

Efficient Portable Parallel Implementation of a Large Scale Air Pollution Model for Sensitivity Studies

Tzvetan Ostromsky

Institute for Parallel Processing –
Bulgarian Academy of Sciences,
Acad. G. Bonchev st., bl. 25-A,
1113 Sofia, Bulgaria

E-mail: ceco@parallel.bas.bg

<http://parallel.bas.bg/~ceco/>



NATO Grant PDD(TC)-(ESP.EAP.CLG 982641)

"Monte Carlo Sensitivity Studies of Environmental Security"

Outline of the talk

- Introduction to the air pollution modeling
- The Danish Eulerian Model
- Numerical treatment of the model
- DEM and UNI-DEM
- Parallel implementation features
- Performance and scalability results
- Sensitivity analysis approach
- Total Sensitivity Indices
- Sobol' Approach
- Calculating sensitivity analysis data
- Concluding remarks

Introduction to the air pollution modeling

A system of PDE for calculating the concentrations of a number of chemical species (pollutants and other components of the air that interact with the pollutants) in a large 3D domain (part of the atmosphere above the studied geographical region).

The main physical and chemical processes (horizontal and vertical wind, diffusion, chemical reactions, emissions and deposition) should be adequately represented in the system.

Application areas:

- Ecosystem mutual relations
- Forestry and wild life protection
- Human health preservation
- Crops production losses and related economical aspects
- Global climate changes

Classification of the air pollution models

With respect to the coordinate system of the domain:

- Eulerian models (fixed with respect to the geographical coordinates);
- Lagrangian models (attached to the moving clouds).

With respect to the size of the domain:

- Large scale (best for global and long-term studies)
- Medium scale (mainly national level)
- Small (regional level, suitable for local and more specific studies)

The Danish Eulerian Model

$$\frac{\partial c_s}{\partial t} = -\frac{\partial(uc_s)}{\partial x} - \frac{\partial(vc_s)}{\partial y} - \frac{\partial(wc_s)}{\partial z} + \frac{\partial}{\partial x} \left(K_x \frac{\partial c_s}{\partial x} \right) + \frac{\partial}{\partial y} \left(K_y \frac{\partial c_s}{\partial y} \right) + \frac{\partial}{\partial z} \left(K_z \frac{\partial c_s}{\partial z} \right) + E_s - (k_{1s} + k_{2s})c_s + Q_s(c_1, c_2, \dots, c_q), \quad s = 1, 2, \dots, q .$$

- q – number of chemical species
(=number of equations in the system),
- c_s – concentrations of the chemical species,
- u, v, w – components of the wind along the coordinate axes,
- K_x, K_y, K_z – diffusion coefficients,
- E_s – emissions in the space domain,
- k_{1s}, k_{2s} – coefficients of dry and wet deposition,
- $Q_s(c_1, c_2, \dots, c_q)$ – non-linear functions that describe the chemical reactions between the species involved.

DEM is described in detail in the books of Zlatev (1995), Zlatev & Dimov (2006).

Numerical treatment of the model – characteristics and requirements

- Complexity of the equations – splitting, efficient algorithms;
- Huge computational domain – high speed and storage requirements;
- Decomposition – parallel tasks;
- Dynamics of the processes – small time-step (to ensure stability);

A challenging task for high-performance parallel computing!

Splitting into submodels

$$\begin{aligned}\frac{\partial c_s^{(1)}}{\partial t} &= -\frac{\partial(uc_s^{(1)})}{\partial x} - \frac{\partial(vc_s^{(1)})}{\partial y} + \frac{\partial}{\partial x} \left(K_x \frac{\partial c_s^{(1)}}{\partial x} \right) + \frac{\partial}{\partial y} \left(K_y \frac{\partial c_s^{(1)}}{\partial y} \right) \\ &= A_1 c_s^{(1)}(t) \quad \text{horizontal advection \& diffusion}\end{aligned}$$

$$\begin{aligned}\frac{\partial c_s^{(2)}}{\partial t} &= E_s + Q_s(c_1^{(2)}, c_2^{(2)}, \dots, c_q^{(2)}) - (k_{1s} + k_{2s})c_s^{(4)} = A_2 c_s^{(2)}(t) \\ &\quad \text{chemistry, emissions \& deposition}\end{aligned}$$

$$\begin{aligned}\frac{\partial c_s^{(3)}}{\partial t} &= -\frac{\partial(wc_s^{(3)})}{\partial z} + \frac{\partial}{\partial z} \left(K_z \frac{\partial c_s^{(3)}}{\partial z} \right) = A_3 c_s^{(3)}(t) \\ &\quad \text{vertical transport}\end{aligned}$$

