

Aerosol Dispersion Modeling with Implicit Probabilistic Particle Tracking Algorithm

Andrei Smirnov, Steven Rowan, James McCormic
Department of Mechanical and Aerospace Engineering
West Virginia University, Morgantown, WV 26506-6106

Abstract

A method is proposed which facilitates express analysis of the results of large distributed simulations of physical systems. The method uses the domain decomposition technique based on the idea of transition probabilities. The method opens the possibility of conducting parallel simulations of strongly coupled systems on loosely coupled distributed computer platforms, such as grid computing environments. The results of the simulations can then be retrieved in a timely manner for express risk analysis. The case of aerosol dispersion in urban environments was considered, and a prototype solver was developed and tested. Fast response of the system is achieved by replacing complex 3D simulations with information retrieval from aerosol dispersion database. The database was created by extensive simulations of different aerosol dispersion scenarios in a grid computing environment. The reduction of system response from a few days to several seconds can be achieved compared to real-time 3D simulations.

Keywords: Physically based modeling, Grid computing, Web Interfaces, Risk analysis, Information Retrieval, Aerosols

1. Introduction

The emerging grid computing technologies offer a great promise for a number of applications, such as distributed databases and transaction processing. However in the field of scientific computing these systems often present formidable obstacles, since they are essentially based on loosely coupled collection of computing nodes. The lack of adequate communication bandwidth prevents their effective usage in solving tightly coupled problems, such as those based on various continuum approximations and partial differential equations (PDE).

It is well recognized that parallel or distributed computing is not well suited for many types of problems, especially in the field of high-performance (HP) com-

puting, because of tight coupling or procedural bottlenecks in the solution methods used.

Computer clusters can be designed to suit the need of HP computations by using high-bandwidth communication switches and fast local area networks (LANs). However, this approach reaches its limits as the number of nodes increases to the point where their cumulative failure rate becomes a problem, as well as the surface-to-volume ratio effect increases the communication overhead. In addition to these, in the limit of a grid computing environment, one can no longer guarantee the required high-bandwidth, since the underlying communication network is inherently slow.

In this study a variant of domain decomposition technique based on *domain transition probabilities* (DTP) was developed and successfully applied to facilitate solutions to some of tightly coupled problems in grid computing environments. The technique was applied to the problem of express analysis of aerosols dispersion inside a city landscape.

This problem closely relates to environmental issues, and has also received a special attention in the United States following the terrorist attacks on September 11, 2001. The U.S. Department of Homeland Security has sponsored planning scenarios in order to help prepare municipalities for the possibility of an event related to the release of biological, chemical, or radioactive material in urban areas. One approach is to conduct computer simulations of these scenarios using computational fluid dynamics (CFD) [7]. While these simulations can provide very accurate results, the time required to run a simulation of a specific incident might very well take from several hours to several days to complete, depending upon the size of the simulation grid and the computing resources at hand. Considering the time critical nature of the need to respond to these incidents, it is all too likely that the information necessary for first responders to carry out their jobs would not be available until after the fact. Realizing this limitation, the Department of Defense has requested a large number of disbursement scenarios in major cities []. In this way, they would be able to respond to a specific event by utiliz-

ing the scenario that most closely matches the incident in question.

In solving problems of this type CFD is not the only method used, and more often it is a combination of different techniques. Another method most commonly used in conjunction with CFD is Lagrangian particle dynamics [2], which provides a model for particle motion. One can find a number of other techniques, suitable for simulating the dispersion of aerosols, such as a Multiple Relaxation Time Lattice Boltzmann Model (MRTLBM) along with a hybrid thermal MRTLBM to model thermal effects upon the flow fields [3], and many more.

