

Genetic Economy, Symmetry, and Computational Complexity for Atomic Virus Structure Determination

Dan C. Marinescu, Yongchang Ji, Gabriela M. Marinescu
School of Electrical Engineering and Computer Science
University of Central Florida
Orlando, Florida,
Email: [dcm, yji, magda]@cs.ucf.edu

Abstract

Genetic economy dictates that complex biological structures like viruses exhibit some degree of symmetry and are composed of sub-units repeated throughout the structure. This symmetry has a profound effect upon the complexity of computations required for structure determination from experimental evidence. In this paper we report on the relationship between the symmetry and computational complexity for 3D atomic structure determination of viruses, from experimental data provided by electron microscopy.

1. Introduction

The knowledge of the 3D atomic structure of viruses is invaluable for the design of anti-viral drugs, [Bra91]. The experimental evidence regarding the atomic structure of viruses can be gathered using crystallographic methods; in this case the biological material sample is exposed to a beam of X-rays and we obtain and analyze an X-ray diffraction image. An alternative approach is provided by electron microscopy; now the sample is exposed to a beam of electrons and we obtain and then analyze a micrograph.

While crystallography is capable to produce high resolution electron density maps, up to 1 – 1.5 Å and routinely at 2.5 Å resolution, the crystallization of viruses is a rather tedious and sometimes impossible task. On the other hand, the samples of biological materials for electron microscopy are easier to prepare, but routinely the resolution is limited to 15 – 20 Å. More recent results [Böt97], [Con97] report reconstructions at 7 – 7.5 Å and there is real hope that this limit can be pushed to the 3 – 5 Å resolution range.

The 3D atomic structure determination using experimental evidence collected through electron microscopy, [Fra96], [Bak99] involves several steps: identification of individual particle projections, determination of the orientation of the

virus particle, for each projection, and finally the 3D reconstruction. The last two steps are performed iteratively at increasing resolutions, up to the limit allowed by the accuracy of the experimental data.

Crystallography as well as electron microscopy are data intensive scientific investigation techniques and require the use of parallel computers. Our group has been involved in the development of parallel algorithms for atomic structure determination of spherical viruses since early 1990's, [Cor93], [Cor95], [Mar97], [Mar98]. More recently we have initiated a program to develop parallel algorithms for the study of asymmetric objects using experimental information from electron microscopy, [Lyn99], [Lyn00], [Mar01] and environments to facilitate the use of parallel programs [Bol00] for structural biology studies.

The study of macromolecules like viruses that do not exhibit symmetry is computationally more demanding; the lack of symmetry implies that we must collect a larger volume of experimental data to determine the 3D atomic structure. A virus exhibiting some form of symmetry is built out of few sub-units and we only need to determine the electron density map of one sub-unit. Instead of reconstructing a sub-unit and then applying the symmetry operators to reconstruct the entire virus, we have to conduct the 3D reconstruction for the entire volume of an asymmetric virus.

In this paper we report on a parallel algorithm for orientation refinement and discuss the computational complexity of orientation refinement for non-symmetric particles. The lack of symmetry and its effect upon the algorithms for 3D reconstruction is discussed elsewhere, [Mar01].

This paper is organized as follows: Section 2 provides a succinct introduction to virus structures and symmetry. In Section 3 we outline the basic methodology for electron density refinement in electron microscopy and in Section 4 we present an algorithm for parallel orientation refinement and evaluate its complexity. In Section 5 we present some experimental results confirming that our algorithm is sound,

it is capable to identify the best match between a view with unknown orientation and a set of views with known orientations. Finally, we present our conclusions.

2. Virus Symmetry and Genetic Economy

Viruses are large macromolecules that cause various human, animal, and plant diseases. Viruses infect healthy cells by attaching to them, then delivering their nucleic acid to the host cell. Instead of expressing the cellular gene, the biosynthetic system of an infected cell is forced to facilitate the replication, transcription, and translation of the viral gene. Then the virus particle created in the infected cell leaves it to infect another host.

