

Algorithmic Issues in Modeling Motion

PANKAJ K. AGARWAL

Duke University

LEONIDAS J. GUIBAS

Stanford University

HERBERT EDELSBRUNNER

Duke University

JEFF ERICKSON

University of Illinois, Urbana-Champaign

MICHAEL ISARD

Compaq Research Labs

SARIEL HAR-PELED

University of Illinois, Urbana-Champaign

JOHN HERSHBERGER

Mentor Graphics

CHRISTIAN JENSEN

University of Aalborg

LYDIA KAVRAKI

Rice University

PATRICE KOEHL

Stanford University

MING LIN AND DINESH MANOCHA

University of North Carolina, Chapel Hill

DIMITRIS METAXAS

University of Pennsylvania

BRIAN MIRTICH

MERL

DAVID MOUNT

University of Maryland

This article is based on the report of the *Workshop on Algorithmic Issues in Modeling Motion*, sponsored by an NSF grant CCR-00-83-033 and an Army Research Office grant DAAD 19-00-1-0478, held on August 6 and 7, 2000 at Duke University, Durham, NC.

Authors' addresses: Pankaj K. Agarwal, Duke University, pankaj@cs.duke.edu; Leonidas J. Guibas, Stanford University, guibas@cs.stanford.edu; Herbert Edelsbrunner, Duke University, edels@cs.duke.edu; Jeff Erickson, University of Illinois, Urbana-Champaign, jeffe@cs.uiuc.edu; Michael Isard, Compaq Research Labs, misard@robots.ox.ac.uk; Sariel Har-Peled, University of Illinois, Urbana-Champaign, sariel@cs.uiuc.edu; John Hershberger, Mentor Graphics john_hershberger@mentorg.com; Christian Jensen, University of Aalborg, csj@borg.cs.auc.dk; Lydia Kavradi, Rice University, kavraki@cs.rice.edu; Patrice Koehl, Stanford University, koehl@csb.stanford.edu; Ming Lin, University of North Carolina, Chapel Hill, ming@cs.unc.edu; Dinesh Manocha, University of North Carolina, Chapel Hill, dm@cs.unc.edu; Dimitris Metaxas, University of Pennsylvania, dnm@central.cis.upenn.edu; Brian Mirtich, MERL, mirtich@merl.com; David Mount, University of Maryland, mount@cs.umd.edu; S. Muthukrishnan, AT&T, muthu@research.att.com; Dinesh Pai, University of British Columbia, pai@cs.ubc.ca; Elisha Sacks, Purdue University, eps@cs.purdue.edu; Jack Snoeyink, University of North Carolina, Chapel Hill, snoeyink@cs.unc.edu; Subhash Suri, University of California, Santa Barbara, suri@cs.ucsb.edu; Ori Wolfson, University of Chicago, wolfson@eecs.uic.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

©2002 ACM 0360-0300/02/1200-0550 \$5.00

S. MUTHUKRISHNAN

AT&T

DINESH PAI

University of British Columbia

ELISHA SACKS

Purdue University

JACK SNOEYINK

University of North Carolina, Chapel Hill

SUBHASH SURI

University of California, Santa Barbara

AND

OURI WOLEFSON

University of Chicago

This article is a survey of research areas in which motion plays a pivotal role. The aim of the article is to review current approaches to modeling motion together with related data structures and algorithms, and to summarize the challenges that lie ahead in producing a more unified theory of motion representation that would be useful across several disciplines.

Categories and Subject Descriptors: F.2.2 **Analysis of Algorithms and Problem Complexity**: Nonnumerical Algorithms and Problems—*Computations on discrete structures, Geometrical problems and computations*; I.6.5 **Simulation and Modeling**: Model Development—*Modeling methodologies*

General Terms: Algorithms

Additional Key Words and Phrases: Computational geometry, computer vision, mobile networks, modeling, molecular biology, motion modeling, physical simulation, robotics, spatio-temporal databases

1. INTRODUCTION

Motion is ubiquitous in the physical world. Computational models of physical objects

and processes, from protein folding in biology to assembly planning in manufacturing, from connectivity in mobile communications networks to location-based

services in spatial databases, and from ecological models in environment science to galactic evolution in astrophysics must deal with representing mobile data and their evolution over time. It is instructive to draw a comparison with another modality that permeates modeling the physical world as much as motion, namely shape. Over the years computational shape representations are intensely studied and a variety of approaches have been developed for modeling the different shapes that arise in applications. Entire disciplines, such as CAGD (computer-aided geometric design), evolved, in this case devoted to the study of the smooth shapes needed by the automotive and aerospace industries. What is striking when we compare motion to shape, however, is how rich and complex the space of representations for motion can be. For one, motion is often described intensionally—by specifying what the motion needs to accomplish, and not what the motion is. For another, motion happens over time, and in the physical world unpredictable events can occur that change the evolution of a system; this on-line character of motion must be modeled.