Related work: Strang (1968); Marchuk (1982);
McRae, Goodin & Seinfeld (1984);
Lancer & Verwer (1998); Dimov, Farago & Zlatev (1999);
Dimov, Farago, Havasi & Zlatev (2001).

The chemical scheme: condensed CBM IV

Chemistry is of primary importance for the model. It is nonlinear and stiff – e.g. **tuff** and **complex** computational task. Without it the model brakes out in a number of independent advection-diffusion subproblems.

Condensed CBM IV (Carbon Bound Mechanism) – features:

- 35 pollutants and 116 chemical reactions;
- – 47 time-independent,
- – 69 time-dependent, including
- – 19 photo-chemical.

Gery et al. (1989)

UNI-DEM, the Unified Danish Eulerian Model

The development and improvements of DEM throughout the years has lead to a variety of different versions with respect to the grid-size/resolution, 2D – 3D layering and the number of species in the chemical scheme. The most prospective of them has been united under a common driver routine, called UNI-DEM. It provides an uniform and easy user access to the available up-to-date versions of the model.

The versions, incorporated in UNI-DEM, and their governing parameters are shown in the table below.

Choosable parameters for selecting an optional UNI-DEM version				
Parameter	Description	Optional values		
NX = NY	Grid size Grid step	96 × 96 50 km.	288 × 288 16.7 km.	480 × 480 10 km.
NZ	# layers (2D/3D)	1		10
NEQUAT	# chem. species	35	56	168

Numerical methods used in the solution of the submodels

- **Advection-diffusion part:** Finite elements, followed by predictor-corrector schemes with several different correctors (Zlatev - 1984)

Hints: Requires overlapping of the spatial subdomains for exchanging boundary conditions. The parallel implementation should take care for synchronization (in the case of shared memory model) or for the necessary communications (in the case of distributed memory model) on each time step.

- **Chemistry-deposition part:** An improved version of the QSSA (Quasi Steady-State Approximation) (Hesstevedt et al. (1978))

Hints: The calculations in each grid-point are small independent tasks, which are grouped in **chunks** for more efficient cache utilization.

- **Vertical transport:** Finite elements, followed by θ -methods

Hints: The calculations within each vertical grid-line is a relatively small task. These can be grouped in **chunks** for more efficient cache utilization, like in the chemical stage.

Distributed memory parallelization scheme

- Based on MPI. For maximum portability only **standard MPI routines** are used.
- Based on domain decomposition of the horizontal grid
 - restrictions on the number of MPI tasks.
- Domain overlapping of the advection-diffusion subproblems:
 - computational overhead, grows up with increasing the number of MPI tasks;
 - requires communication on each time step (**communication stage**).
- Improving the data locality for more efficient **cache utilization**
 - by using **chunks** to group properly the small tasks in the chemistry-deposition and vertical exchange stages.
- Additional **pre-processing** and **post-processing** stages are needed for scattering the input data and gathering the results. These are cheap, but their relative weight grows up with increasing the number of MPI tasks (*affects the total speed-up and efficiency*).