A virus consists of a nucleic acid genome, a protein shell or *capsid*, and sometimes a membrane called an *envelope* enclosing the capsid. For example, in Figure 1 (a) we see a schematic diagram of a spherical virus, called *picona*, with a diameter of about 300 Å, and the radius of the RNA of about 110 Å. In Figure 1 (b) we show a schematic representation of a large enveloped virus called *sindbis*, with several spiked proteins protruding out of the surface. The external diameter of the *sindbis* virus is of about 700 Å, the diameter of the neurocapsid core of about 400 Å, [Bra91].




Figure 1. Viruses differ in terms of genome, molecular mass, relative weights of the protein shell and the genome, shape, and size as shown by the schematic diagrams of two spherical viruses. (a) A *picona* virus with a diameter of about 300 Å. (b) A *sindbis* virus with a diameter of about 700 Å.

Viruses differ in terms of genome, molecular mass, relative weights of the protein shell and the genome, shape, and size. The genome consists of either RNA or DNA. The molecular weight differs; for example the molecular weight

of the *picona* viruses is of about 8.5×10^6 daltons; about 30% of the weight is due to one long RNA molecule with about 8,000 nucleotides. There are spherical, cylindrical, as well as other shapes of viruses. The size of a virus ranges from say 150 to 900 Å.

The protein shell of a spherical virus is made out of protein sub-units that are symmetrically arranged, while the nucleic acid genome lacks symmetry. We have more detailed knowledge of the viral capsids than about the nucleic acid genome because it is easier to collect experimental evidence about the structure of a symmetric object than about an asymmetric object. For example, X-ray diffraction is used for 3D atomic structure determination of the protein shell, but the nucleic acid genome is not seen in virus X-ray diffraction images.

A virus has to locate a specific docking site on a host cell. The information about the structure of the capsid is critical in determining the compounds that can block the virus binding site. This explains the interest of pharmaceutical companies in the atomic structure of viruses.

The symmetry of the protein shell means that we have a relatively small number of identical building blocks that recognize each other and are able to assemble together spontaneously. The principle of genetic economy requires that the shell be built out of a few copies of identical units; the amount of genetic information, thus the size of the genome is considerably smaller for a symmetric virus particle. For example, one of the smallest known viruses, the *satellite tobacco necrosis* virus has a diameter of 180 Å, a protein shell of 60 sub-units, and its RNA is very small, has about 1,120 nucleotides [Bra91].

The protein shell of spherical viruses have icosahedral symmetry. An icosahedron, see Figure 2 (a), has 12 vertices and an equal number of 5-fold axes, 20 faces and the same number of 3-fold axes; it has 30 edges and the same number of 2-fold axes. An object placed on an icosahedron is repeated 60 times by combining the symmetry operations of the 5-fold, 3-fold and 2-fold axes. Indeed:

$12 \text{ vertices} \times 5 = 60$ – copies of the object relative to the 5-fold axes;

$20 \text{ faces} \times 3 = 60$ – copies of the object relative to the 3-fold axes;

$30 \text{ edges} \times 2 = 60$ – copies of the object relative to the 2-fold axes.

The capsid consists of $T \times 60$ sub-units where T can be 1 or more. Caspar and Klug have shown that the *triangulations* number $T = h^2 + hk + k^2$, with h and k integers, [Bra91]. Possible values for T are 1, 3, 4, 7, 9, 12, 13, ...

For $h = k = 1$, $T = 3$ and the protein shell consists of $3 \times 60 = 180$ sub-units. Each asymmetric unit consists of three proteins A, B, and C, see Figure 2 (b). In this case the RNA contains the genetic blueprint of the three proteins, A, B, and C.




Figure 2. The Protein shell of spherical viruses have icosahedral symmetry. (a) An icosahedron has twelve, (12), 5-fold axes of symmetry one through each vertex, twenty, (20), 3-fold axes, each going through the center of each face, and thirty, (30), 2-fold axes, each going through the half point of an edge. (b) For $T = 3$ each asymmetric unit (shown with tick lines) consists of three proteins A,B, and C

3. 3D Atomic Virus Structure Determination in Electron Microscopy

The reconstruction of an n -dimensional object from a set of $(n - 1)$ -dimensional projections is a well researched mathematical problem with numerous applications. For example, in medical imaging, a 3D representations of the human body can be obtained from a set of 2D images collected at known relative orientations of the X-rays source, through a procedure called Computer-Aided Tomography, or CAT-scan.