However, running a time-critical simulation based on a physical model can be very time-consuming. An alternative to this approach adopted in this study is to run a series of simulations using a cluster or grid-computing network to generate a set of data from which accurate predictions of the dispersion of a contaminant could be made in a timely manner [7]. In other words, utilizing a computing grid, it is possible to run simulations of multiple dispersion situations on parallel systems in order to quickly develop data sets for a wide range of scenarios. With these data sets, a fast response of the system can be achieved by replacing complex 3D simulations with information retrieval from a database. An easy to use web-based application will be needed to submit queries to the database and display the results. This application would allow the user to select a location from which a contaminant has been released and obtain accurate results in less time than it would take to run the simulation. In this way, not only can the issue of accurate dispersion scenario be addressed accurately and quickly, it would also allow emergency first responders to immediately access this information stored on a common media (CD, DVD) or via a laptop computer and a wireless Internet connection from close proximity to the actual contaminant release location.

2. Method

From the previous discussion it follows that there are two issues involved in solving a particular problem of aerosol dispersion in a city on a grid computing system: (1) tightly coupled nature of the problem does not fit with the loosely coupled parallelism of the grid, and (2) the simulations need to be performed in a timely manner, which is hard to accomplish given the physical complexity of the problem.

In a method suggested in this study both issues are addressed at the same time using the idea of *domain transition probabilities* (DTPs). Formally speaking, the DTP is defined as a probability for an event A occurring at the boundary of a domain to cause an event B at some other point on the boundary.

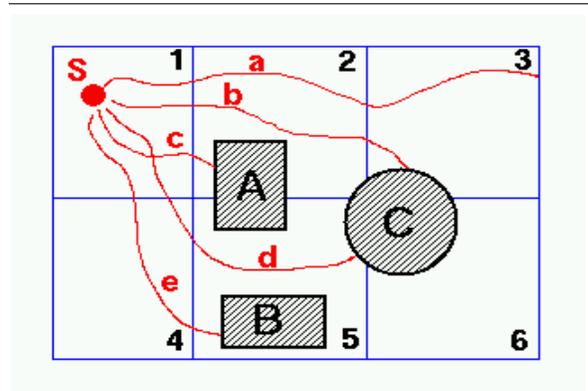


Figure 1. Particle passage through a decomposed domain: S - particle source, A...C - physical objects where the particle fallout occurs, a...e - particle trajectories, 1..6 - domains of a decomposed computational space.

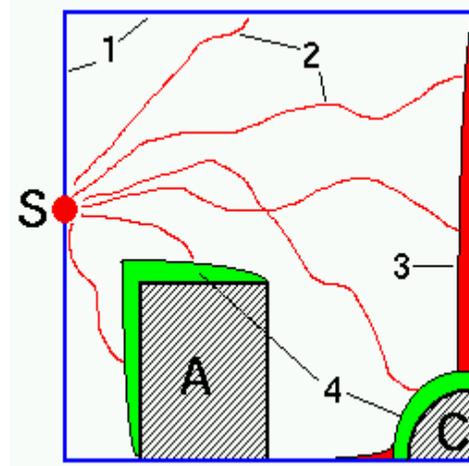


Figure 2. Assembling domain transfer probabilities: S - boundary particle source, A-C - objects, 1 - domain boundaries, 2 - particle trajectories, 3 - source-to-boundary transfer probability, 4 - source-to-object transfer probabilities.

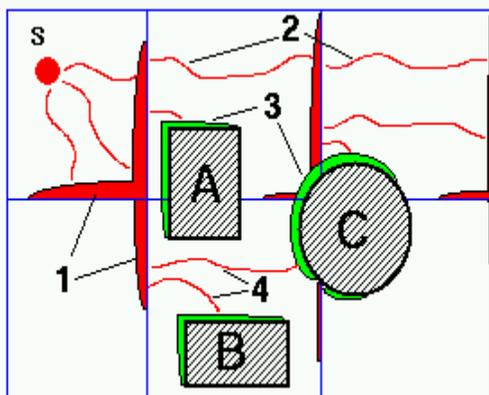


Figure 3. Implicit particle tracking using DTPs: 1 - boundary particle sources, 2 - boundary-to-boundary transfer probabilities, 3 - object sinks, 4 - boundary-to-object transfer probabilities,