Both military and civilian applications (digitized battlefields, automatic target recognition, mobile communications, virtual environments, animation, physics-based simulation, animation of deformable objects, to name a few) have raised a wide range of algorithmic issues in modeling motion and in coping with continuously moving objects. Although a vast amount of work has been done on modeling, analyzing, searching, processing, and visualizing geometric objects, most of the work has focused by and large on handling static objects and on dynamic data structures, which can update the information as objects are inserted or deleted at discrete times. These approaches are, however, not suitable for handling moving or deforming objects because either the algorithm has to be executed continuously, or the solution returned will be at times obsolete. The on-line nature of motion makes it difficult to consider time simply as another dimension and regard the moving object as a static object in space-time. Algorithms

that can handle continuous motion, what we call *kinetic algorithms*, are needed in such applications.

Despite a flurry of activity in several areas in the last few years on modeling motion and on developing efficient algorithms that deal with moving objects, most of the work to date has been scattered among a number of fields, and a unified algorithmic theory of motion that permeates across multiple disciplines and applications is missing. In Sections 2–9, we survey the state of the art in motion modeling across many areas—computational geometry, mesh generation, physical simulation, biology, computer graphics, computer vision, robotics, spatio-temporal databases, and mobile wireless networks—and suggest that a synergy among these areas is needed, since many issues are common to all of them; furthermore, many real-world applications require the integration of techniques from a wide range of disciplines. We conclude in Section 10 by highlighting a number of challenges that lie ahead in creating a firm algorithmic theory for the processes of acquiring, modeling, reasoning about, planning, manipulating, and executing motion.

2. COMPUTATIONAL GEOMETRY

Earlier work on moving objects in computational geometry, initiated by a paper by Atallah [1985], focused on bounding the number of combinatorial changes in geometric structures as the objects move along prescribed trajectories [Agarwal and Sharir 2000]. In the late 1980s, algorithms were developed for computing the changes in geometric structures as objects undergo motion [Kahan 1991]. However, these results assume an off-line model of motion in which the trajectories are given in advance, a relatively uncommon situation. In practice, researchers were using fixed-rate sampling methods to handle moving objects, in which a structure is updated at fixed time intervals and therefore the fastest moving object gates the time step for the entire system. The kinetic data structure (KDS) framework introduced by

Guibas et al. [Basch et al. 1997] alleviated many of the problems with the off-line and fixed-rate-sampling methods. The main idea in the kinetic framework is that even though objects move continuously, the relevant *combinatorial* structure changes only at certain discrete times and therefore one does not have to update the data structure continuously. The validity of the combinatorial structure is certified by a number of elementary assertions about the state of the system called *certificates*. *Kinetic updates* are performed on the data structure only when certain *kinetic events* occur, corresponding to certificate failures, which can be monitored at different rates.

Kinetic data structures have led to efficient algorithms and data structures for a wide range of problems. However, in order for KDS and other motion models to be made more realistic, several issues need to be addressed.

(1) *Motion Sensitivity.* The motions of objects in the world are often highly correlated, therefore, it is desirable to find representations and data structures that exploit such motion coherence. It is also important to find mathematical measures that capture how coherent motions are and then use this coherence as a parameter to quantify the performance of motion algorithms. If we do not do this, our algorithm design may be aimed at unrealistic worst-case behavior, without capturing solutions that exploit the special structure of the motion data that actually arise in practice. Most kinetic algorithms to date have been analyzed under the model in which each object follows an independent smooth trajectory defined by a small set of parameters, and the algorithm's efficiency is measured by the worst-case behavior over all such possible motions. A challenging issue is the development of a class of kinetic *motion-sensitive* algorithms, whose performance can be expressed as a function of how coherent the motions of the underlying objects are.