The same basic idea is used to reconstruct the 3D electron density map from 2D micrographs. The major difference is that in Cryo-TEM the actual reconstruction uses m different copies of the object, assumed to be identical with each other, but each of them in a different orientation. The sample observed with Cryo Transmission Electron Microscope, Cryo-TEM, consists of a frozen solution with m identical copies of the virus, each frozen in a orientation

that is not known.

We are only concerned with the process of refining the electron density map; we assume that we do have a low resolution, possibly inaccurate electron density map and our goal is to obtain a high resolution electron density map.

We have a 3D lattice with a size determined by the dimensions of the virus and with a variable number of grid points, depending upon the resolution. For example assuming a spherical virus with a radius of 500 Å, the number of grid points is 100^3 at 10 Å resolution and $1,000^3$ at 1 Å resolution. To determine the electron density map means to determine the value of a real valued function at all grid points.

The procedure for the atomic structure determination consists of the following steps:

[Step A] Extract individual particle projections from micrographs and identify the center of each projection.

[Step B] Determine the orientation of each projection.

[Step C] Carry out the 3D reconstruction of the electron density of the virus.

[Step D] Dock an atomic model into the 3D density map.

Steps B and C are executed iteratively until the 3D electron density map cannot be further improved at a given resolution. The number of iterations for 3D reconstruction is in the 10–20 range and one step of 3D reconstruction for a medium size virus may take several hours on a sequential computer and minutes on a parallel system. It typically takes weeks or even months to obtain an electron density map using sequential programs for the orientation determination and for the 3D reconstruction.

Algorithms for Step A, which include automatic identification of particle projections, the determination of the center and orientation of each virus particle projection are discussed elsewhere [Mar97]. The basic algorithm for the reconstruction of asymmetric objects in Cartesian coordinates is presented in [Lyn99]. An improvement of the algorithm that reduces the number of arithmetic operations by two orders of magnitude is outlined in [Lyn00], and a detailed presentation and analysis of the algorithms together with preliminary experimental results are given in [Mar01].

4. Parallel Orientation Refinement

Given: (1) a set of m views and (2) the electron density map, the goal is to find out the orientation of each view. Figure 3 (a) shows an icosahedral virus and defines the three angles determining the orientation of each particle, θ, ϕ, ω .

Several methods including the method of "common lines", [Bak99] can be used to this end. Here we describe a procedure for the refinement of orientations that is less sensitive to the noise caused by experimental errors. The basic idea is to project the electron density at known angles

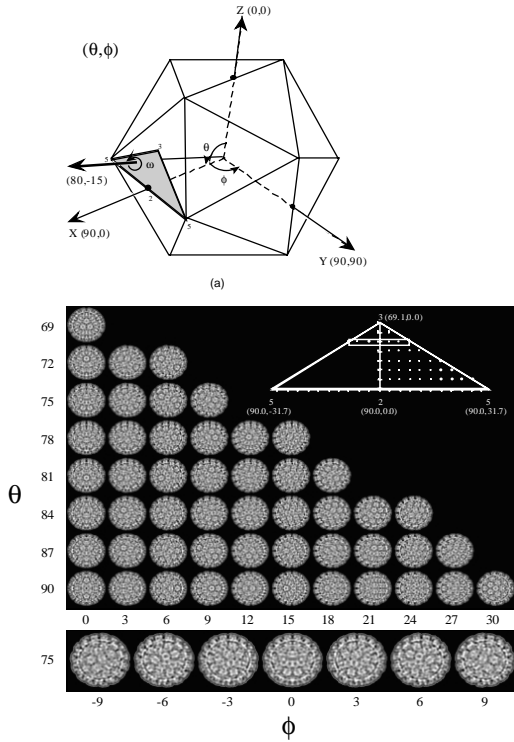


Figure 3. (a) The three angles used to characterize the orientation of a view. (b) The set of calculated views for an icosahedral virus at 3 deg. resolution.

and then to compare each experimental view with the calculated projection. Once we define the distance between an experimental view and a calculated projection, the goal of the search is to identify the calculated projection at the minimum distance from the experimental view. The procedure does not make any assumptions about the symmetry, but is capable to recognize the symmetry if one exists.