In this work the issues discussed above have been addressed and implemented in a prototype risk assessment system for the analysis of aerosol release and spreading in urban environments. The system consists of a simulator and front-end interface. The simulator incorporates a physically based model of aerosol transport and is executed on high-performance computing facilities, such as computer clusters or grid environments. It collects the data of possible aerosol release scenarios. The front-end interface runs as a web-based application and retrieves the data for particular scenarios and their outcomes. There are several aspects required to implement this method:

1. Conceptual model of possible scenarios.
2. Physical model enabling accurate simulations of events of interest.
3. The possibility to conduct massive and possibly exhaustive sets of simulations.
4. Data reduction and efficient storage.
5. Data retrieval and visualization techniques.

2.1. Conceptual Model

A comprehensive answer to the problem of aerosol dispersion can be provided by either conducting experiments or computer simulations by means of CFD. Both may require a long time to complete for a realistic city landscape. The outcome is usually given as three dimensional time dependent probability density functions of

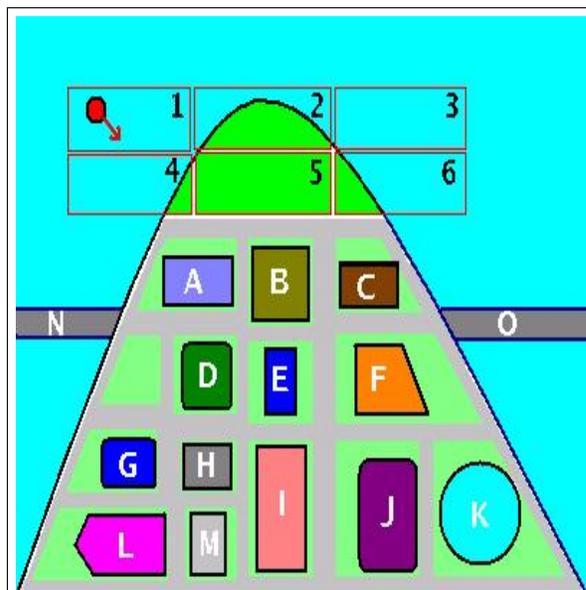


Figure 4. Aerosol dispersion in a city: 1..6 - available aerosol release data; A-O - affected objects.

aerosol distribution. However, from a practical perspective the important question is not the level of contamination at each location, but rather to what degree each particular object is affected. Since the number of identified objects can be significantly less than the total number of discrete locations, this formulation can lead to considerable reduction of the size of data. As a consequence, it will result in an increase in the speed of the data retrieval process, which is an important factor for express risk assessment, especially when the database is accessed remotely.

Figure 4 shows a generic city map where different objects of importance, such as buildings and bridges are marked by letters (A,B,C...). The results of aerosol simulations are kept in a remote database of hit counts on all objects, corresponding to all possible release scenarios. Since the whole database is too large to be transferred over the network at once, the data are split into chunks, or *data-sets*, corresponding to limited areas, marked by numbered squares (1,2,...) in Fig. 4. The size of each data set is selected so as to minimize the time of data Transfer, as well as to provide a representative set of release scenarios in a given location (see Sec.3). For each a particular scenario to be investigated, the corresponding data set is retrieved, including all of the scenarios from the local area. Thus, multiple scenarios from the same location can be investigated without the need of subsequent data transfers.

Data corresponding to each release scenario contain the degree of fallout on every object in the domain of interest (i.e. downtown area). In this format the size of the data set is much smaller than it would be if three-dimensional probability density functions were used.

2.2. Physical Modeling

A simple aerosol transport model can be built from the equation of particle motion, expressed in terms of particle velocity, $\mathbf{v}(\mathbf{x}, t)$, in a given mean flow field, $\mathbf{u}(\mathbf{x}, t)$:

$$\frac{d\mathbf{v}}{dt} = C_D(\mathbf{u} - \mathbf{v}) + C_T\mathbf{u}' - \mathbf{g} \quad (1)$$

where C_D is the drag coefficient, C_T the turbulent diffusivity, \mathbf{u}' is the instantaneous turbulent fluctuation vector, and \mathbf{g} is the gravity acceleration vector [1].

