases. *Proc. EUROSAM 1979*, Lecture Notes in Computer Science 72, Springer Verlag, Berlin-Heidelberg-New York, 3-21.
- [6] Buchberger, B. (1983). Gröbner bases: An Algorithmic Method in Polynomial Ideal Theory. *Technical report*, RISC-LINZ.
- [7] Buchberger, B. (1987). The Parallelization of Critical-Pair/Completion Procedures on the L-Machine. *Proc. The Japanese Symposium on Functional Programming*, February, 54-61.
- [8] Chikayama, T. (1992). Operating System PIMOS and Kernel Language KL1. *Proc. International Conference on Fifth Generation Computer Systems*, 73-88.
- [9] Clarke, E.M., Long, D.E., Michaylov, S., Schwab, S.A., Vidal, J.P. and Kimura, S. (1990). Parallel Symbolic Computation Algorithms. *Technical Report CMU-CS-90-182*, School of Computer Science, Carnegie Mellon University, October.
- [10] Colmerauer, A. (1987). Opening the Prolog III Universe: A new generation of Prolog promises some powerful capabilities. *BYTE*, August, 177-182.
- [11] Czapor, S.R. (1991). A Heuristic Selection Strategy for Lexicographic Gröbner Bases? *Proc. ISSAC'91*, ACM, July, 39-48.
- [12] Gebauer, R. (1985). A collection of examples for Groebner calculations. IBM Thomas J. Watson Research Center, Yorktown Heights, NY 10598.
- [13] Gebauer, R. and Möller, H.M. (1988). On an Installation of Buchberger's Algorithm. *J. Symbolic Computation*, 6(2 and 3), 275-286.
- [14] Giovini, A., Mora, T., Niesei, G., Robbiano, L. and Traverso, C. (1991). "One sugar cube, please" or selection strategies in the Buchberger algorithm. In: (Watt, S.M., editor) *Proc. ISSAC'91*, ACM, July, 49-54.
- [15] Hawley, D.J. (1991). The Concurrent Constraint Language GDCC and Its Parallel Constraint Solver. *Technical Report TR-713*, Institute for New Generation Computer Technology.

- [16] Hollman, J. and Langemyr, L. (1991). Algorithms for Non-linear Algebraic Constraints. *Proc. WCLP 91*, Marseille, January.
- [17] Hollman, J. (1992). Theory and Applications of Gröbner bases. Ph.D. Thesis, Department of Numerical Analysis and Computing Science, Royal Institute of Technology, Stockholm, Sweden.
- [18] Jaffar, J. and Lassez, J.-L. (1987). Constraint Logic Programming. In *4th IEEE Symposium on Logic Programming*.
- [19] Katsura, S. (1986). Theory of spin glass by the method of the distribution function of an effective field. *Progress of Theoretical Physics, Supplement*, No. 87, 139-154.
- [20] Matsumoto, Y. (1990). Text for KLI class: Basic Data for Hackers. *Technical Memorandum TM-898*, Institute for New Generation Computer Technology.
- [21] Melenk, H. and Neun, W. (1988). Parallel Polynomial Operations in the Large Buchberger Algorithm. *Computer Algebra and Parallelism*, Workshop at the TIM3 Laboratory, University of Grenoble, France, Academic Press, London, June. 143-158.
- [22] Nakashima, H., Nakajima, K., Kondo, S., Takeda, Y., Inamura, Y. and Onishi, S. (1992). Architecture and Implementation of PIM/m. *Proc. International Conference on Fifth Generation Computer Systems*, 425-435.
- [23] Ponder, C. G. (1988). Evaluation of "Performance Enhancements" in Algebraic Manipulation Systems. Ph.D. Thesis, Computer Science Division, University of California, September.
- [24] Sawada, H., Terasaki, S. and Aiba, A. (1994). Parallel Computation of Gröbner Bases on Distributed Memory Machines. *Technical Report TR-898*, Institute for New Generation Computer Technology.
- [25] Senechaud, P. (1989). Implementation of a parallel algorithm to compute a Gröbner basis on boolean polynomials. *Computer Algebra and Parallelism*, Academic Press, 159-166.
- [26] Siegl, K. (1990). Gröbner basis computation in STRAND: A case study for concurrent symbolic computation in logic programming languages. DIPLOMA Thesis, RISC-LINZ, November.
- [27] Taki, K. (1992). Parallel Inference Machine PIM. *Proc. International Conference on Fifth Generation Computer Systems*, 50-72.


- [28] Terasaki,S., Hawley,D.J., Sawada,H., Satoh,K., Menju,S., Kawagishi,T., Iwayama,N. and Aiba,A.(1992). Parallel Constraint Logic Programming Language GDCC and its Parallel Constraint Solvers. *Proc. International Conference on Fifth Generation Computer Systems*, 330-346.
- [29] Ueda,K. and Chikayama,T.(1990). Design of the Kernel Language for the Parallel Inference Machine. *The Computer Journal*, Vol.33, No.6, 494-500.
- [30] Vidal,J.-P.(1990). The computation of Gröbner bases on a shared memory multiprocessor. *Proc. DISCO'90*.



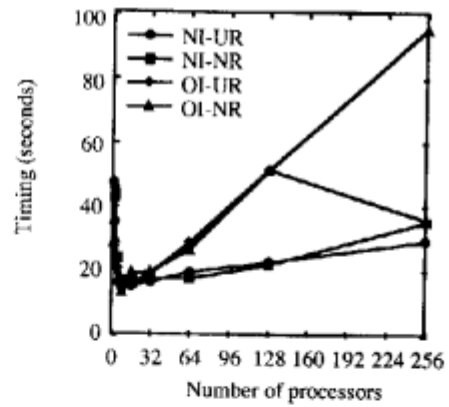
(a) Timing for Katsura-4



(b) Timing for Katsura-5




(c) Timing for Katsura-6

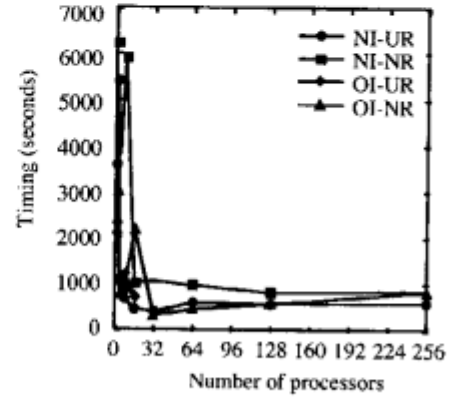


(d) Timing for Butcher


Figure 5: Timing data-1




(a) Timing for Modified Hairer-2



(b) Timing for Modified Hairer-3







(c) Timing for Modified Gerdt



(d) Timing for Cyclic 4-roots

Figure 6: Timing data-2



(c) Timing for T-6

(d) Timing for Ex-17

Figure 7: Timing data-3