We describe now an algorithm where the search is conducted in the Fourier domain. We first perform a Discrete Fourier Transform, DFT, of the experimental view and then compare it with a cut at a with a precise orientation through the 3D DFT of the electron density map. In this algorithm the distance between two $l \times l$ arrays of complex numbers:

$$A^{exp} = [a_{j,k}^{exp} + i \times b_{j,k}^{exp}]_{1 \leq j,k \leq l} \quad \text{and}$$

$$A^{cal} = [a_{j,k}^{cal} + i \times b_{j,k}^{cal}]_{1 \leq j,k \leq l}$$

is computed as:

$$d(A^{exp}, A^{cal}) = \frac{1}{l^2} \sum_{j=1}^l \sum_{k=1}^l \sqrt{([a_{j,k}^{exp} - a_{j,k}^{cal}]^2 + [b_{j,k}^{exp} - b_{j,k}^{cal}]^2)}.$$

To compute the distance between the DFTs of two views requires

$$\mathcal{O}(l^2)$$

arithmetic operations. Given a resolution r , the computation to determine the distance, $d(A^{exp}, A^{cal})$ will use only the Fourier coefficients up to r . In this case the number of operations can be reduced accordingly.

Step a. Construct the 3D Discrete Fourier Transform, DFT, of the electron density map. Assuming that the size of electron density map is l^3 , this step requires

$$\mathcal{O}(l^3 \times \log_2(l^3))$$

arithmetic operations.

Step b. Interpolate in the 3D Fourier domain and construct a set of 2D planes, $\mathcal{C} = \{C_{\theta,\phi,\omega}\}$ for different θ, ϕ, ω angles. The set \mathcal{C} spans the search domain,

$$0 \leq \theta_{min} \leq \theta \leq \theta_{max} \leq \pi,$$

$$0 \leq \phi_{min} \leq \phi \leq \phi_{max} \leq \pi, \text{ and}$$

$$0 \leq \omega_{min} \leq \omega \leq \omega_{max} \leq \pi.$$

At angular resolution α the cardinality of the set \mathcal{P} is:

$$|\mathcal{P}| = \frac{\theta_{max} - \theta_{min}}{\alpha} \times \frac{\phi_{max} - \phi_{min}}{\alpha} \times \frac{\omega_{max} - \omega_{min}}{\alpha}.$$

This step requires

$$\mathcal{O}(l^2 \times |\mathcal{P}|)$$

arithmetic operations. For example, if $\alpha = 0.1$ deg. and the search range is from 0 to π for all three angles, then the size of the search space is very large: $|\mathcal{P}| = (1800)^3 = 5.832 \times 10^9$.

Figure 3 (b) shows that the corresponding size for an icosahedral particle at 3 deg. resolution consists of only 51 calculated views; at 0.1 deg. the size of the search space is less than 4,000 calculated views, [Bak99]. Thus for an asymmetric particle the size of the search space increases by six (6) orders of magnitude compared with an icosahedral particle !! Moreover, when comparing two views of an icosahedral virus particle, a calculated and an experimental one, we could use only a shell of thickness corresponding to the capsid, rather than the entire 2D image.

Step c. Given a set of m experimental views, $\mathcal{E} = \{E_1, E_2, \dots, E_i, \dots, E_m\}$ construct the set of 2D DFTs:

$$\mathcal{F} = \{F_1, F_2, \dots, F_i, \dots, F_m\}.$$

This step requires

$$\mathcal{O}(m \times l^2 \times \log_2(l^2))$$

arithmetic operations.