The position of the particle at each time step is computed using second order Runge-Kutta time-stepping scheme:

$$\mathbf{x} = \mathbf{x}_0 + \frac{\mathbf{v} + \mathbf{v}_0}{2} dt$$

where $(\mathbf{x}_0, \mathbf{v}_0)$ and (\mathbf{x}, \mathbf{v}) are the old and new coordinates and velocities of a particle. The effects of turbulent dispersion on the particles encapsulated in the second term of (1) can be accounted for by an appropriate sub-grid scale turbulence model. In particular, the RFG technique developed earlier by the authors enables one to simulate realistic effects of turbulence on the aerosols and bubbles [4]. A particular advantage of this approach is that the effects of turbulence can be completely decoupled from the flow solver. Thus, in a first approximation the effect of turbulent dispersion can be simulated by applying turbulent fluctuations to aerosol particles, which is superimposed on the flow velocity field, \mathbf{u} , which is set constant inside the domain.

Each aerosol particle is convected in a velocity field \mathbf{u} , and is traced inside the computational domain until it crosses the domain boundary. In the second event the hit count for that object is incremented. The geometry of the landscape is given by a 3D Cartesian grid, where each cell of the grid is represented by a single integer number, identifying an object, which occupies the corresponding region of space.

2.3. Distributed Simulations

Given the large size of the grid that is required to represent a realistic urban Environment, as well as a variety of other factors such as wind strengths and direction, aerosol volatility, turbulence level, etc., a large number of aerosol release scenarios, N_s , need to be played. A

distributed computing facility, such as a Beowulf cluster or a grid computing environment will be the best choice for these simulations.

The number of particles released for each realization will determine the accuracy of predictions, since the variance of the result, σ , is inversely proportional to the square root of the number of particles $n(x)$ deposited at position x :

$$\sigma(x) \approx 1/\sqrt{n(x)} \quad (2)$$

Consequently, one may expect low accuracy of predictions in the regions where only few particles have reached. Given the already large number of scenarios, the number of released particles for each scenario, N_p should be selected as a trade-off between the desired accuracy and the available computing time and resources.

If one wants to keep the statistical uncertainty in (2) below a certain level, σ_{max} for most parts of the domain, then with the N_n grid nodes the appropriate number of injected particles, N_p , can be estimated from the condition $n(x) = C N_p/N_n$, where factor C determines the volume fraction of the affected area to the total area of the domain. When applied to (2) this condition provides the estimate for the total number of particles needed:

$$N_p = \frac{C N_n}{2\sigma_{max}^2} \quad (3)$$

If the average time for computing a single trajectory, τ , is given, then the time required for computing all N_s scenarios will be:

$$t = N_p N_s \tau / N_c \quad (4)$$

where N_c is the number of computing nodes on a cluster. It should be noted that this problem is especially suitable for parallel simulations in cluster of grid computing environments, since it involves statistical data sampling for completely independent realizations. Therefore, no inter-process communications are necessary, and the problem belongs to a so-called *embarrassingly parallel* category.

2.4. Data Reduction and Compression

Additional consideration should be given to the storage of the resultant data sets. Since particle deposition counts are recorded for each scenario, the resultant data set will consist of N_s numbers, which with a four-byte storage for an integer number will bring the size of the data set for a single object to:

$$D_o = 4N_s \quad (5)$$

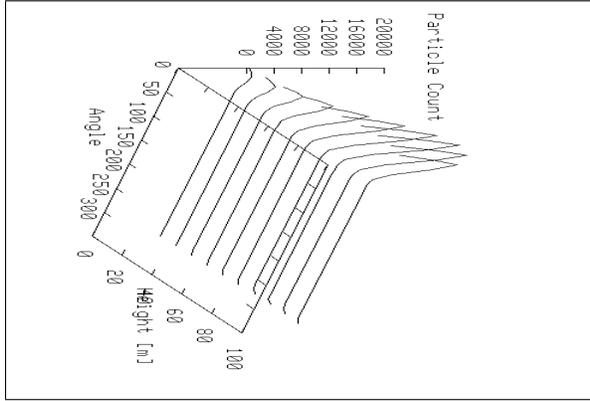


Figure 5. Typical particle count on an object for different space locations and angles.

