

On-the-fly Optimization of Parallel Computation of Symbolic Symplectic Invariants

Joseph Ben Geloun, *Camille Coti*, Allen D. Malony

LIPN, CNRS UMR 7030, Université Sorbonne Paris Nord, France
University of Oregon, USA

Séminaire LATECE, UQAM
April 7th, 2022

Roadmap

Symplectic invariants

Symbolic computation

Parallel computation

Performance evaluation

- Comparison of the different schemes

- Combining the algorithms

Conclusion

Outline

Symplectic invariants

Symbolic computation

Parallel computation

Performance evaluation

- Comparison of the different schemes

- Combining the algorithms

Conclusion

Group invariants

Group invariants

- ▶ define quantum field theory interaction

To quantize gravity

- ▶ tensor models
- ▶ address classical Lie group (unitary and orthogonal) invariants in their construction.

Classical Lie groups :

- ▶ orthogonal
- ▶ unitary
- ▶ real **symplectic group** $Sp(2N)$

Goal : enumerate symplectic invariants

- ▶ Applications : e.g., problems in condensed matter, black hole physics

Matrix :

a_{14}	a_{24}	a_{34}	a_{44}
a_{13}	a_{23}	a_{33}	a_{43}
a_{12}	a_{22}	a_{32}	a_{42}
a_{11}	a_{21}	a_{31}	a_{41}

Tensor :

a_{003}	a_{103}	a_{203}	a_{303}	
a_{002}	a_{102}	a_{202}	a_{302}	
a_{001}	a_{101}	a_{201}	a_{301}	
a_{000}	a_{100}	a_{200}	a_{300}	a_{313}
				a_{312}
a_{010}	a_{110}	a_{210}	a_{310}	a_{311}
				a_{323}
a_{020}	a_{120}	a_{220}	a_{320}	a_{321}
				a_{333}
a_{030}	a_{130}	a_{230}	a_{330}	a_{331}

Calculation of symplectic invariants

Relies on **tensor contraction**

- ▶ the tensor contains formal, real variables
- ▶ result : a polynomial made of these variables
- ▶ **is the invariant equal to 0 ?**

Example : rank $d = 3$

- ▶ contraction of 4 tensors
- ▶ called *complete graph contraction*

Goal : show that for $N > 3$, the invariant is not identically null

Calculation of symplectic invariants (1/2)

Symplectic matrix J

- ▶ size $2N \times 2N$

$$J = \begin{pmatrix} 0 & I_N \\ -I_N & 0 \end{pmatrix}, \quad J^2 = -I_{2N}, \quad (1)$$

with :

- ▶ I_N , for all N the identity matrix of $M_N(\mathbb{R})$
- ▶ A matrix $K \in Sp(2N)$ obeys $KJK^T = J$, and $K^T JK = J$

Interactions of $Sp(2N)$ tensor models :

- ▶ contractions of an even number of tensors T
- ▶ contraction metric : the matrix J

Calculation of symplectic invariants (2/2)

Physical applications :

- ▶ in particular, contraction of 4 tensors as follows :

$$T^4 = \sum_{\substack{a_1, \dots, a_6 \\ \bar{a}_1, \dots, \bar{a}_6}} \left[\prod_{i=1}^6 J_{a_i \bar{a}_i} \right] T_{a_1, a_2, a_3} T_{a_4, a_5, \bar{a}_3} T_{\bar{a}_4, \bar{a}_2, a_6} T_{\bar{a}_1, \bar{a}_5, \bar{a}_6} \quad (2)$$

- ▶ for all c , a_c and $\bar{a}_c \in [1, 2N]$
- ▶ T^4 denotes the invariant

Algorithm 1 : “Naive”

```

1:  $Tens = 0$ 
2: for  $a1 \leftarrow 0, size$  do
3:   for  $a2 \leftarrow 0, size$  do
4:     for  $a3 \leftarrow 0, size$  do
5:        $A = T[a1][a2][a3]$ 
6:       for  $b1 \leftarrow 0, size$  do
7:          $TAB = J[a1][b1]$ 
8:         for  $b2 \leftarrow 0, size$  do
9:           for  $b3 \leftarrow 0, size$  do
10:             $TABB = TAB * A * T[b1][b2][b3]$ 
11:            for  $c1 \leftarrow 0, size$  do
12:              for  $c2 \leftarrow 0, size$  do
13:                 $TABC = TABB * J[a2][c2]$ 
14:                for  $c3 \leftarrow 0, size$  do
15:                   $TABCC = TABC * T[c1][c2][c3] * J[b3][c3]$ 
16:                  for  $d1 \leftarrow 0, size$  do
17:                     $TABCD = TABCC * J[c1][d1]$ 
18:                    for  $d2 \leftarrow 0, size$  do
19:                       $TABCDD = TABCD * J[b2][d2]$ 
20:                      for  $d3 \leftarrow 0, size$  do
21:                         $Tens = Tens + TABCDD * T[d1][d2][d3] * J[a3][d3]$ 
22:                      end for
23:                    end for
24:                  end for
25:                end for
26:              end for
27:            end for
28:          end for
29:        end for
30:      end for
31:    end for
32:  end for
33: end for