Numerical experiments on various parallel supercomputers

- Characteristics of the parallel machines used in the experiments:
 - A SUN cluster **SunFire E25k** at DTU, Lyngby, Denmark (72 dual-core CPU-s UltraSparc-IV / 1350 MHz),
 - A SUN cluster **SunFire E15k** at EPCC, the University of Edinburgh (48 “back-end” + 4 “front-end” Ultrasparc III processors, 900 MHz, 2 level cache), currently out of service;
 - An IBM high-performance cluster “HPCx”, based on IBM Power5 processors (64-bit RISC CPU / 1.5 GHz). The cluster is built out of 160 IBM P5 eServer 575 LPAR nodes (each with 16 CPU, 32 GByte RAM, 3-level cache).
- Mainly results of experiments with the 3D version of the model on (96×96) grid are presented in the tables. The chemical scheme with $q = 35$ species has been selected.
- **CHUNKSIZE=32** is used in the experiments. This parameter needs tuning in dependence with the size of the cache memory of the target machine (for its optimal utilization).

Performance and scalability of UNI-DEM (2D) on a SunFire E25k

The 2D UNI-DEM on a SunFire E25k						
PEs	TOTAL		Advection		Chemistry	
	time [sec]		(speed-up)			
1	1798		307		1374	
2	902	(1.99)	155	(1.98)	702	(1.96)
4	454	(3.96)	78	(3.94)	346	(3.97)
8	247	(7.28)	41	(7.49)	178	(7.72)
12	181	(9.93)	31	(9.90)	120	(11.45)
16	152	(11.82)	24	(12.79)	91	(15.10)
24	107	(16.80)	17	(18.06)	60	(22.90)
32	87	(20.67)	14	(21.93)	46	(29.87)

Results of parallel execution of the 2D UNI-DEM for one month on a SunFire E25k supercomputer at DTU. The job user time and the times of the main computational stages in seconds (followed by the corresponding speed-up) are shown in the corresponding columns.

Performance and scalability of UNI-DEM (3D) on a SunFire E25k

The 3D UNI-DEM on a SunFire E25k									
PEs	TOTAL		Advection		Chemistry		Ver. transport		
	time [sec]				(speed-up)				
1	17714		2786		13490		978		
2	9062	(1.95)	1402	(1.99)	6918	(1.95)	494	(1.98)	
4	4625	(3.83)	705	(3.95)	3407	(3.96)	246	(3.97)	
8	2499	(7.09)	374	(7.44)	1754	(7.69)	125	(7.84)	
12	1871	(9.47)	283	(9.84)	1181	(11.42)	84	(11.60)	
16	1428	(12.40)	220	(12.69)	901	(14.98)	63	(15.45)	
24	1027	(17.25)	156	(17.82)	590	(22.85)	42	(23.11)	
32	826	(21.46)	126	(22.14)	449	(30.04)	32	(30.84)	

Results of parallel execution of the 3D UNI-DEM for one month on a SunFire E25k supercomputer at DTU. The job user time and the times of the main computational stages in seconds (followed by the corresponding speed-up) are shown in the corresponding columns.

Performance and scalability of UNI-DEM (3D) on a SunFire E15k

3D UNI-DEM on a Sunfire E15000 (96 × 96 × 10) grid, 35 species, CHUNKSIZE=32									
PEs	Advection		Chemistry		Vertical transport		TOTAL		
	time [sec]				(speed-up)		<i>E</i> [%]		
1	4180	—	15685	—	1500	—	21590	—	—
2	2155	(1.9)	7882	(2.0)	744	(2.0)	11058	(2.0)	98%
4	1037	(4.0)	3924	(4.0)	339	(4.4)	5453	(4.0)	99%
8	509	(8.2)	1976	(7.9)	165	(9.1)	2728	(7.9)	99%
16	265	(15.8)	975	(16.1)	86	(17.4)	1377	(15.7)	98%
24	187	(22.4)	648	(24.2)	58	(25.9)	920	(23.5)	98%
32	139	(30.1)	476	(31.7)	47	(31.0)	702	(30.8)	96%
48	127	(32.9)	317	(47.5)	32	(46.9)	690	(31.3)	65%

Results of parallel execution of the 3D UNI-DEM for one month on the SUN cluster (“Lomond”) at EPCC (SunFire E15k). The job user time and the times of the main computational stages in seconds (followed by the corresponding speed-up) are shown in the corresponding columns.