(2) *Trade-Offs between Complexity Measures.* The performance of a kinetic algorithm can be measured by the number of events it processes, the time spent in up-

dating the structure at each event, and the size of the structure. Although optimal or near-optimal algorithms under these criteria have been proposed for a number of kinetic problems, including convex hull of a point set in the plane and closest pair, no optimal algorithms are known for a wide range of other problems, for example, convex hulls in \mathbb{R}^3 , triangulation of points in \mathbb{R}^2 , range searching, etc. In such cases, one would like to obtain a trade-off between different complexity measures. A natural trade-off would be between the compactness and the efficiency of a KDS. This is similar to the trade-off between the size and the query time of a static data structure—a query can be answered more efficiently by allowing more space [Agarwal and Erickson 1999; Matoušek 1994]. A similar issue is how to improve the efficiency of a kinetic algorithm by maintaining a geometric structure implicitly. For example, we know of techniques for maintaining the triangulation of a point set in the plane implicitly by structures that process only quadratic number of events, but no such bounds are known if one is required to maintain a triangulation explicitly [Agarwal et al. 2000].

Another possible trade-off is between efficiency and accuracy. How much efficiency can be gained by maintaining a geometric structure *approximately*? For example, it is known that the diameter of a set S of points (the maximum distance between a pair of points in S) moving in the plane can change quadratically many times [Agarwal et al. 1997], but recently it was shown [Agarwal and Har-Peled 2001] that if we maintain the diameter approximately, the number of events only depends on the approximation factor (and is independent of the number of points). Although some preliminary work has been done on trading off efficiency with accuracy, a more general theory is needed for approximating motion.

(3) *Flexible Scheduling.* The current KDS implementations maintain a global priority queue that stores the future events at which the structure may need

to be updated. The events are processed in the sorted temporal order. The correctness of these methods relies on the assumption that all the events have been processed in this order. Under algebraic motion, an event is typically a root of a polynomial of fixed degree, and thus we have to compute the roots of a polynomial. If we use floating-point arithmetic or numerical methods to estimate the roots, we compute them only approximately and cannot ensure that the events have been processed in the sorted order. On the other hand, using exact arithmetic and algebraic methods (such as Sturm sequences) for computing or isolating roots is quite expensive. Another source of difficulty in processing events in sorted order is degeneracies in the input—a degeneracy causes multiple events to occur simultaneously. Little attention has been paid so far to these problems.

In many physical simulations, the motion law of the objects is specified by an ordinary or partial differential equation. In such contexts, polynomial trajectories can be at best short-term approximations of the actual trajectories, useful for conservative estimates of certificate failure times. The correctness of the approximation needs to be certified and tracked as well.

It is desirable to develop a framework where event-based scheduling can be intermixed with fixed-time-sampling. The latter can be used whenever explicit motion prediction, apart from full integration of the equations of motion, is difficult. The current KDS repair mechanism strongly depends on the assumption that it is invoked to repair *a single* certificate failure. In a time-step-based scenario, multiple certificates may fail; we need techniques that can repair a structure after multiple certificate failures.

(4) *Canonical vs. Noncanonical Structures.* The complexity measures mentioned earlier are more suitable for maintaining *canonical* geometric structures, those that are uniquely defined by the position of the data, for example, convex hull, closest pair, and Delaunay triangulation. In these cases, the notion of

external events (those that change the combinatorial structure of the attribute of interest) is well defined and is independent of the algorithm used to maintain the structure. On the other hand, suppose we want to maintain a triangulation of a moving point set. Since the triangulation of a point set is not unique, the external events depend on the triangulation being maintained, and thus depend on the algorithm. This makes it hard to analyze the efficiency of a kinetic triangulation algorithm. Most of the current approaches for maintaining noncanonical structures artificially impose canonicity and maintain the resulting canonical structure. But this typically increases the number of events. Another drawback of imposing a canonicity condition is that it might make the structure global, which in turn would make the task of designing a distributed algorithm considerably harder.

In general, current techniques are too weak to analyze the performance of non-canonical structures and better techniques are needed. In Agarwal et al. [2000], Kirkpatrick et al. [2000], and Kirkpatrick and Speckmann [2000, 2002], we made some progress in this direction, but these techniques are very problem specific. A better model for analyzing the lower bounds is needed. Some of the analysis techniques developed for analyzing the on-line algorithms could prove useful.