And this much data should be given for each object in the domain, i.e. building, bridge, etc. Considering all objects of interest, N_{obj} , the size of the data set will be:

$$D = N_{obj} D_o \quad (6)$$

which can become a very large number even in a simplified test case (Sec.3). One way to deal with large data sets is to arrange a fast database access to these data. However, this will require a database retrieval for each value of position coordinate and angular direction, which may involve many database accesses. Secondly, with an increase in the size of the domain the size of the database files will become prohibitively large. A distributed database on a cluster may be a way out. However, there is a cheaper and simpler solution using reduced data sets and compressed files.

The size of the data set can be reduced significantly considering that the data on angular distribution are usually sparse, that is, many wind directions carry particles away from the objects causing zero deposition counts. Figure 5 shows a typical distribution of the fallout data on a particular object as a function of angle and one of the three spatial positions (height). The dependence in two other spatial directions shows similar features. It can be seen from this figure that if one stores only positive counts and discards zero counts, the size of the data set can be reduced. On the average, a reduction by about a factor of 4 is possible. Another size reduction can be achieved by using data compression. Standard data compression methods, such as LZW, LZ77, and others can easily reduce the size of the file by a factor of 3 or more. Thus, the size of the data set can be reduced to:

$$D_r \approx \frac{D}{4 \cdot 3} = \frac{N_{obj} D_o}{12} \quad (7)$$

2.5. Data Retrieval and Visualization

From the perspective of express risk analysis a preferred situation would be a possibility of direct access to the simulation data on a local workstation, laptop, or via the Internet. In addition to this the application should be platform independent, and should run equally well on different operation systems. A simple solution will be an efficient web interface to aerosol dispersion data. This interface should provide a reasonably short access time to the data no matter how large the total data set can be, and at the same time it should be simple and intuitive to use. Java technology provides an appropriate environment for the development of such interface. Java applications are platform independent, they can be used to access local or remote data sets, provide intuitive user interfaces and capable of 3D graphics, and can run as local or web-based applications (applets). Thus Java is a natural choice for data retrieval and visualization.

3. Results and Discussion

A generic city landscape based upon the Pittsburgh downtown area was constructed using the voxel-based 3D graphics system [6, 5]. The whole domain was discretized on the grid with spacial resolution of $N_x = 92$, $N_y = 92$, $N_z = 32$, which corresponds to the total of $N_n = N_x N_y N_z = 270848$ nodes or possible particle release points. The landscape was populated with characteristic features like rivers, hills, bridges, park area, pavements and buildings (Fig.4). To examine all possible scenarios of aerosol release, the injection of airborne particles was simulated at each grid node with a total number of $N_a = 36$ aerosol release scenarios for each node, which corresponds to 36 wind directions spanning the whole circle of 360° . Thus, to exhaustively cover all the possible scenarios for this particular domain size and angular resolution the total of $N_s = N_n N_a = 9750528$, or about ten million realizations of positions and wind directions has to be analyzed on particle dispersion.

To keep the average statistical error close to 10%, and estimating the typical area coverage coefficient in (3) as $C \approx 0.01$, we arrive at an estimate of required particle count as: $N_p \approx 10^5$.

On a single 1GHz node of a Beowulf cluster the computation of one particle trajectory with our scheme in a given landscape takes on the average $\tau = 4.5 \times 10^{-5}$ sec. Thus computing all the $N_s = 9750528$ scenarios with $N_p = 10^5$ particles will require according to (4): $t = N_p N_s \tau \approx 5000$ days, or 14 years of computing time. If one uses a computer cluster with 100

nodes this time can be reduced to under 50 days of simulations, which will still be prohibitively large.