Performance and scalability of UNI-DEM (3D) on an IBM P5 cluster

3D UNI-DEM on the IBM "HPCx" (96 × 96 × 10) grid, 35 species, CHUNKSIZE=32									
PEs	Advection		Chemistry		Vertical transport		TOTAL		
	time [sec]		(speed-up)					<i>E</i> [%]	
1	958	—	8216	—	385	—	9598	—	—
4	246	(3.9)	2064	(4.0)	97	(4.0)	2438	(3.9)	98%
8	136	(7.0)	1035	(7.9)	48	(8.0)	1244	(7.7)	96%
16	76	(12.6)	518	(15.9)	25	(15.8)	637	(15.1)	94%
24	50	(19.2)	345	(23.8)	16	(24.0)	441	(21.8)	91%
32	43	(22.3)	259	(31.7)	12	(32.0)	339	(28.3)	88%
48	34	(28.2)	173	(47.5)	8	(48.1)	242	(39.7)	83%

Results of parallel execution of the 3D UNI-DEM for one month on the IBM high-performance cluster "HPCx" at EPCC, Edinburgh. The job user time and the times of the main computational stages in seconds (followed by the corresponding speed-up) are shown in the corresponding columns.

Introduction to the sensitivity analysis (SA)

Aims of the research

- To propose a new mechanism for sensitivity studies of the concentrations levels of some important pollutants, like ozone (O_3), in real-life scenarios of the air pollution transport over Europe, calculated with Unified Danish Eulerian Model (UNI-DEM).
- To evaluate the output, taking into account the perturbation influences or the relative importance of each input parameter on the output variable.

Mathematical background of the SA study

- The mathematical representation:

$$u = f(\mathbf{x}), \quad \text{where } \mathbf{x} = (x_1, x_2, \dots, x_d) \in U^d \equiv [0, 1]^d$$

is a vector of input parameters with a joint probability distribution function $p(\mathbf{x}) = p(x_1, \dots, x_d)$.

- Local sensitivity analysis: $\mathbf{x} = \mathbf{x}^* \in U^d, \quad u^* = f(\mathbf{x}^*)$
- Global sensitivity analysis: $u = f(\mathbf{x}), \quad \mathbf{x} \in U^d$
(subject to this study)

Total sensitivity index

- Total Sensitivity Index (TSI) of an input parameter x_i , $i \in \{1, \dots, d\}$:

$$S_{x_i}^{tot} = S_i + \sum_{l_1 \neq i} S_{il_1} + \sum_{l_1, l_2 \neq i, l_1 < l_2} S_{il_1 l_2} + \dots + S_{il_1 \dots l_{d-1}},$$

where $S_{il_1 \dots l_{j-1}}$ – j^{th} order sensitivity index for the parameter x_i ($1 \leq j \leq d$),

$j = 1$: S_i – the "main effect" of x_i .

- Classification of the input parameters:

–very important:	$0.8 < S_{x_i}^{tot}$	
–important:	$0.5 < S_{x_i}^{tot} < 0.8$	
–unimportant:	$0.3 < S_{x_i}^{tot} < 0.5$	
–irrelevant:	$S_{x_i}^{tot} < 0.3$	

Sobol' approach

Analysis of Variances (ANOVA) of a square integrable function $f(\mathbf{x})$

High Dimensional Model Representation (HDMR):

$$f(\mathbf{x}) = f_0 + \sum_{s=1}^d \sum_{l_1 < \dots < l_s} f_{l_1 \dots l_s}(x_{l_1}, x_{l_2}, \dots, x_{l_s}),$$

where

- f_0 - constant,
- $\int_0^1 f_{l_1 \dots l_s}(x_{l_1}, x_{l_2}, \dots, x_{l_s}) \times_{l_k} = 0, \quad 1 \leq k \leq s, \quad s = 1, \dots, d.$

Sobol (1969), Sobol (1993)

Sobol' approach (cont.)