Step d. Compare the DFT of each experimental view $F_i \in \mathcal{F}$ with every $C_j \in \mathcal{C}$ and compute $d_{i,j} = d(F_i, C_j)$ the corresponding distance function.

This step requires

$$\mathcal{O}(|\mathcal{P}| \times m \times l^2)$$

arithmetic operations. If we are only using the DFT coefficients up to resolution r , this becomes:

$$\mathcal{O}(|\mathcal{P}| \times m \times r^2).$$

Step e. For every experimental view $E_i \in \mathcal{E}$ identify the minimum distance

$$d_i^{min} = \min\{d_{i,1}, \dots, d_{i,j} \dots d_{i,|\mathcal{P}|}\}.$$

Assign to view E_i the orientation of the plane C_{min} .

This step requires

$$\mathcal{O}(|\mathcal{P}| \times m)$$

arithmetic operations.

The practical implementation of the algorithm is considerably more intricate than the brief description presented above could suggest. The algorithm is embarrassingly parallel, each experimental view can be processed independently by a different processor. But the 3D electron density map and its DFT can be very large; the database of calculated views could require several TB for storage. We have implemented a shared virtual memory where 3D "bricks" of electron density or its DFR are brought on demand in each node when they are needed, a strategy presented in [Cor93].

A more sophisticated and efficient version of the algorithms should be based upon a two-step approach. First, perform a low resolution search to determine the approximate orientation and then carry out the high angular resolution search.

5. Experimental Results

An obvious question raised in connection with the previous algorithm is if the distance between a view whose orientation we want to determine and a set of calculated views with known orientations, each indexed by the three orientation angles, has a monotonic variation and exhibits a clear minimum. If this distance represented as a function of the orientation of the calculated view oscillates, rather than have a monotonic variation, then we cannot use this algorithm to identify the best match between a view with unknown orientation and calculated views with known orientations.

Another question is related to the maximum angular resolution of the search process. To distinguish between the detailed features of a virus structure we need to know the

orientation of a view with a precision ranging from 1 deg to 0.1 deg, depending upon the resolution of the reconstruction process.

To answer these questions we carried several tests. We took the 3D electron map ED_{map} of the *sindbis* virus at 15.4 Å resolution and projected it at several orientations to obtain distinct reference views, V_i^{exp} , with the corresponding orientations $(\theta_i, \phi_i, \omega_i)$.

Then we performed the 3D DFT of the electron density map and obtained $DFT(ED_{map})$, a 3D lattice of complex numbers. We constructed a set of views by calculating several cuts into this lattice for a range of angles around the known orientations of the reference views.

In Figure 4 we illustrate the results for the reference view with orientation $(\theta = 10, \phi = 20, \omega = 30)$, when ϕ is constant, $\phi = 20$ deg and the other two angles, θ and ω cover a range of 60 values. First, in Figures 4 (a) we show the case when the angular resolution is 3 deg and θ and ω take values in the range: $-26 \leq \theta \leq 154$ and $-6 \leq \omega \leq 174$. We see some local minima, but the global minimum is well defined and located at $(\theta = 10, \omega = 30)$. This minimum is obvious at 1 deg angular resolution, in Figure 4 (b). Figure 4 (c) shows the results of the search with the angular resolution set to 0.1 deg and again the minimum is well defined.

In Figure 5 we illustrate the results for the same reference view with orientation $(\theta = 10, \phi = 20, \omega = 30)$, when ω is constant, $\omega = 30$ deg and the other two angles, θ and ϕ cover a range of 60 values. First, in Figures 5 (a) we show the case when the angular resolution is 3 deg and θ and ϕ take values in the range: $-26 \leq \theta \leq 154$, and $-16 \leq \phi \leq 164$. Again, we see some local minima, but the global minimum is well defined and located at $(\theta = 10, \phi = 20)$. This minimum is obvious at 1 deg angular resolution, in Figure 5 (b). Figure 5 (c) shows the results of the search with the angular resolution set to 0.1 deg and again the minimum is well defined.