```


Some algebraic properties

The matrix J has some *symmetries* :

$$J = \begin{pmatrix} 0 & I_N \\ -I_N & 0 \end{pmatrix}, \quad J^2 = -I_{2N}, \quad (3)$$

Hence the tensor contraction becomes :

$$T^4 = \sum_{I \subset \{1,2,\dots,6\}} (-1)^{6-|I|} \prod_{l \in I} \left[\sum_{a_l, \bar{a}_l} \delta_{\bar{a}_l, a_l + N} \right] \\ \times \prod_{l \notin I} \left[\sum_{a_l, \bar{a}_l} \delta_{a_l, \bar{a}_l + N} \right] T_{a_1, a_2, a_3} T_{a_4, a_5, \bar{a}_3} T_{\bar{a}_4, \bar{a}_2, a_6} T_{\bar{a}_1, \bar{a}_4, \bar{a}_6}.$$

Algorithm 2 : exploiting the aforementioned properties

```

1:  $Tens = TE = T_1 = T_2 = T_3 = T_4 = T_5 =$ 
    $T_{12} = T_{13} = T_{14} = T_{16} = T_{23} = T_{24} = T_{26} =$ 
    $T_{123} = T_{126} = T_{134} = 0$ 
2:  $N = size/2$ 
3: for  $a_4 \leftarrow, N$  do
4:    $A_4 = a_4 + N$ 
5:   for  $a_2 \leftarrow 0, N$  do
6:      $A_2 = a_2 + N$ 
7:     for  $a_6 \leftarrow 0, N$  do
8:        $A_6 = a_6 + N$ 
9:        $W_1 = T[a_4][a_2][a_6]$ 
10:       $W_2 = T[a_4][a_2][a_6]$ 
11:       $W_3 = T[a_4][a_2][A_6]$ 
12:       $W_4 = T[a_4][a_2][a_6]$ 
13:       $W_5 = T[a_4][a_2][A_6]$ 
14:       $W_6 = T[a_4][a_2][A_6]$ 
15:       $W_7 = T[a_4][a_2][A_6]$ 
16:      for  $a_1 \leftarrow 0, N$  do
17:         $A_1 = a_1 + N$ 
18:        for  $a_5 \leftarrow 0, N$  do
19:           $A_5 = a_5 + N$ 
20:           $Z_1 = T[a_1][a_5][a_6]$ 
21:           $Z_2 = T[a_1][a_5][a_6]$ 
22:           $Z_6 = T[a_1][a_5][A_6]$ 
23:           $T_5 = W_3 * T[a_1][A_5][a_6]$ 
24:           $TE = W_4 * T[a_1][A_5][A_6]$ 
25:           $T_1 = W_3 * Z_2$ 
26:           $T_{13} = T_1$ 
27:           $T_2 = W_5 * Z_1$ 
28:           $T_{23} = T_2$ 
29:           $T_3 = W_3 * Z_1$ 
30:           $T_4 = W_6 * Z_1$ 
31:           $T_{12} = W_5 * Z_2$ 
32:           $T_{14} = W_6 * Z_2$ 
33:           $T_{134} = T_{14}$ 
34:           $T_{16} = W_1 * Z_6$ 
35:           $T_{24} = W_7 * Z_1$ 
36:           $T_{26} = W_2 * T[a_1][a_5][A_6]$ 
37:           $T_{123} = W_5 * Z_2$ 
38:           $T_{126} = W_2 * Z_6$ 
39:          for  $a_3 \leftarrow 0, N$  do
40:             $A_3 = a_3 + N$ 
41:             $TE+ = TE * T[a_1][a_2][a_3] * T[a_4][a_5][a_3]$ 
42:             $T_5+ = T_5 * T[a_1][a_2][a_3] * T[a_4][a_5][a_3]$ 
43:             $X_7Y_5 = T[a_1][a_2][a_3] * T[a_4][A_5][a_3]$ 
44:             $T_{1+} = T_1 * X_7Y_5$ 
45:             $T_{16+} = T_{16} * X_7Y_5$ 
46:             $T_{2+} = T_2 * T[a_1][a_2][a_3] * T[a_4][A_5][a_3]$ 
47:             $T_{3+} = T_3 * T[a_1][a_2][a_3] * T[a_4][A_5][a_3]$ 
48:             $T_{4+} = T_4 * T[a_1][a_2][a_3] * T[a_4][A_5][a_3]$ 
49:             $T_{12+} = T_{12} * T[a_1][a_2][a_3] * T[a_4][A_5][a_3]$ 
50:             $T_{13+} = T_{13} * T[a_1][a_2][a_3] * T[a_4][A_5][a_3]$ 
51:             $T_{14+} = T_{14} * T[a_1][a_2][a_3] * T[a_4][A_5][a_3]$ 
52:             $T_{23+} = T_{23} * T[a_1][a_2][a_3] * T[a_4][A_5][a_3]$ 
53:             $T_{24+} = T_{24} * T[a_1][a_2][a_3] * T[a_4][A_5][a_3]$ 
54:             $T_{26+} = T_{26} * T[a_1][a_2][a_3] * T[a_4][A_5][a_3]$ 
55:           $T_{123+} = T_{123} * T[a_1][a_2][a_3] * T[a_4][A_5][a_3]$ 
56:           $T_{126+} = T_{126} * T[a_1][a_2][a_3] * T[a_4][A_5][a_3]$ 
57:           $T_{134+} = T_{134} * T[a_1][a_2][a_3] * T[a_4][A_5][a_3]$ 
58:          end for
59:        end for
60:      end for
61:    end for
62:  end for
63: end for
64:  $Tens = 4 * (TE + T_{12} + T_{13} + T_{14} + T_{16} + T_{23} + T_{24} +$ 
    $T_{26} - (T_1 + T_2 + T_3 + T_4 + T_5 + T_{123} + T_{126} + T_{134}))$ 