Therefore

- $\int_{U^d} f_{i_1, \dots, i_s} f_{j_1, \dots, j_l} \mathbf{x} = 0, \quad (i_1, \dots, i_s) \neq (j_1, \dots, j_l), \quad s, l \in \{1, \dots, d\}$

and the functions in the right-hand side are defined in a unique way:

- $f_0 = \int_{U^d} f(\mathbf{x}) \mathbf{x}$

- $f_{l_1}(x_{l_1}) = \int_{U^{d-1}} f(\mathbf{x}) \prod_{k \neq l_1} \mathbf{x}_k - f_0, \quad l_1 \in \{1, 2, \dots, d\}$

- $f_{l_1 l_2}(x_{l_1}, x_{l_2}) = \int_{U^{d-2}} f(\mathbf{x}) \prod_{k \neq l_1, l_2} \mathbf{x}_k - f_0 - f_{l_1}(x_{l_1}) - f_{l_2}(x_{l_2}),$
 $l_1, l_2 \in \{1, 2, \dots, d\}$

Global (Sobol') sensitivity indices

Definition (Sobol):

$$S_{l_1 \dots l_s} = \frac{\mathbf{D}_{l_1 \dots l_s}}{\mathbf{D}}, \quad s \in \{1, \dots, d\},$$

where

- variances $\mathbf{D}_{l_1 \dots l_s} = \int f_{l_1 \dots l_s}^2 \mathbf{x}_{l_1} \dots \mathbf{x}_{l_s}$,
- total variance $\mathbf{D} = \int_{U^d} f^2(\mathbf{x}) \mathbf{x} - f_0^2$,

and the following properties hold:

- $S_{l_1 \dots l_s} \geq 0$, $\sum_{s=1}^d \sum_{l_1 < \dots < l_s} S_{l_1 \dots l_s} = 1$.

Methods for evaluating global sensitivity indices

Method	Cost (Model runs)	Sensitivities
FAST (1973)	$O(d^2)$	$S_i, \forall i$
Sobol (1993)	$N(2d + 2)$	$S_i, S_{x_i}^{tot}, \forall i$
EFAST (1999)	dN	$S_i, S_{x_i}^{tot}, \forall i$
Saltelli (2002)	$N(d + 2)$	$S_i, \forall i, S_{lj}, \forall l, j, l \neq i$

Sobol' Monte Carlo algorithm

Let $\mathbf{x} = (\mathbf{y}, \mathbf{z}) \in \mathbb{R}^d$, $\mathbf{y} = (x_{k_1}, x_{k_2}, \dots, x_{k_m}) \in \mathbb{R}^m$, $K = (k_1, k_2, \dots, k_m)$.

Variance of the subset \mathbf{y} :
$$\mathbf{D}_{\mathbf{y}} = \sum_{s=1}^m \sum_{(k_1, \dots, k_m) \in K} \mathbf{D}_{l_1 \dots l_s}.$$

Theorem (Sobol).
$$\mathbf{D}_{\mathbf{y}} = \int f(\mathbf{x}) f(\mathbf{y}, \mathbf{z}') \mathbf{x} \mathbf{z}' - f_0^2 \quad \left(f_0 = \int_{U^d} f(\mathbf{x}) \mathbf{x} \right).$$

Monte Carlo algorithm:

$$\begin{aligned} \frac{1}{N} \sum_{j=1}^N f(\xi_j) &\xrightarrow{P} f_0 & \frac{1}{N} \sum_{j=1}^N f^2(\xi_j) &\xrightarrow{P} \mathbf{D} + f_0^2 \\ \frac{1}{N} \sum_{j=1}^N f(\xi_j) f(\eta_j, \zeta_j') &\xrightarrow{P} \mathbf{D}_{\mathbf{y}} + f_0^2, & \xi &= (\eta, \zeta). \end{aligned}$$

$$S_{l_1} = S_{(l_1)} = \frac{\mathbf{D}_{(l_1)}}{\mathbf{D}}, \quad S_{(l_1 l_2)} = S_{l_1} + S_{l_2} + S_{l_1 l_2} = \frac{\mathbf{D}_{(l_1 l_2)}}{\mathbf{D}} \implies S_{l_1 l_2}.$$

Numerical results – Preliminary stage

The input data:

$$r_s(\alpha) = \frac{c_s^\alpha(a_s, b_s)}{c_s(a_0, b_0)},$$

where $\alpha = \{\alpha_i\}_{i=1}^d$, $\alpha_i \in \{0.1, 0.2, \dots, 2.0\}$,

$d = 116$

– all the chemical reactions (full analysis),

$d = 69$

– the time-dependent chemical reactions,

$d = 47$

– the constant chemical reactions,

or any subset of the above 3 sets;

$s = 1, \dots, 35$

– the chemical species (pollutants),

$$\begin{aligned} c_k^{max} &= c_k(a_0, b_0) = \\ &= \max_{(a,b) \in G} \{c_k(a, b)\} \end{aligned}$$

– the maximal mean monthly concentration of a fixed chemical specie k over the domain G .