In Figures 6 (e) and (f) we show the results at 1 and 0.1 deg angular resolution for the $(\theta = 0, \phi = 0, \omega = 0)$ special case. The distance function is symmetric in ϕ and ω due to conventions defining the three angles.

In Figure 6 we first illustrate the results for the same reference view with orientation $(\theta = 10, \phi = 20, \omega = 30)$, when θ is constant, $\theta = 10$ deg and the other two angles, ϕ and ω cover a range of 60 values. First, in Figures 6 (a) we show the case when the angular resolution is 3 deg and ϕ and ω take values in the range: $-16 \leq \phi \leq 164$ and $-6 \leq \omega \leq 174$. Once again, we see some local minima, but the global minimum is very well defined. This minimum is obvious at 1 deg angular resolution, in Figure 6 (b). Figure 6 (c) shows the results of the search with the angular resolution set to 0.1 deg and again the minimum is well defined. Figure 6 (d) shows a close up for the 0.1 deg resolution.









Figure 4. Plot of the distance between the DFT of a reference view and calculated cuts in the 3D DFT of the electron density map. The reference view has orientation ($\theta = 10, \phi = 20, \omega = 30$). Search when $\phi = 20$ deg and θ and ω cover a range of values at angular resolutions from 3 deg to 0.1 deg. (a) Angular resolution 3 deg, $-26 \leq \theta \leq 154$ and $-6 \leq \omega \leq 174$. (b) Angular resolution 1 deg, $-2 \leq \theta \leq 58$ and $18 \leq \omega \leq 78$. (c) Angular resolution 0.1 deg, $8.8 \leq \theta \leq 14.8$ and $28.8 \leq \omega \leq 34.8$.

Figure 5. Plot of the distance between the DFT of a reference view and calculated cuts in the 3D DFT of the electron density map. The reference view has orientation ($\theta = 10, \phi = 20, \omega = 30$). Search when $\omega = 30$ deg and θ and ϕ cover a range of values at angular resolutions from 3 deg to 0.1 deg. (a) Angular resolution 3 deg, $-26 \leq \theta \leq 154$ and $-16 \leq \phi \leq 154$. (b) Angular resolution 1 deg, $-2 \leq \theta \leq 58$ and $8 \leq \phi \leq 68$. (c) Angular resolution 0.1 deg, $8.8 \leq \theta \leq 14.8$ and $18.8 \leq \phi \leq 24.8$.