```

"Only" 6 nested loops

Outline

Symplectic invariants

Symbolic computation

Parallel computation

Performance evaluation

Comparison of the different schemes

Combining the algorithms

Conclusion

Symbolic computation

Let's compute polynomial expressions by hand :

▶ $P_1 = x_1, P_2 = x_2$

▶ $P_1 + P_2 = x_1 + x_2$

Symbolic computation

Let's compute polynomial expressions by hand :

▶ $P_1 = x_1, P_2 = x_2$

▶ $P_1 + P_2 = x_1 + x_2$

▶ $P_1 = x_1 + x_2 + x_3 + x_4, P_2 = x_2$

▶ $P_1 + P_2 = x_1 + 2 * x_2 + x_3 + x_4$

Symbolic computation

Let's compute polynomial expressions by hand :

▶ $P_1 = x_1, P_2 = x_2$

▶ $P_1 + P_2 = x_1 + x_2$

▶ $P_1 = x_1 + x_2 + x_3 + x_4, P_2 = x_2$

▶ $P_1 + P_2 = x_1 + 2 * x_2 + x_3 + x_4$

▶ $P_1 = x_1 + x_2 + x_3 + x_4, P_2 = x_2 + x_4 + x_6 + x_8$

▶ $P_1 + P_2 = x_1 + 2 * x_2 + x_3 + 2 * x_4 + x_6 + x_8$

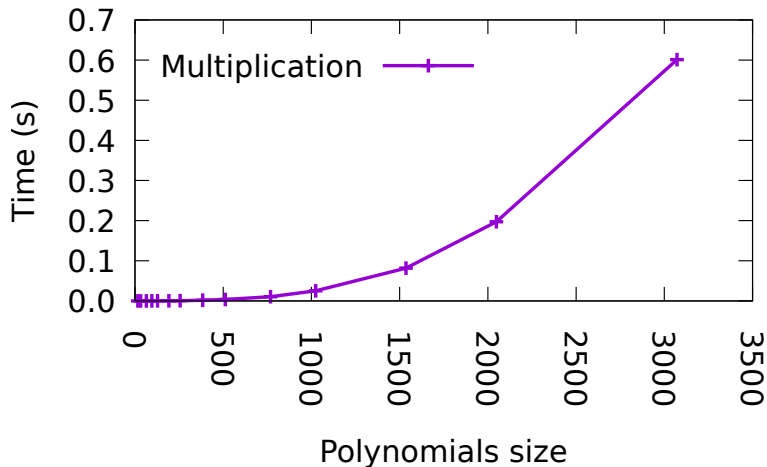
Symbolic computation

Let's compute polynomial expressions by hand :

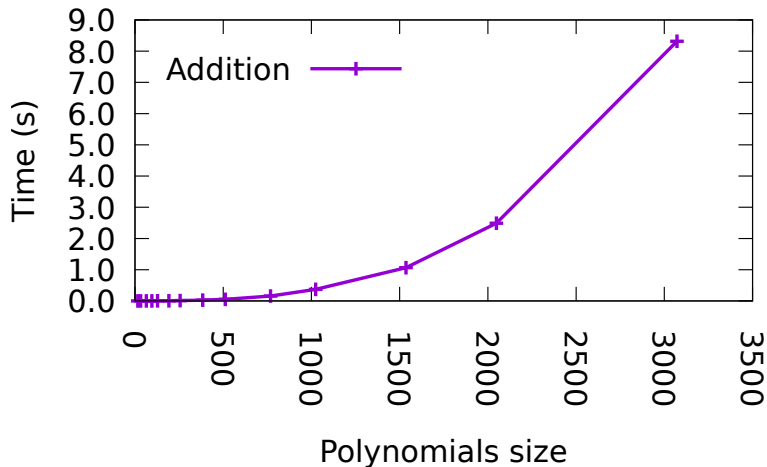
- ▶ $P_1 = x_1, P_2 = x_2$
 - ▶ $P_1 + P_2 = x_1 + x_2$
- ▶ $P_1 = x_1 + x_2 + x_3 + x_4, P_2 = x_2$
 - ▶ $P_1 + P_2 = x_1 + 2 * x_2 + x_3 + x_4$
- ▶ $P_1 = x_1 + x_2 + x_3 + x_4, P_2 = x_2 + x_4 + x_6 + x_8$
 - ▶ $P_1 + P_2 = x_1 + 2 * x_2 + x_3 + 2 * x_4 + x_6 + x_8$
- ▶ $P_1 = x_1 + x_2 + x_3 + x_4, P_2 = -x_1 - x_2 - x_3 - x_4$
 - ▶ $P_1 + P_2 = x_1 - x_1 + x_2 - x_2 + x_3 - x_3 + x_4 - x_4$
 - ▶ $P_1 + P_2 = 0 + 0 + 0 + 0$