Numerical results – Parameter values

$$d = 3$$

$$i_1 = 3$$

$$i_2 = 6$$

In our case: $i_3 = 22$



$$k = 13 \quad (\text{ozone concentrations})$$

C_s – mean monthly concentrations for July 1998

Mean monthly concentrations of O_3 and NO_2 for July 1998


Distribution of ozone concentrations in some parts of Europe for July 1998

Results obtained by UNI-DEM (35 species)
(96x96) grid / 50 km x 50 km surface cells
Emissions and meteorological data: July 1998
UNITS: ppb, File: Basic (k=1.0)
Maximal value in the domain: 97.4
Minimal value in the domain: 1.9

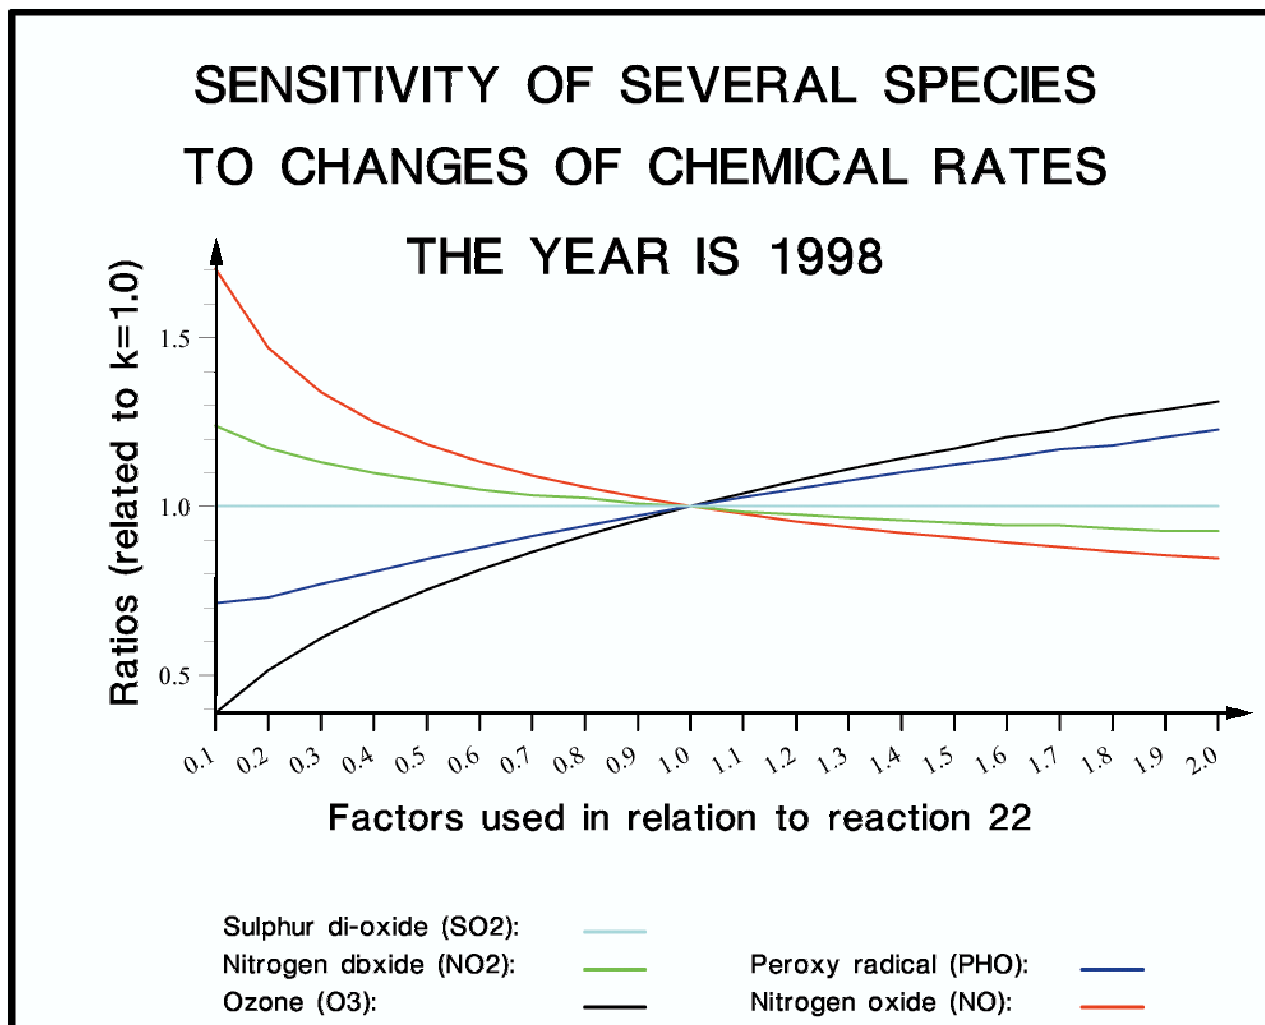


Distribution of NO_2 concentrations in some parts of Europe for July 1998

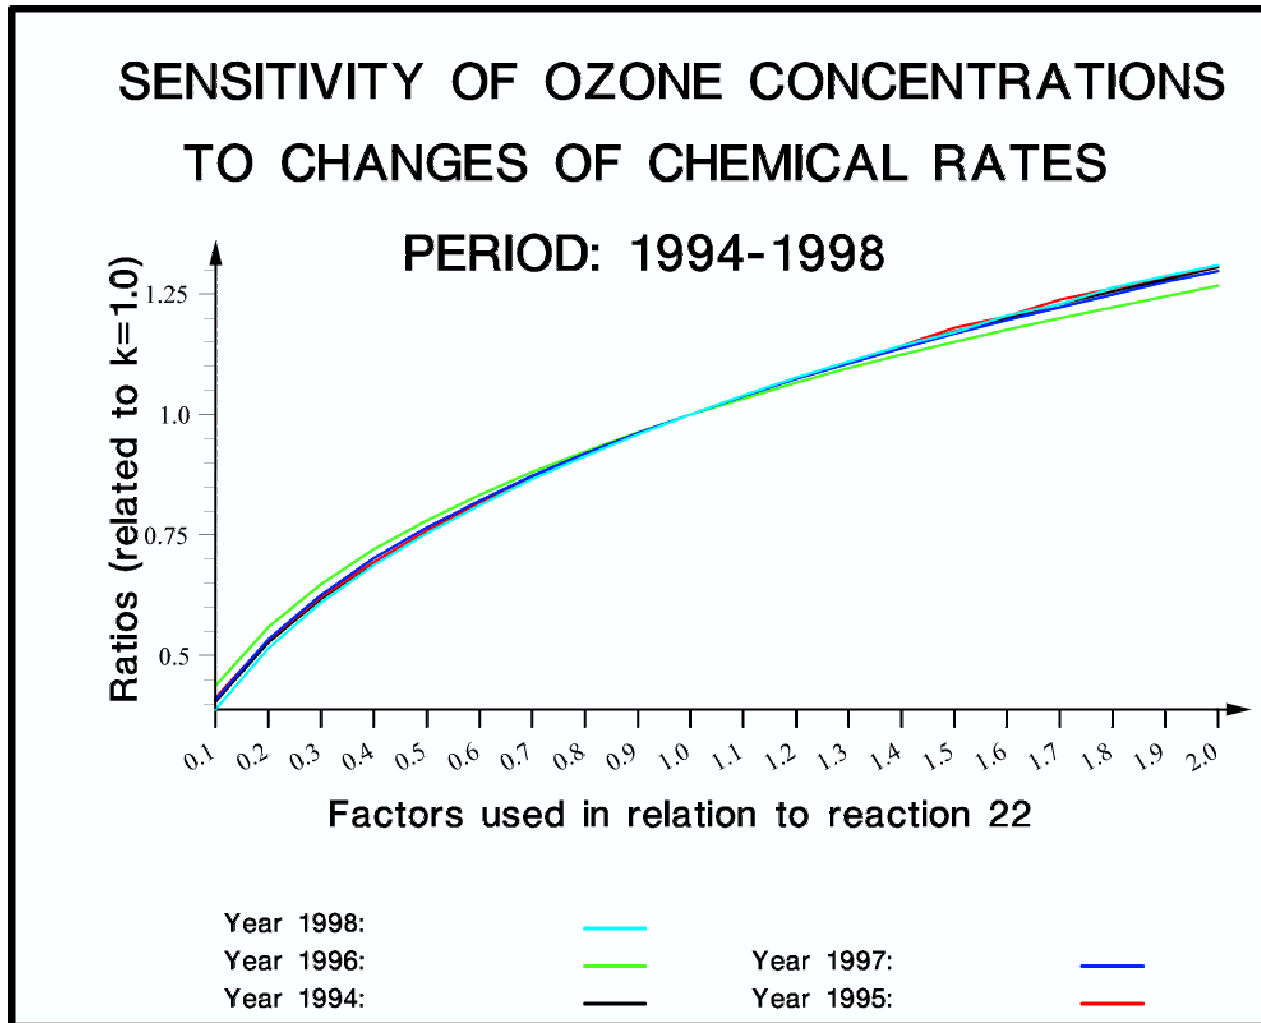
Results obtained by UNI-DEM (35 species)
(96x96) grid / 50 km x 50 km surface cells
Emissions and meteorological data: July 1998
UNITS: ppb, Reaction 8 with factor 1.0
Maximal value in the domain: 13.90
Minimal value in the domain: 0.02
C(27,41): 7.66
C(31,39): 9.10