(5) *Decentralization.* When we want to simulate motion over a large scale, for example, the flow patterns of cars in a city, mobile agents in an ad-hoc network, or a chemical reaction among several macromolecules, the motion computation may have to be distributed over several processors/machines. Furthermore, if we are confirming certificates from sensor data, the sensors themselves may be geographically distributed over a large area and may be also constrained by communication bandwidth and power requirements. In such settings, the normal centralized way of implementing KDSs cannot be used, as the communication, processing, and power capabilities of the central node will severely limit the scalability of the system. Most of

the existing approaches assume the existence of a central event queue that keeps track of all the changes. In a distributed environment, reasoning about the state of the system and tracking the attributes of interest must take place in a distributed manner. Each processing node will be responsible for only some of the mobile data and will communicate information with other nodes as necessary. How to develop such algorithms even for simple problems remains a challenging issue, despite a few recent attempts in this direction [Gao et al. 2001].

(6) *The Number of Combinatorial Changes.* Although optimal or near-optimal bounds are known on the number of changes in the basic geometric structures such as convex hull, closest pair, diameter, and width of a point set, such bounds have remained elusive on many other structures, including Delaunay triangulation, minimum spanning tree, and alpha shapes. Moreover, most of the work has focused on bounding the number of changes in the worst case, which rarely occurs in practice. There has been some work on bounding the number of changes when the trajectory of each point is randomly chosen from a reasonable distribution. It would be useful to study the number of changes under more realistic assumptions on motion, for example, when the motion of the points is correlated, or to develop a model in which the bounds are proved as a function of the complexity of the motion of objects.

3. MESH GENERATION

Many physical and engineering problems are modeled by defining a set of partial differential equations over a continuous domain. However, it is usually infeasible to perform computer simulations over a continuous space. Therefore, the domain is discretized by decomposing its interior into a mesh of simple and well-shaped elements such as boxes or simplices, the differential equations are approximated over the discretized domain using finite difference, element, or volume techniques, and

then solved numerically. The error of a solution is estimated, and if necessary, the mesh is refined and the numerical simulation repeated. Not all meshes are equally good; the error introduced by discretization depends on the size and shape of the elements of the mesh, and the computational complexity of the simulation depends on the number of elements in the mesh and their overall quality. This has led to vast literature on mesh generation; see Bern and Plassmann [2000] for a recent survey.

In many applications, ranging from weather modeling to seismic analysis, from protein folding to heart modeling, the domain evolves with time, which raises the problem of mesh generation over time-varying domains. Depending on the application, either the mesh is updated dynamically over time or time is considered as another dimension, and the mesh is generated over the space-time domain.

3.1. Dynamic Meshes

To maintain the mesh, we must trace changes of the shape and repair the mesh at places where it becomes inadequate. We see at least three types of changes:

- (i) Changes of *shape*, which can usually be accommodated by moving mesh vertices (and thus implicitly cells).
- (ii) Changes of *curvature*, which cause distortions of the homeomorphism between the mesh and the underlying domain. In the case of triangulations, we can repair such distortions by locally inserting or deleting simplices.
- (iii) Changes of *topology*, which require a local reorganization of the mesh connectivity.

In all cases, we require good mechanisms to predict or detect when these changes happen, as we already discussed for kinetic certificates in Section 2.

Algorithms that maintain meshes under all three types of changes are rare [Cheng et al. 2001], and they are sensitive to how quickly changes can happen. This gives rise to a hierarchy of motions or deformations, ordered by

difficulty in maintaining a mesh algorithmically. A deformation may be best understood by considering the trajectory of the domain through a higher-dimensional space-time. For example, a deforming 2-dimensional surface in \mathbb{R}^3 sweeps out a 3-dimensional manifold in \mathbb{R}^4 . The speed of the deformation is then captured by the slope of the manifold in the direction of time.

3.2. Meshing Space-Time Domains

A possible alternative to dynamic meshing is meshing directly in space-time, which trades the dynamic nature of the problem for an extra dimension. *Space-time meshing* algorithms are becoming more desirable now as numerical algorithms for space-time problems are being developed.