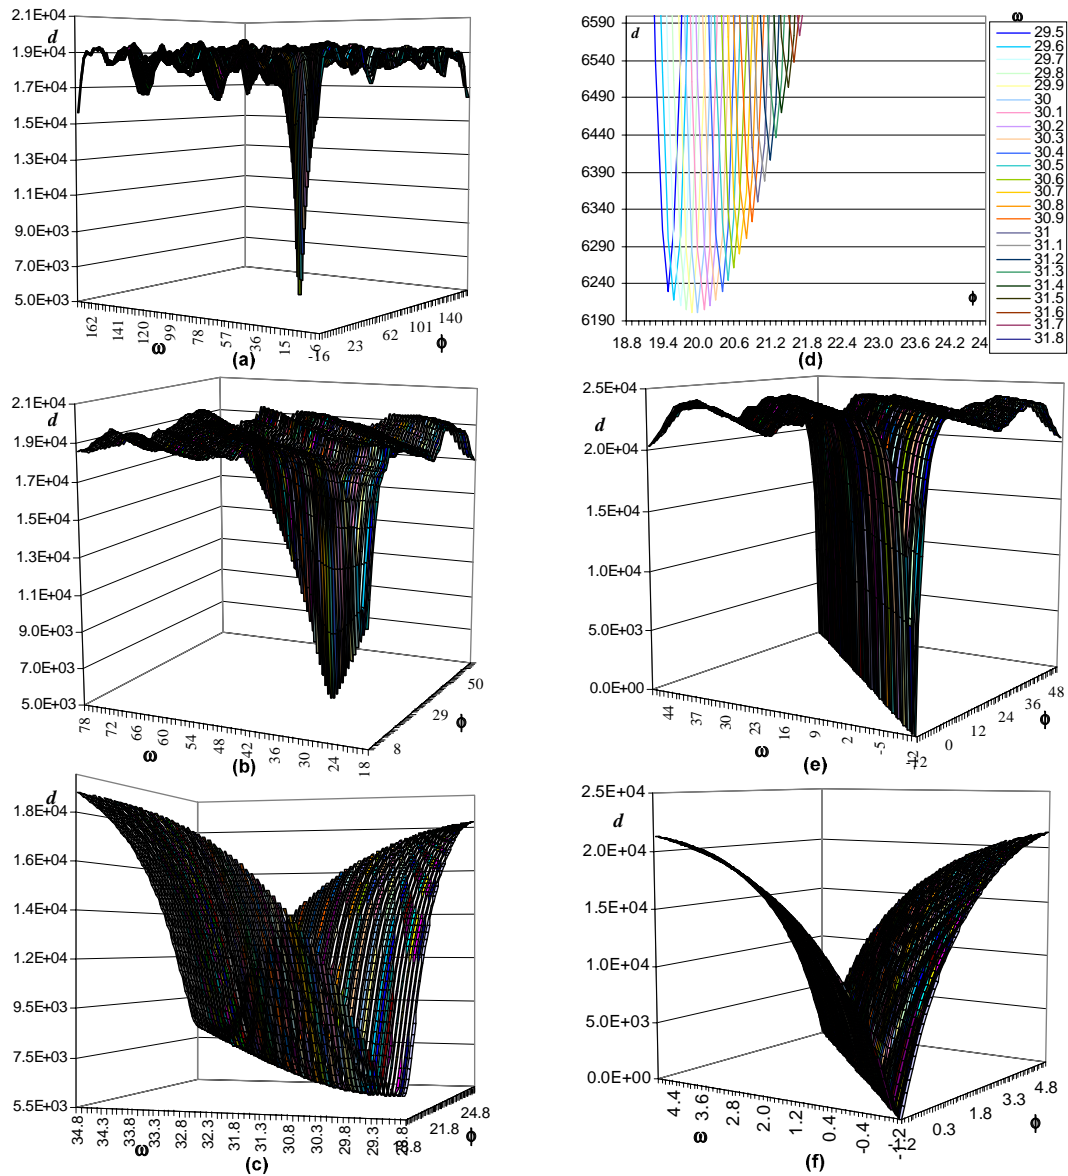


Figure 6. Plot of the distance between the DFT of a reference view and calculated cuts in the 3D DFT of the electron density map. The reference view has the orientation ($\theta = 10, \phi = 20, \omega = 30$) in (a)-(d). Search when $\theta = 10$ deg and ϕ and ω cover a range of values at angular resolutions from 3 deg to 0.1 deg. (a) Angular resolution 3 deg, $-16 \leq \phi \leq 154$ and $-6 \leq \omega \leq 174$. (b) Angular resolution 1 deg, $8 \leq \phi \leq 68$ and $18 \leq \omega \leq 78$. (c) Angular resolution 0.1 deg, $18.8 \leq \phi \leq 24.8$. $28.8 \leq \omega \leq 34.8$. (d) Angular resolution 0.1 deg, $18.8 \leq \phi \leq 24.8$. $28.8 \leq \omega \leq 34.8$, detail. The reference view has the orientation ($\theta = 0, \phi = 0, \omega = 0$) in (e)-(f). Search when $\theta = 0$. (e) Angular resolution 1 deg, $-12 \leq \phi, \omega \leq 18$. (f) Angular resolution 0.11 deg, $-1.2 \leq \phi, \omega \leq 4.8$.

One way to improve the results of the algorithm, suggested by Tim Baker is to give more weight to higher frequency components at higher angular resolution.

6. Conclusions

In this paper we discuss a novel algorithm for the refinement of orientations of individual virus particle projections for asymmetric particles and analyze its computational complexity.

The symmetry of the protein shell means that we have a relatively small number of identical building blocks that recognize each other and are able to assemble together spontaneously. The principle of genetic economy requires that the shell be built out of a few copies of identical units; the amount of genetic information. Thus the size of genome is considerably smaller for a symmetric virus particle.