→ **The computation time depends on the number of elements in the polynomial**

Computation time : multiplication

Operation type : $P_1 = 2 * P_2$ 

Computation time : addition

Operation type : $P_1 = P_2 + P_3$ 

Outline

Symplectic invariants

Symbolic computation

Parallel computation

Performance evaluation


Comparison of the different schemes

Combining the algorithms


Conclusion

Parallelizing the computation : domain decomposition

Naive approach : domain decomposition



Problem : not all the subparts of the matrix will take the same time



Parallelizing the computation : load balancing

Other problem : remember the computation time of our polynomials


- ▶ Process 1 : $x_1 + x_2 - x_1 - x_2$
- ▶ Process 2 : $x_1 + x_2 + x_1 + x_3$
- Process 1 will compute a lot faster than process 2

We need automatic, dynamic load balancing

- Use a master-worker scheme.

Master-worker ?

result = 0



Master

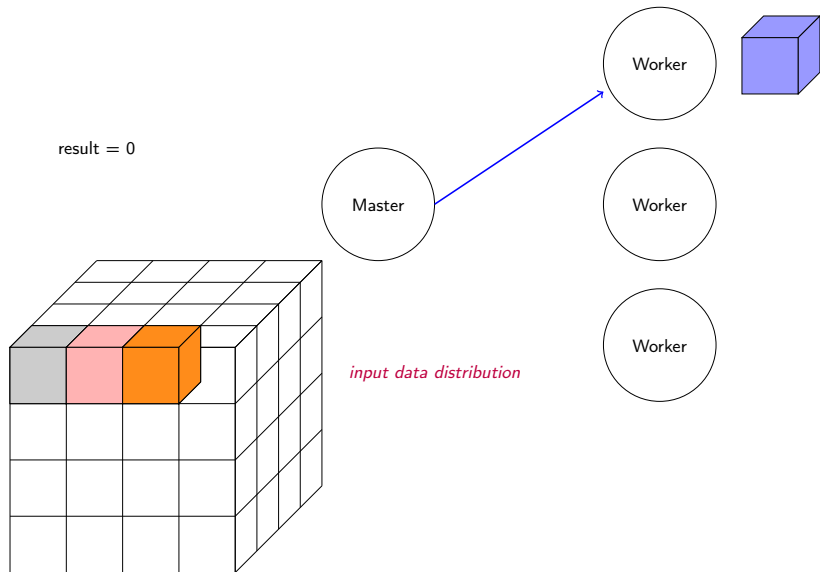
input data distribution

Worker

Worker

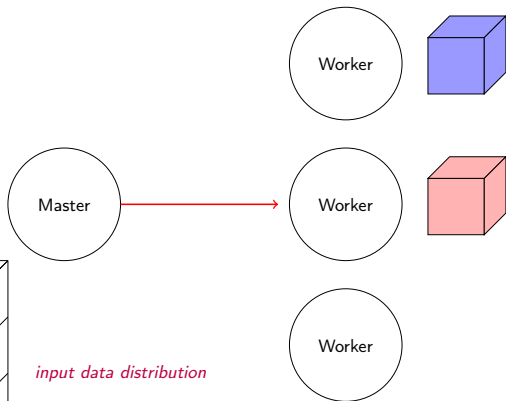

Worker

Master-worker ?





Master-worker ?

result = 0





Master-worker ?


result = 0




Master

*input data distribution*

Worker



Worker





Worker



Master-worker ?



formation of the global result

result = P0



Master-worker ?

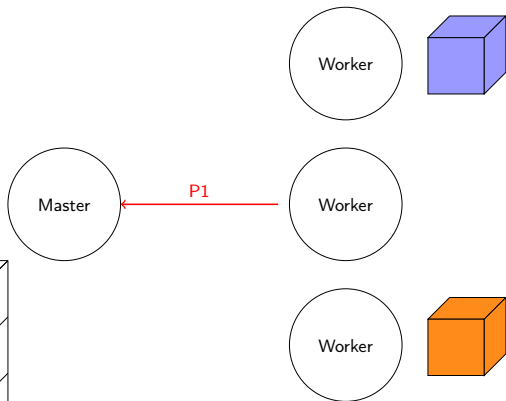

result = P0



Master-worker ?



formation of the global result

result = $P_0 + P_1$








Master-worker ?

result = P0+P1



Master-worker ?