Space-time discontinuous Galerkin methods are a relatively new family of techniques for solving systems of nonlinear hyperbolic partial differential equations [Cockburn et al. 2000; Lowrie et al. 1998; Richter 1994; Yin. et al. 2000]. These methods rely on the *domain of influence* for dynamic data. The domain of influence for a point p in the space-time domain is a cone with apex p ; the cone specifies points in space-time whose relevant physical parameters (temperature, pressure, etc.) depend on the corresponding parameters of p . A finite-element mesh satisfies the so-called *cone constraint* if the domain of influence of any point on any interior mesh face is disjoint from that face; intuitively, this implies that information flows in only one direction between any two adjacent elements. If the cone constraint is satisfied, space-time discontinuous Galerkin methods can compute a numerical solution one element at a time, avoiding the need to solve a large system of equations.

Most space-time meshing algorithms divide time into fixed-length intervals and construct a mesh within each layer [Lowrie et al. 1998; Üngör et al. 2002]. The length of the global time step is imposed by the smallest element in the underlying space mesh; this increases

both the size of the mesh and the numerical error in the computed solution. The only exception to date is the “Tent Pitcher” algorithm of Üngör and Sheffer [2002] and its recent extension by Erickson et al. [2002]. These algorithms construct space-time meshes using an advancing front method. At each step, one node on the front is advanced through time to create a new space-time element. The time step taken by each node depends on the size and quality of the nearby spatial elements. This advancing front strategy is necessary for problems in which the parameters of the cone constraint change over time and must be computed as part of the numerical solution. Despite this recent progress, the problem of meshing over general space-time domains remains largely open.

- (i) Unlike fixed time-step methods, advancing front methods require a constant underlying space mesh; the fixed discretization of time has been traded for a fixed discretization of space. Methods to coarsen, refine, and otherwise locally remesh the advancing front would be useful not only to avoid numerical error but also to make the meshing process itself more efficient.
- (ii) In many problems, the boundary of the simulation domain changes continuously over time. Tracking moving boundaries requires new techniques to move nodes of a space mesh continuously through space-time.
- (iii) In problems involving material flow, cone constraints can be skewed or anisotropic, but current techniques assume uniform cone constraints. This brings in the issue of hierarchically modeling increasingly complex motion, as in Section 2.
- (iv) Finally, further work is required to determine useful quality measures for space-time meshes and to construct space-time meshes with guaranteed size and quality bounds. This is in stark contrast to the traditional setting, where appropriate measures

of quality and algorithms to compute high-quality meshes are well known [Bern and Plassmann 2000].

4. PHYSICAL SIMULATION

The applications of physical simulation can be divided into two broad categories. In one category are applications requiring physical simulation to reason about the real world for engineering, verification, and prediction. Here, accuracy is key, and many of the challenges involve the development of accurate models. Many of these, such as finite-element or boundary-element models, require the use of meshing techniques as discussed in Section 3. The second category of applications require physical simulation for producing “plausible” motion; these applications include animation, games, content creation, and some kinds of education and training. Here, since the constraints on accuracy are more lax, modeling is less of a problem, and efforts have focused on other issues such as unstructured interaction and real-time speeds. Below are four challenges related to motion in physical simulation; the first three are relevant to both categories of applications, and the final one is relevant to the nonengineering applications.

One grand challenge of physical simulation is the combination of different state representations to solve complex, heterogeneous problems. Consider the different ways motion is represented through velocities. The velocities of all points on a rigid body are simultaneously described by the linear velocity of a single point on the body and an angular velocity. On the other hand, the velocity of a deformable body is often represented by the linear velocity of a number of node points attached to the body. Finally, the velocity of fluid might be represented by linear velocities at lattice points fixed in space. How are these representations to be harmonized when studying, for example, the turbulence induced by a rigid body moving through a fluid, or the wrinkling of the clothes on an articulated human figure? Some progress has been made [Baraff

and Witkin 1998; van den Doel et al. 2001; Meyer et al. 2001], but many problems remain unresolved in modeling interactions and constraints between objects with different underlying state and motion representations.

