

Exploiting efficiently different parallel architectures

1. Need for *parallel* computations
2. Importance of using *standard tools and templates*
3. Optimizing the computations *on one processor*
4. Parallel runs on *shared memory* computers
5. Parallel runs on *distributed memory* computers
6. Parallel runs on *more advanced* architectures
7. Numerical results
8. Conclusions

The competition between computers and scientists: the computers are always behind!

“Although the fastest computers can execute **millions** of operations in one second **they are always too slow**. This may seem a paradox, but the heart of the matter is: the bigger and better computers become, the larger are the problems scientists and engineers want to solve”.

Arthur Jaffe:

“Ordering the universe: The role of mathematics”,
SIAM Review, Vol. 26 (1984), p. 478.

References

- 1. I. Dimov, K. Georgiev, Tz. Ostromsky, R. van der Pas and Z. Zlatev:** “*Computational Challenges in Numerical Treatment of Large Air Pollution Models*”, The Mathematical Preprint Server, <http://www.mathpreprints.com>, 2001.
- 2. K. Georgiev and Z. Zlatev:** “*Parallel Sparse Matrix Algorithms for Air Pollution Models*”. *Parallel and Distributed Computing Practices*, **Vol. 2 (1999)**, 429-443.
- 3. W. Owczarz and Z. Zlatev:** “*Parallel Matrix Computations in Air Pollution Modelling*”, *Parallel Computing*, **Vol. 28(2002)**, 355-368.

Two more papers in the Mathematical Preprint Server

Existing versions of DEM

<u>Species</u>	<u>32x32x10</u>	<u>96x96x10</u>	<u>288x288x10</u>	<u>480x480x10</u>
1	10240	92160	829440	2304000
2	20480	184320	1658880	4608000
10	102400	921600	8394400	23040000
35	358400	3225600	29030400	80640000
56	573440	5169960	-	-
168	1720320	15482880	-	-

Corresponding **2-D** versions exist

PC (or workstations) vs parallel computers

“Non-optimized” code

<u>Module</u>	<u>Comp. time</u>	<u>Percent</u>
Chemistry	16147	83.09
Advection	3013	15.51
Initialization	1	0.00
Input operations	50	0.26
Output operation	220	1.13
<u>Total</u>	<u>19432</u>	<u>100.00</u>

It is important to optimize the chemical part for this problem

2-D version on a 96x96 grid (50 km x 50 km)

The time-period is one month

The situation changes for the fine resolution models (the advection becomes very important)

The computing time is measured in seconds

One processor is used in this run

“Optimized” chemistry for the fine resolution 2-D code

<u>Module</u>	<u>Comp. time</u>	<u>Percent</u>
Chemistry	43.01	29
Advection	95.78	64
Initialization	0.59	0
Input operations	1.18	0.6
Output operations	9.82	6
Total	149.80	100

It is important to optimize the advection part for this problem

2-D version on a 480x480 grid (10 km x 10 km)

Time-period: **one year**

The computing time is measured in **hours**

Sixteen processors are used in this run

Ordering the computations

1. Cost of arithmetic operations **vs** cost of loading and storing the involved in the arithmetic computations quantities in old days and in the modern hierarchical memory architectures
2. It is important to work as long as possible with data which are in cache (preferably in the **fastest** cache when several levels of cache memory are available)
3. Difficult task (because we are not able to tell the computer “**put these data in cache and hold them there until we tell you to exchange them with other data**”)
4. While we have not a **direct** control of the contents of the cache memory, reordering the arithmetic operations can give us some possibility to do this **indirectly**.

Traditional ways of carrying out the chemical reactions

- Using a “box” subroutine

```
DO I=1,N                                ! Template 1
    Call the box routine to perform all chemical reactions at point I
END DO
```

- Vectorizing the computations

```
DO J=1,NSPECIES                          ! Template 2
    DO I=1,N
        Perform the chemical reactions involving species J at point I
    END DO
END DO
```

- The array containing the concentrations: $C(N,NSPECIES)$

Effect of implementing chunks

Size of chunks	Fujitsu	SGI Origin 2000	Power Mac G4	IBM SMP
1	76964	14847	6952	10313
48	2611	12114	5792	5225
9216	494	18549	12893	19432

- The choice **NSIZE=1** is a disaster on a vector computer, but not so bad on the parallel computers
- The maximal size of **NSIZE** is the best choice on a vector computer, but the worst choice on the parallel computers
- Medium sizes of **NSIZE** give good results on all parallel computers
- The optimal choice of **NSIZE** will depend on the computer used

Parallel runs on shared memory computers

Parallel tasks arise in a natural way after the splitting procedure

Number of the parallel tasks

- Advection sub-model: $NS*NZ$ very large tasks
- Chemistry sub-model: $NX*NY*NZ$ small tasks
- Vertical exchange: $NX*NY*NS$ small tasks

Achieving portability: by using only OpenMP directives

Computers used: **SGI Origin 2000** (up to 32 processors)
SUN (up to 16 processors)