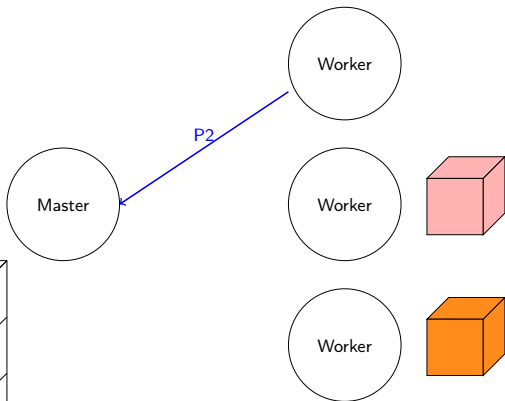

result = P0+P1



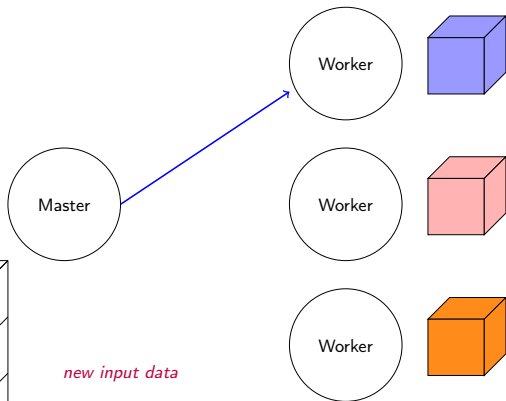

Master-worker ?

formation of the global result

$$\text{result} = P_0 + P_1 + P_2$$








Master-worker ?

 $\text{result} = P_0 + P_1 + P_2$ 

Master-worker ?

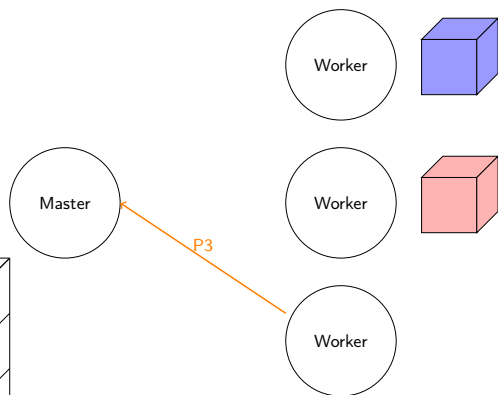

$$\text{result} = P_0 + P_1 + P_2$$





Master-worker ?

formation of the global result

result = $P_0 + P_1 + P_2 + P_3$








Master-worker ?

$$\text{result} = P_0 + P_1 + P_2 + P_3$$


Master-worker ?

result = $P_0 + P_1 + P_2 + P_3$



Master-worker ?

Traditional master-worker :

- ▶ The master maintains two queues : data and results
- ▶ The master sends chunks of data to the workers
- ▶ The workers compute partial sums
- ▶ The master gets results from the workers, **combines them to form the global result**

In our case :

- ▶ The workers send partial sums
- ▶ The master adds them to form the **global sum** (the invariant)

Problem : this global sum gets bigger and bigger

→ Bottleneck on the master, busy adding polynomials



Delegate the sum

Bottleneck on the master

- ▶ Ask a worker to compute this sum
- ▶ The workers can have two different types of tasks :
 - ▶ Compute a partial sum (inner loops)
 - ▶ Add partial sums to form the global sum

However :


- ▶ Adds interactions between the master and a worker
 - ▶ A worker computing the global sum is not computing any partial sum
- Switch to this scheme when the global sums are too expensive



Hierarchical master-worker

Another reason why the master can become a bottleneck :

- ▶ Workers' requests are too frequent
 - ▶ Granularity is too small, too many workers
- Use a hierarchical pattern
- ▶ The workers request work from a foreman
 - ▶ The foremen request work from the master
 - ▶ The master sends a bigger chunk of work to each foreman
 - ▶ The foremen split this chunk into smaller chunks
 - ▶ The foremen compute intermediate sums, or delegate to a worker



If the master becomes a bottleneck, how to tell why?

- ▶ **Use measurements**
- ▶ Which time proportion does the master spend outside communications?

Stateful worker

The main challenge is the computation of the global sum

- ▶ Do not centralize it on the master
- ▶ **Keep the partial sums on the workers**, add them while waiting for more data
- ▶ Only at the end, add them to form the global polynomial (tree)

Algorithm 1 Master

```

/* prepare parameter sets */
for a4 ← 0, N do
  for a2 ← 0, N do
    for a6 ← 0, N do
      params.push_back({ a4,a2,a6})
    end for
  end for
end for
/* distribute parameter sets */
while !parameters.empty() do
  src = recv( request, ANY_SOURCE )
  p = params.pop()
  send( src, p, TAG_WORK )
/* wait for all the workers */
while running() do
  src = recv( request, ANY_SOURCE )
  send( src, 0, TAG_END )
/* global sum */
Tens = reduction_sum()