A second challenge is to design simulation systems that can effectively choose appropriate motion models and algorithms for the given situation. Perhaps a rigid body model is a reasonable approximation, or maybe a linear or even nonlinear deformable model is needed. Inertial effects may or may not be negligible. There are a variety of ways to model friction with different properties and computational costs. Even after the model has been chosen, there might be alternative algorithms. For example, there are at least three families of methods for modeling rigid body contact (analytic/LCP formulations [Lotstedt 1981; Baraff 1989], penalty methods which are popular in haptics [Zilles and Salisbury 1995], and impulse-based methods [Mirtich 1996]), each with different strengths and weaknesses. By and large, today’s simulation systems use a single model and a single algorithm for a given class of dynamic objects. A better approach would be to choose among the models and algorithms, depending on dynamic objects. To do this, a system must know the important results of the simulation, the allowable tolerances, the sources of error in each algorithm, the bounds on these errors, and the error propagation over time.

A third challenge is to parallelize simulation algorithms so that they can be scaled to bigger problems by adding more computational units. This is already possible for certain types of simulations involving finite element methods and computational fluid dynamics. It is harder for more unstructured systems such as those involving rigid bodies. Often, such systems can not be statically partitioned into independent subsystems, and furthermore the parts of the system that affect another part vary over time. A partitioning among CPUs needs to be a dynamic one. Better bounds on the motion of objects over time could make maintaining this

partition easier. Also to be considered are load balancing issues: equally sized partitions are most useful.

A final challenge in physical simulation is to give better control to the user. Particularly when used for nonengineering applications, simulation is praised for producing complex motion while cursed for being an unwieldy, unpredictable tool [Laszlo et al. 2000]. Despite living in a physical world, humans are poor at designing simulations to give them “what they want”; in many cases, the desired outcomes are not even physically possible. Despite a lot of work in this area, users still need better handles and hooks into the physical simulation process in order to steer its course. Often the user’s desires are expressible in terms of motion constraints. Representations for these constraints are needed that are not only intuitive to humans but that also lead to tractable computer solutions. Also needed are algorithms that can let go of the true physics just enough to meet the constraints.

5. BIOLOGY

Molecules move and deform. Chemical processes essential to life critically depend on the ability of biomolecules to adopt different shapes over time. If we could realistically describe molecular motions, we would greatly improve our understanding of these temporal processes, including protein folding and molecular interactions (such as the assembly of protein complexes, protein-ligand docking, etc.), representing some of the most important unsolved problems in biology. Simulating the motions of molecules poses its own challenges, beyond those covered in earlier sections.

5.1. Molecular Dynamics Simulations

At the atomic level, the physical principles underlying molecular motions are fairly well understood: ligands and proteins have primarily torsional degrees of freedom, and the forces on an individual atom can be expressed as sums of well-

known potentials [Leach 1996]. Yet molecular dynamics simulations of a molecular system are still far too expensive computationally to allow us to follow a process over its complete time period. For example, a protein folds within a few seconds, while the longest published molecular dynamics simulation of a protein covers one microsecond. The reason that accurate molecular dynamics simulations remain a major challenge in computational biology is that they involve thousands of degrees of freedom in the molecule of interest, require many evaluations of potential energy functions that can be expensive to compute, and in addition need to take in account the solvent environment. In order to make progress, we must develop novel ways to represent deforming shapes and track their evolution over time, find techniques for describing the physics in terms of higher-level units than individual atoms, and use efficient approximations whenever justified to do so. More specifically, we need the following.

(1) *Implicit Solvent Potential Energy Functions.* Molecular dynamics simulations that include a large number of water molecules around the solute remain the state of the art in this field. They are, however, inefficient, since a large fraction of the computing time is spent calculating a detailed trajectory of the solvent molecules, even though it is primarily the solute behavior that is of interest. Several semi-analytical implicit treatments of solvent have been proposed. They all rely on an energy term that is related to the solvent exposed surface area of each atom of the solute. Inclusion of such terms in a molecular dynamics simulation requires the calculation of accurate surface areas, as well as their analytical derivatives with respect to atomic position. Fast analytical methods are needed for these calculations [Bryant et al. 2002].

(2) *Efficient Updating Procedures.* Molecular dynamics simulations are based on solving the Newton equations of motion for all atoms of the system considered. These equations are solved numerically, using very short time steps

of the order of a fraction of a femtosecond. Each step requires the evaluation of the total energy function of the system, as well as its derivatives with respect to atomic position. For standard energy terms, such as Lennard–Jones and electrostatics interactions, only a fraction of all pairs of atoms need be considered. A crucial aspect for fast molecular dynamics simulation is to have access to fast algorithms for building and updating these interacting pairs of atoms.