In addition to this, according to (5) the size of the data for a single object will be also too large for efficient Internet transfer, and close to: $D_o = 4N_s \approx 40MB$. The size of the whole data set as determined by (6) can easily approach $D \approx 1GB$, which makes it also impractical for Internet access. Using data reduction technique described in Sec.2.4 (See (7)), the size of this data set can be reduced to $D_r = D/4/3 = N_{obj}D_o/4/3 \approx 20 * 40/4/3 \approx 67MB$. This is a considerable reduction, but still not quite enough to achieve reasonably short transfer times.

A faster Internet access to this information can be achieved if the data set is split among several data files, one for each object. In our case with $N_{obj} = 20$ objects this will reduce the data size for each object to: $D_{or} = D_r/N_{obj} \approx 3.3MB$.

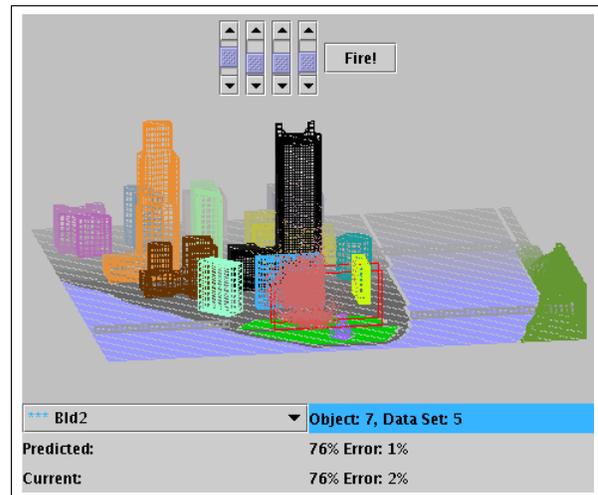
An even greater reduction of retrievable data size can be achieved if the data for the whole domain is split into chunks or sub-domains. Such sub-division makes sense if one considers that usually the express information on the contaminant's release is needed for a specific area where the release has occurred. From this perspective, the size of the sub-domains can be selected so as to minimize the file transfer time, as well as to provide access to a representative sample of data.

To support a reasonably fast web access it was decided to maintain the size of each data file within $F \approx 100KB$. It was found by experimentation that the data file of this size approximately corresponds to the sample size of around $N_i = 10^3$ injection points. Strictly speaking, the size of data sample F for a fixed number of injection points, N_i may vary, since the data arrays with deposition counts for different angles are represented by variable size lists to conserve space as discussed above. However, this variation is not significant. Thus, by subdividing the domain into chunks of $N_i = 10^3$ nodes we obtain the total of $N_{chunks} = N_n/N_i \approx 270$ sub-domains. A complete simulation of all possible scenarios with $N_p = 10^5$ particles in each realization for one such sub-domain will take about:

$$t_{chunk} = N_i Na Np\tau \approx 45 \text{ hrs}$$

on a single cluster node. The computation of the whole domain on a 100-node cluster will take $N_{chunks}t_{chunk}/N_{cluster} \approx 122 \text{ hrs}$ of cluster time. This estimate presumes that several sub-domains can be executed on a single processing node.

In the current study we used $N_{chunks} = 12$ sub-domains for possible aerosol release locations, with each sub domain containing $N_i = 1000$ aerosol release locations. In parallel simulations one sub-domain domain was assigned per computing node. The hit count data for each sub-domain were split between N_{obj} files,



<http://mulphys.com/sim/demo>

Figure 6. Web interface to simulate aerosol release in a city

each corresponding to one object in the domain. With the number of objects $N_{obj} = 20$ this resulted in 240 files, which were compressed with a standard gzip compressor and stored on a remote server.

The simulator of aerosol transport and dispersion was implemented in C++ language and executed on a Beowulf cluster with 1GB, 1GHz computing nodes. The simulations were performed for different number of particles released per realization: 10^2 , 10^3 , 10^4 , and 10^5 . Thus, the total of 4 runs were conducted. Total processor time required per each run is given in Table 1.