```

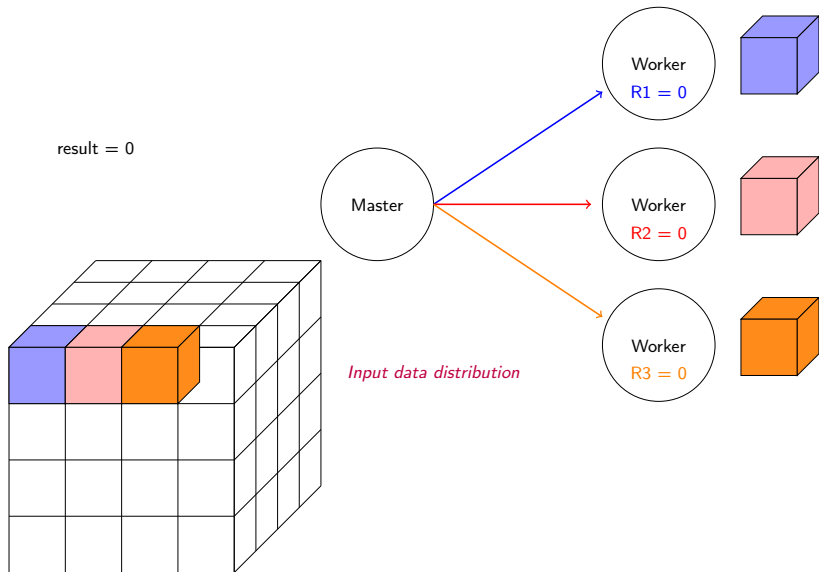
Algorithm 2 Worker

```

Tens = 0
T = 0
while true do
  /* ask for some work */
  send( root, 0, TAG_REQ )
  /* as I wait for a parameter set, add my polynomials */
  req = irecv( ROOT, ANY_TAG )
  if T != 0 then
    Tens += T
    p, tag = wait( req )
    if tag == TAG_END then
      break
    /* compute a polynomial for the parameters I have received */
    T = compute( p )
  /* global sum */
  reduction_sum(Tens)