Sensitivity of 5 chemical species to the rate of chemical reaction # 22 (July 1998)



Sensitivity of ozone concentrations to changes of chemical rates in 5 consecutive years



Conclusions

Scalability experiments on several state-of-the-art supercomputers has been presented. The following conclusions can be drawn from the results:

- The parallel MPI implementation of DEM, based on horizontal partitioning of the domain, is highly portable, well balanced and proved its efficiency on some of the most powerful parallel supercomputers in Europe.
- When using the full capacity of the SUN cluster “Lomond” (all its 48 back-end processors), a significant breakdown in the efficiency (to about 65% from more than 95% for up to 32 processors) is observed. The bottleneck appeared to be the communication network, which shows certain slowdown in its work on full load.
- The efficiency and speed-up is relatively higher in the computationally-intensive stages. In particular, the chemistry stage has almost linear speed-up (and it takes up to 85% of the total time on the “HPCx”).

The so-called preliminary sensitivity analysis experiments with perturbation of the reaction coefficients have created the necessary input data for some of the most advanced variance-based approaches for sensitivity analysis. Such methods have been applied to study the influence of chemical rates variation over the concentration levels of air pollutants using UNI-DEM. This study would possibly lead to improvements in the chemical scheme of the model as well as to its more efficient use in some time-critical applications.

- Applicability of the results
 - for a verification and an improvement of the model;
 - for a reliable prediction of the final output.
- Future plans
 - other approximation tools;
 - computations with 3D version of UNI-DEM;
 - studies of model sensitivity with respect to the the emission levels and boundary conditions.

Acknowledgments

This research was supported in part by the Bulgarian IST Centre of Competence in 21 Century – BIS-21++ (contract # INCO-CT-2005-016639), by the NATO project "Monte Carlo Sensitivity Studies of Environmental Security" (PDD(TC)-ESP.EAP.CLG 982641), and by the HPC-Europa project (RII3-CT-2003-506079) with the support of the European Community - Research Infrastructure Action under the FP6 "Structuring the European Research Area" Programme.