N_p	10^2	10^3	10^4	10^5
CPU [min]	3	27	266	2695

Table 1. Wall clock time of parallel runs for different number of particles.

To demonstrate the viability of this method a prototype web interface was developed, which enabled us to test different scenarios of aerosol release and dispersion (mulphys.com/sim/demo). The interface is written in Java language and provides a 3D representation of a city landscape with the possibility of navigating through the landscape, arbitrary positioning of the aerosol source, and setting wind direction (Fig.6). The applet also performs a real time simulation of aerosol propagation and dispersion for a limited number of particles.

The count of particles deposited on the selected object is constantly updated as the simulation is progress-

ing. At the same time, the original count of particles retrieved from the cluster simulation data is displayed alongside the real-time simulation count. For the purposes of better interpretation, the count of particles, n , for each object is converted into the deposition probability by dividing it by the total number of particles released, N . The statistical error of the results is shown alongside with the probability count as being proportional to $n^{-1/2}$, according to (2).

The simulation in the applet is performed with a limited number of particles, and serves the purpose of comparing its execution time with a fast retrieval of the pre-assembled simulation data. The accuracy of the applet results is also less than that obtained from the data, since the latter were performed with a much higher number of particles.

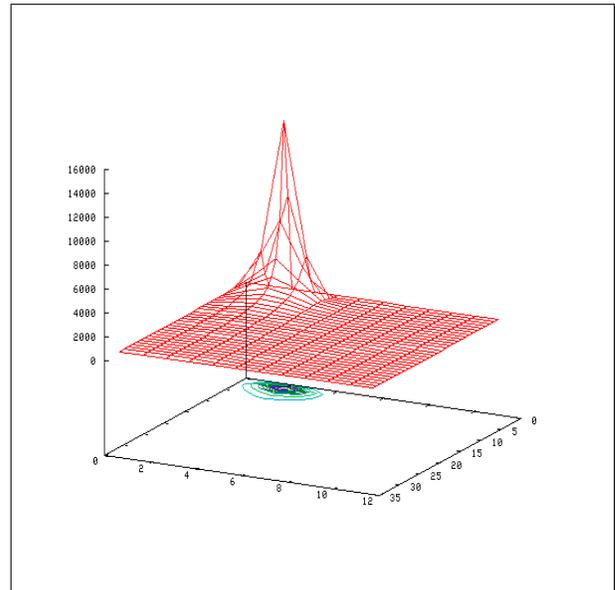
When tested on a 1GB 1HG workstation with 100MB Internet connection the speed of retrieval of each dataset was on the average 10 to 20 times faster than the Java execution of the simulation for 10^4 particles. Considering that the retrieved data set consisted of 1000 realizations, it leads to a much greater speedup if multiple local scenarios need to be analyzed.

3.1. Discussion

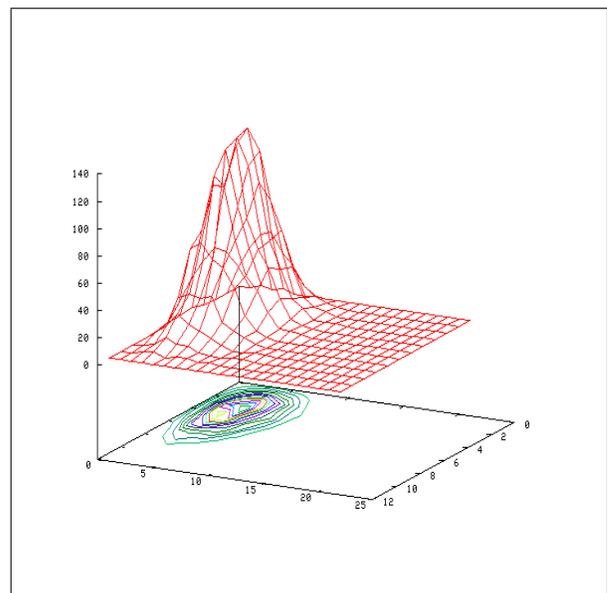
high order linear extrapolations can be made that would accurately predict the results of scenarios that are similar to, but not specifically modeled by the previous simulations.