```

Stateful worker




Stateful worker


result = 0




Master

Local computations

Worker

 $R1 = 0$ 

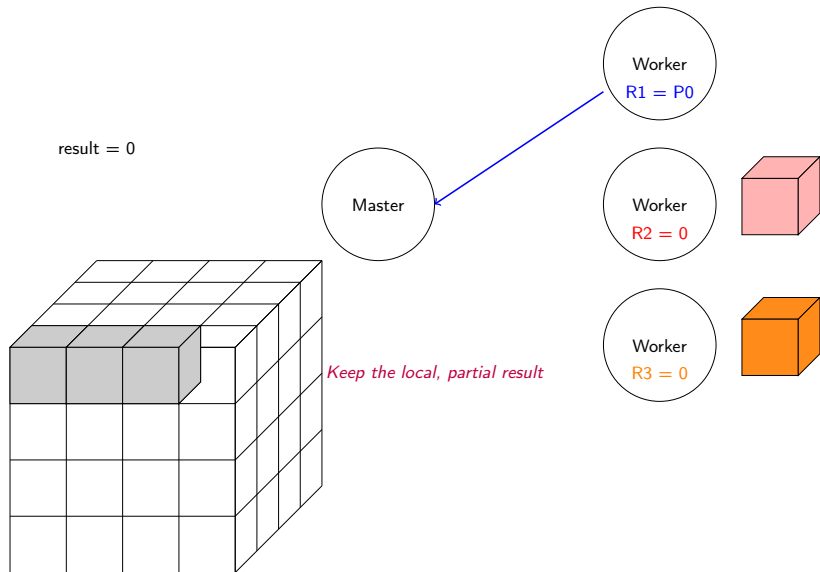
Worker

 $R2 = 0$ 

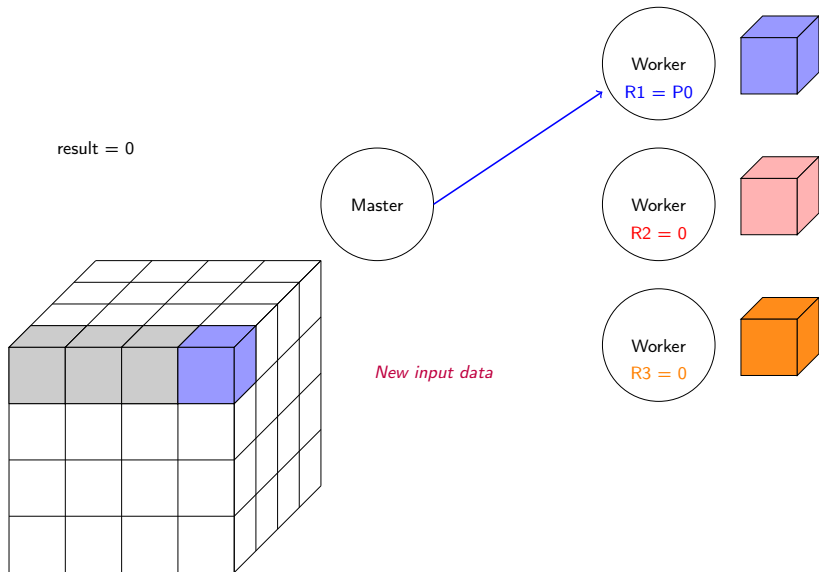
Worker

 $R3 = 0$ 

Stateful worker








Stateful worker

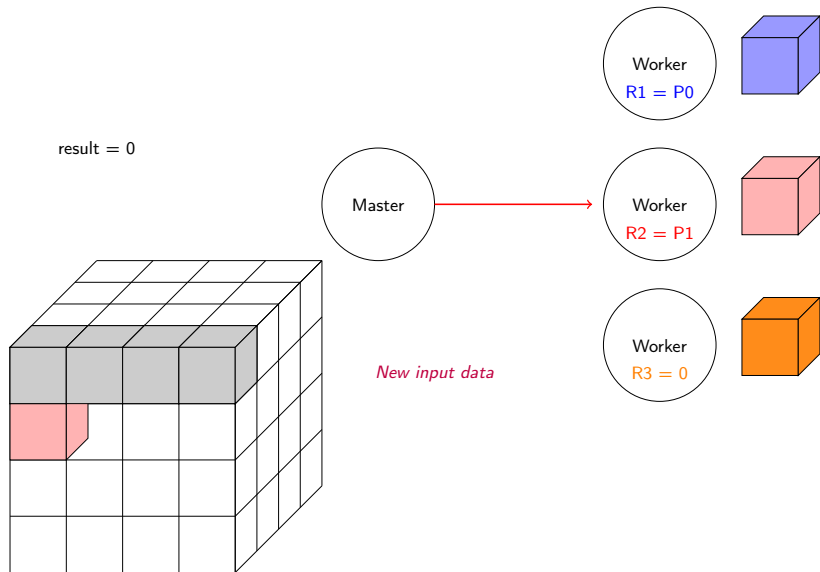


Stateful worker

result = 0

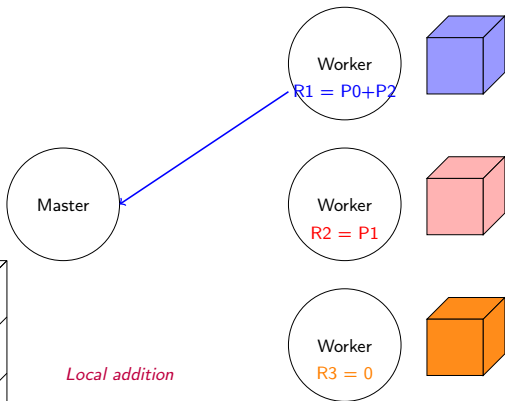

*Keep the local, partial result*

Stateful worker

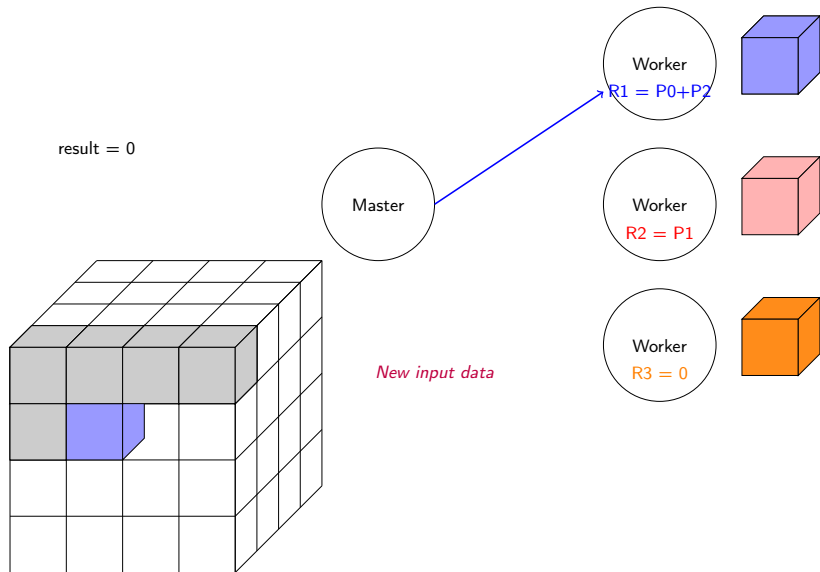


Stateful worker

result = 0








Stateful worker



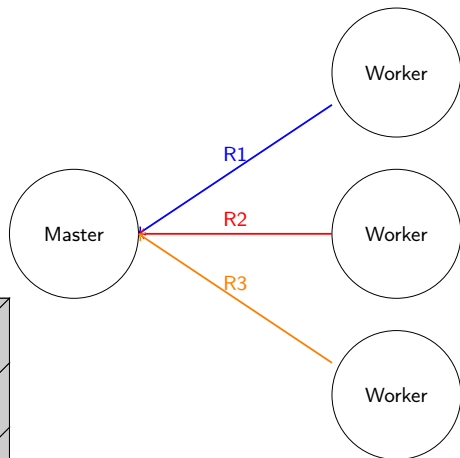

Stateful worker

result = 0



Stateful worker


result = 0



Stateful worker

Formation of the global result

$$\text{result} = R1+R2+R3$$



Master

Worker

Worker

Worker

Outline

Symplectic invariants

Symbolic computation

Parallel computation

Performance evaluation

Comparison of the different schemes

Combining the algorithms

Conclusion

Experimental setup

Software environment

- ▶ OpenMPI 4.1, Linux kernel 4.9.0, Debian 9.8, g++ 8.3.0

Hardware

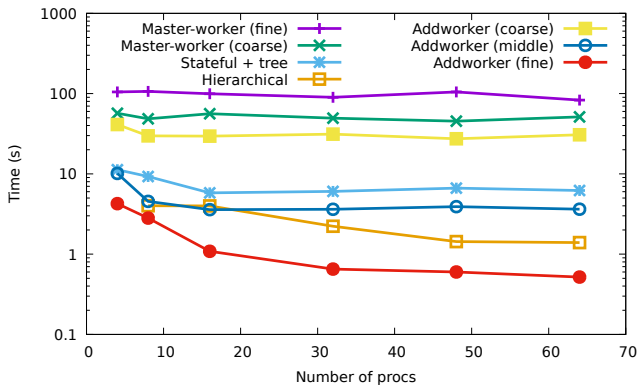
- ▶ Grid'5000 cluster : Parapide (Rennes)
- ▶ 20 nodes, 2x Intel Xeon X5570 CPUs (4 cores/CPU), 24 GB of memory
- ▶ 20 Gb InfiniBand + GigaEthernet