(3) *Hierarchical Representation of Proteins.* Very long simulation of molecular dynamics require approximations. One approach is to simplify the representation of the molecule of interest: simplified models of proteins have proved both popular and effective until the present time. A significant weakness of such models, however, is their inability to deal with topological features of the molecule in a meaningful way. We need new methods that build simplified models with geometrical persistence [Edelsbrunner et al. 2000].

5.2. Motion-Planning and Protein Folding

Molecular dynamics simulations follow molecular motions by solving a deterministic or stochastic system of equations. They provide very detailed analysis of short and usually small amplitude motions, with reasonable agreement with experimental data coming from structural biology. While it is expected that improvements of the computer technology and in the algorithms applied for solving these equations will increase the time span covered by these simulations, time scales on the order of milliseconds and seconds are still out of reach. As we already remarked, unfortunately these are the time scales of most interest for many biological processes. Different approaches specific to these problems need to be explored. See Dill et al. [1995], Levitt et al. [1997], and Onuchic et al. [1997] for a review of known techniques on protein folding.

There are large and ongoing research efforts whose goal is to determine the native folds of proteins (the protein structure prediction problem). Most approaches

proceed in two steps. First, a large collection of possible conformations (or decoys) for the protein of interest is built. Second, these conformations are screened using potential energy functions and the best models define the predicted structures for the protein. In these procedures, the focus is set on thermodynamics (definition of low energy states) and not on the kinetics of the protein folding process (i.e., how the protein folds to its native conformation). A better understanding of the latter should not only facilitate protein structure prediction but also provide insights on how proteins function. Motion planning and more specifically probabilistic roadmap methods provide one approach to solving this problem [Amato et al. 2002; Song and Amato 2002; Apaydin et al. 2002]. Application of such techniques to the protein folding problem requires the following:

(1) *Falsifiability Studies.* Motion planning usually starts by sampling the moving object's configuration space (C -space), and retaining those conformations that satisfy certain feasibility requirements. Roadmap nodes are then connected by finding the path of minimum energy between the starting point and the goal. Application of such a strategy to the protein folding problem is not intuitive and should be extensively tested. Interestingly, recent experimental results suggest that protein folding mechanisms and landscapes are largely determined by the topology of the native state and are relatively insensitive to details of the interatomic interactions [Alm and Baker 1999]. This dependence on low-resolution structural features suggest that it should be possible to describe the physics of the folding process using relatively low resolution models. It remains that initial studies should focus on proteins for which kinetic data are available, as well as on proteins whose dynamics have been described.

(2) *Algorithmic Developments.* In the case of protein folding, the C -space is large, since the protein has a very large number of degrees of freedom. Fast algorithms for extensive sampling of the C -space of proteins must be developed.

The protein folds based on geometric constraints (i.e., no self-collision) but also energetic constraints. Realistic energy functions are therefore needed, such as those used for molecular dynamics simulations—a point already mentioned. Algorithms are also needed to select roadmap nodes from possible configurations of the protein in its C -space. Such selection can be based on energetic and/or geometric criteria [Apaydin et al. 2001, 2002]. While several algorithms have been developed for generating the path of minimal energy among road nodes in robotics, they have to be adapted to the specific case of protein folding, in order to cope with a large number of nodes.

6. COMPUTER VISION

6.1. Shape and Motion

Vision processes data that represents motion in the physical world. Sensor noise, calibration, and uncertainty are novel issues that have to be addressed. Motion analysis in computer vision can be divided into scene-based and object-based approaches.

Scene-based analysis based on generic, model-free methods has been studied since the 1980s and has led to successful systems for camera pose estimation and image stabilization. This method usually exploits either rigid-body scene assumptions or optic flow to estimate gross motion characteristics. A common approach in this area is to extract structure from motion, assuming a single moving camera in a static scene [Longuet-Higgins 1981; Hartley and Zisserman 2000]. While challenges remain in these areas, commercial codes are starting to become available to solve such problems (e.g., <http://www.2d3.com/>).