It should be noted that the current method works under the assumption of a static landscape. Each change in the landscape will require recompilation of the database, or the parts of it which were directly affected by the changes.

It should be noted that the method proposed here is not a panacea for all types of tightly coupled problems. It is supposed to work well for extensive parametric studies of the type described here. However, it may have very little advantage if any in the cases when just a single aerosol dispersion scenario needs to be computed. Even though a tightly coupled simulation may be excruciatingly slow in a grid computing environment, a completely decoupled computation of all DTPs followed by implicit particle tracking will not be efficient either. In this case the preferred method will be determined by the size of the system and the number of nodes engaged in a simulation. The advantage of the current method is that the generated sets of DTPs can be repeatedly reused to analyze a variety of different scenarios.



(a) Narrow



(b) Wide

Figure 7. Typical DTPs of aerosol dispersion

4. Conclusions and Future Work

The method presented in this study serves double purpose by (1) providing a solution to tightly coupled problems in grid computing environments and (2) facilitating express risk analysis of multiple scenarios for a complex physical event.

The method was applied to a problem of express analysis of aerosol contamination in the city. The results show that this method provides a viable and efficient tool for fast analysis of different contamination scenarios.

A significant saving of retrieval time and data space was achieved by querying objects rather than particular space locations for fallout data. If a more differentiated approach is needed, this approximation can easily be refined by splitting large objects into smaller ones, like buildings can be split into floors, etc.

The method is especially effective when implemented in multi-processor and distributed computing environments. Grid computing environment offers a particularly elegant implementation, which can be done based on the Monte-Carlo sampling. In this case random realizations of aerosol release scenarios are played on each processing node. The results are then combined in a database of aerosol distributions for each scenario. The accuracy of the distributions will constantly improve as the simulations progress and more data points are added. Eventually histograms can be constructed from these aerosol distributions with a desired resolution and accuracy. The method will be particularly efficient in grid computing environments, since it presumes an independent statistical processing on each node, and thus enables a completely decoupled execution.

Simulations of aerosol propagation in a generic city landscape were conducted on a Beowulf cluster and sampled data sets were obtained. After the appropriate reduction and compression of these data sets it was possible to demonstrate the feasibility of fast retrieval of these data via the Internet for express analysis of different scenarios. A prototype web interface for express simulations of particle deposition scenarios was developed and tested (<http://mulphys.mae.wvu.edu/sim/demo>).

References

- [1] C. Crowe, M. Sommerfeld, and Y. Tsuji. *Multiphase Flows with Droplets and Particles*. CRC Press, 1998.
- [2] C. Crowe, M. Sommerfeld, and Y. Tsuji. *Multiphase Flows with Droplets and Particles*. CRC Press, 1999.
- [3] F. Qiu, Y. Zhao, Z. Fan, X. Wei, H. Lorenz, J. Wang, S. Yoakum-Stover, A. Kaufman, and K. Mueller. Dispersion simulation and visualization for urban security.

In *VIS '04: Proceedings of the conference on Visualization '04*, pages 553–560, Washington, DC, USA, 2004. IEEE Computer Society.

- [4] A. Smirnov, S. Shi, and I. Celik. Random flow generation technique for large eddy simulations and particle-dynamics modeling. *Trans. ASME. Journal of Fluids Engineering*, 123:359–371, 2001.
- [5] A. Smirnov, H. Zhang, A. Burt, and I. Celik. Fuel-cell simulator interface. *Journal of Power Sources*, 138/1-2:187–193, 2004.
- [6] A. Smirnov, H. Zhang, and B. Sowers. Voxel-based volume graphics system for multi-physics modeling. In *The 8th World Multi-Conference on Systemics, Cybernetics and Informatics*, volume 5 of *Computer Science and Engineering*, pages 144–149. International Institute of Informatics and Systemics, 2004.
- [7] J. Thilmen. Harm's way. *Mechanical Engineering*, 127:22–24, 2005. bioterror CFD.