This symmetry has a profound effect upon the complexity of computations required for structure determination from experimental evidence. For an asymmetric particle the size of the search space for orientation refinement increases by six (6) orders of magnitude compared with an icosahedral particle.

7. Acknowledgments

The research reported in this paper was partially supported by the National Science Foundation grants MCB 9527131 and DBI 9986316, by the Scalable I/O Initiative, and by a grant from the Intel Corporation. The authors are indebted to Prof. Timothy Baker and to Dr. Wei Zhang from the Biological Sciences Department of Purdue University for sharing with us the intricacies of electron microscopy, supplying the data for testing our algorithms, and for their insights into various aspects of orientation refinement.

References

- [Bak99] T. S. Baker, "Adding a third dimension to virus life cycles: three-dimensional reconstruction of icosahedral viruses from cryo-electron micrographs". *Microbiology and Molecular Biology Reviews*, December, 862–922, 1999.
- [Bol00] L. Bölöni, D. C. Marinescu, J. R. Rice, P. Tsompanopoulou, and E. A. Vavalis, "Agent-based scientific simulation and modeling", *Concurrency Practice and Experience*, vol 12, 845–861, 2000.
- [Böt97] B. Böttcher, S. A. Wynne, and R. A. Crowther, "Determination of the fold of the core protein of hepatitis B virus by electron cryomicroscopy", *Nature (London)* 386, 88–91, 1997.
- [Bra91] C. Branden and J. Tooze, "Introduction to protein structure", Garland Publishing, New York and London, 1991.
- [Con97] J. F. Conway, N. Cheng, A. Zlomick, P. T. Wingfield, S. J. Stahl, and A. C. Steven, "Visualization of a 4-helix bundle in the hepatitis B virus capsid by cryo-electron microscopy", *Nature (London)* 386, 91–94, 1997.
- [Cor93] M. A. Cornea-Hasegan, D. C. Marinescu, Z. Zhang, J. R. Rice, R. E. Lynch, and M. G. Rossmann, "Macromolecular electron density averaging on distributed memory MIMD systems", *Concurrency: Practice and Experience*, Vol. 5, No. 8, pp. 635–657, 1993.
- [Cor95] M. A. Cornea-Hasegan, D. C. Marinescu, Z. Zhang, A. Hadfield, J. Muckelbauer, R. E. Lynch, and M. G. Rossmann, "Phase refinement and extension by means of non-crystallographic symmetry averaging using parallel computers", *Acta Crystallographica*, Vol. D51, pp. 749–759, 1995.
- [Fra96] J. Frank, "Three-dimensional electron microscopy of macromolecular assemblies", Academic Press, 1996.
- [Lyn99] R. E. Lynch, D. C. Marinescu, H. Lin, and T. S. Baker, "Parallel algorithms for 3D reconstruction of asymmetric objects from electron micrographs", *Proc. IPPS/SPDP*, IEEE Press, 632–637, 1999.
- [Lyn00] R. E. Lynch, H. Lin, and D. C. Marinescu, "An efficient algorithm for parallel 3D reconstruction of asymmetric objects from electron micrographs," *Proc. Euro-Par 2000, Lecture Notes in Computer Science*, vol. 1900, 481–490, 2000.
- [Mar97] I. M. Martin, D. C. Marinescu, T. S. Baker, and R. E. Lynch, "Identification of spherical particles in digitized images of entire micrographs", *J. of Structural Biology*, 120, 146–157, 1997.
- [Mar98] I. M. Martin and D. C. Marinescu "Concurrent computations and data visualization for spherical virus determination", *IEEE Computational Science & Engineering*, October-December, pp. 40–51, 1998.
- [Mar01] D. C. Marinescu, Y. Ji, and R. E. Lynch "Space-time tradeoffs for parallel 3D reconstruction algorithms for virus structure determination", *Concurrency and Computation: Practice and Experience*, vol. 13, pp. 1083–1106, 2001.