Numerical results achieved on shared memory computers

<u>Processors</u>	<u>Comp. time</u>	<u>Speed-up</u>	<u>Efficiency</u>
1	42907	-	-
32	2215	19.37	61%

(96x96x10) version on **SGI Origin 2000** (NSIZE=48)

<u>Processors</u>	<u>Comp. time</u>	<u>Speed-up</u>	<u>Efficiency</u>
1	56615	-	-
16	9637	14.67	92%

(96x96x10) version on **SUN** (NSIZE=48)

Parallel runs on distributed memory computers

Portability: achieved by using **MPI** (Message Passing Interface); **PVM** can also be used

The space domain is divided into **p** sub-domains (**p** being the number of processors). **Each processors works on its own sub-domain**

Pre-processing

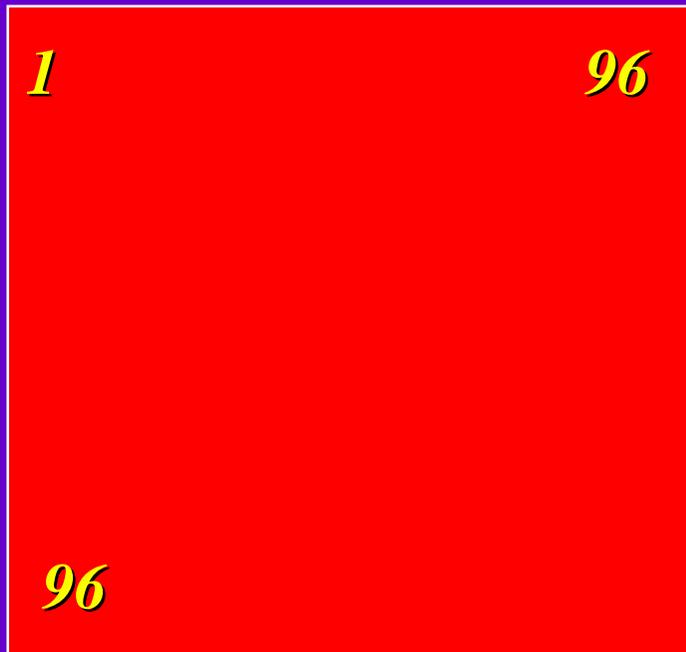
Post-processing

Reduction of the communications during the actual computations

Owczarz and Zlatev (2002)

Division into domains

Concentrations for an arbitrary
chemical compound



Dividing into sub-domains
Chemistry Advection



Distributed memory computers used in the runs

- IBM SP (up to 32 processors)
- Cray T3E (up to 64 processors)
- Macintosh Power PC Cluster (up to 8 processors)

Cray T3E in Edinburgh, Mac PC Cluster in Sofia, IBM SP in Copenhagen

Numerical results achieved on distributed memory computers

<u>Processors</u>	<u>Comp. time</u>	<u>Speed-up</u>	<u>Efficiency</u>
8	54978	-	-
32	15998	3.44	86%

(480x480) version on IBM SP (NSIZE=48)

<u>Processors</u>	<u>Comp. time</u>	<u>Speed-up</u>	<u>Efficiency</u>
32	18306	-	-
64	9637	1.90	95%

(480x480) version on T3E (NSIZE=48)

Numerical results achieved on distributed memory computers

<u>Processors</u>	<u>Comp. time</u>	<u>Speed-up</u>	<u>Efficiency</u>
1	5792	-	-
8	787	7.36	92%

(96x96) version on Macintosh Power PC Cluster
(NSIZE=48)

Parallel runs on more advanced parallel computers

- Computers which combine properties of shared memory computers and distributed memory computers
- Typical representative: **IBM SMP**
- Several nodes are available. Each node contains certain number of processors (4, 8 or 16)
- Shared memory mode (**OpenMP**) can be used within a node
- Distributed memory mode (**MPI**) has to be used across the nodes

Numerical results achieved on more advanced computers

<u>Processors</u>	<u>Comp. time</u>	<u>Speed-up</u>	<u>Efficiency</u>
1	5225	-	-
16	424	12.32	72%

(96x96) version on IBM SMP with 2 nodes
(8 processors per node)

NSIZE=48 was used in this run

Do we need OpenMP?

<u>Process</u>	<u>OpenMP version</u>	<u>MPI version</u>
Start	0.1	12.4
Wind + Sinks	5.8	2.2
Advection	101.2	30.1
Chemistry	232.6	161.9
Output	54.2	4.1
Communications	0.0	46.9
Pre-processing	0.0	11.1
Post-processing	0.0	12.0
<u>Total time</u>	<u>394.1</u>	<u>270.5</u>

Run on 16 processors of SGI Origin 2000

The 2-D version on 96x96 grid has been used

The reduction of the **advection time** for the MPI version was **expected**, while the reduction of the **chemical time** was **a big surprise**

Conclusions and open problems

- The use of **standard** tools is important
- The use of **templates** may facilitate the search for better numerical methods (it is possible to apply the same template with different numerical methods; the differences then are caused by the numerical methods only)
- It is important to **optimize** the computations **on one processor**
- **More powerful computers are needed in our field**