Symbolic computing libraries :

- ▶ GiNaC 1.7.6 (not Gignac!)
- ▶ Obake : successor of Piranha, better on multivariate polynomials

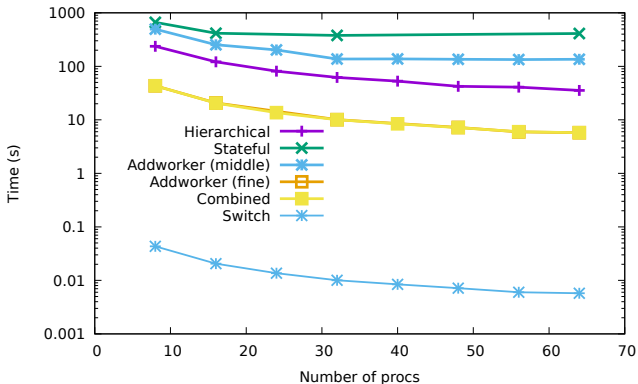
Small tensor ($N=4$, size=8)

Using Obake.



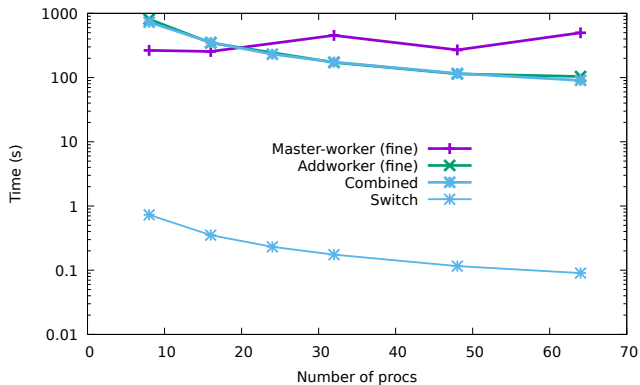
Medium tensor (N=6, size=12)

Using Obake. Blue bottom line : when the switch between *master-worker* and *addition on a worker* happens.



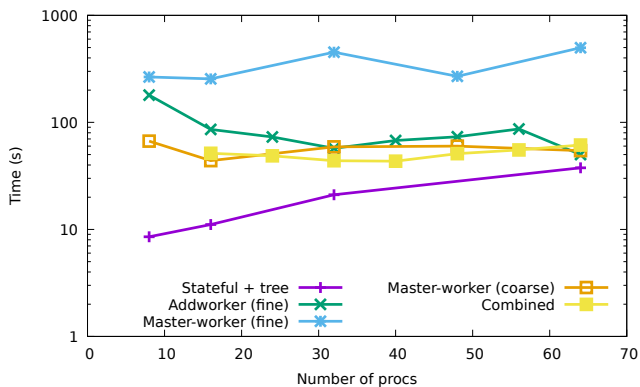
Large tensor ($N=8$, size=16) 1/2

Using Obake. At small scale : we switch too late.



Large tensor ($N=8$, size=16) 2/2

Using GiNaC : the polynomial operations do not take the same time.



Hierarchical ?

We have **never seen the algorithm switch to the hierarchical scheme**

- ▶ Policy : when the master is overloaded by requests → switch to the hierarchical scheme
- ▶ Maybe because when a lot of requests are received, a lot of additions are needed (intermediate and partial)

Outline

Symplectic invariants

Symbolic computation

Parallel computation

Performance evaluation

- Comparison of the different schemes

- Combining the algorithms

Conclusion

Conclusion 1/2

In this problem, the computational work **varies** during the computation

- ▶ We were not sure it would (annulling terms → reduced computation time)
- ▶ Increases in particular in the **critical path** (global sum)
- ▶ Non-linear

Goal : **get as much as we can away from the critical path**

Granularity of the computation :

- ▶ Increase the number of workers → refine the granularity to keep them busy
 - ▶ Too small grain → computation time too short wrt communications

Scalability :

- ▶ Increase the **size of the problem**
- ▶ Workers have more work
- ▶ **More (expensive) polynomial additions** (in the critical path)

Conclusion 2/2

Polynomial additions to for the global sum

- ▶ Become expensive quickly
 - ▶ **Switch to a pattern that computes them on a worker**
 - ▶ Good choice most of the times, switch quickly
- ▶ Stateful workers : much faster... **except to form the global polynomial**
 - ▶ most of the times its cost is higher than the gain during the computation.

Hierarchical scheme

- ▶ Never encountered a case where the switch policy applies
- ▶ The workload on each worker **increases faster than the congestion on the master** (as the size increases to scale)
- ▶ Larger problem → larger polynomials to add

Dynamic workload, evolving (roughly) monotonically : advantage of **run-time performance measurements** to make decisions.