Most of the deep motion-related research problems in vision arise when the scene must be interpreted in terms of objects undergoing independent motion. Consequently, doing motion analysis successfully for computer vision often relies on having sufficiently expressive shape (or, more generally, appearance) models to

describe the objects being viewed. Certain classes of shapes have been successfully modeled. These include rigid, near planar shapes, and those whose contours or intensities can be modeled using linear combinations of principal components [Black and Yacoob 1997]. There have also been successes with 3D jointed structures as well as 2.5D “cardboard cutout” representations [Black and Jepson 1998; Torr et al. 2001]. However, there remain large classes of objects that we do not know how to model well, including pragmatically important ones such as people wearing loose clothing. Much of the difficulty in object tracking comes as much from the challenge of devising a suitable appearance model as from the motion analysis itself.

In the past, most of the shape models used have been hand-crafted, parameterized models [DeCarlo and Metaxas 2000] and the methods developed for motion analysis have been deterministic in nature (e.g., physics-based methods). Recently, statistical methods have been widely used to overcome the obvious limitations of hand-crafted models, to generalize the representation of motion in terms of parameters that follow distributions, and to improve the coupling between motion estimation and motion recognition [Isard and Blake 1998; Blake et al. 1999]. Such methods promise new capabilities for segmentation, parameterized shape representation, motion estimation and motion recognition.

Within this new class of statistical methods, the following open research problems need to be addressed.

6.2. Segmentation, Grouping and Initialization

While there has been a significant improvement in recent years in our ability to track objects with complex appearance and motion models, the problem of initialization—the matching of an object model to this first frame of a sequence as a precursor to tracking—has lagged behind. Initialization is typically treated as a separate algorithm distinct from the subsequent tracking, and often relies

on generic segmentation techniques. Furthermore, initialization is often attempted using a single image whereas it might be much easier to do given several frames; it is clear that we need a better understanding of how to merge the problems of initialization and tracking to arrive at a continuum between them. Initialization is also clearly related to grouping methods; for example, initializing a jointed-limb model can be viewed as the task of grouping the individually moving limbs into a coherent whole. No good algorithms are known for incorporating this grouping into the initialization/tracking process, perhaps in order to select the form of a shape model appropriate to an unknown object. Finally, the problem of coupling shape models with statistical methods in an effort to automate the estimation of the priors for improved segmentation and model initialization has only recently begun to be studied [Chen and Metaxas 2000].

6.3. Statistics and Learning

As in other fields, such as speech recognition, the computer-vision community is increasingly realizing the importance of learning-based statistical models for developing robust methods. Although there are learning-based methods for estimating low-level motion parameters, basic research needs to be done on coupling the low-level motion parameters estimated by the above methods with the semantic interpretation of these parameters. An example of such needed integration is the problem of Gesture and Sign Language Recognition [Vogler and Metaxas 2000].

While motion analysis has so far been studied under the assumption of fairly well lit objects that follow the optical flow assumption, very little research has been done on motion analysis under varying lighting conditions. This is clearly a very important problem for robust motion tracking.

6.4. Multiple Scales

Objects in the world are often naturally represented at multiple scales. This is true of object shape, which may lend it-

self to hierarchical description; object texture: for example, a piece of paper may be well represented as having uniform texture from a distance and as containing discrete letters close up; and object motion, which is typically well approximated by a simple linear process at short timescales, while these linear motions compose over longer timescales to form complex behaviors which may be best described using grammars over discrete spaces. A few algorithms for dealing with multiple scales for shape exist [Lindeberg 1994], but no good overall theory for shape or motion has emerged. In particular, we do not yet fully understand the conditional dependencies, which relate the different scales. The study of principled statistical algorithms to deal correctly with inference of motion over multiple scales is also in its infancy.

6.5. Summary

In some ways, motion analysis is one of the success stories of recent computer vision research. Around ten years ago, memory sizes and computing speeds became adequate to deal with lengthy image sequences, and since then tracking has made great strides: simple “blob” trackers that follow the gross motion of nonoverlapping objects against relatively static backgrounds are now a standard component of robust systems, which can run unattended for months. As soon as the complexity of the problem is increased in any dimension, however, current solutions break down, so there is plenty of research left to do, from following the motion of complex nonrigid objects to interpreting and summarizing extended behaviors, to coupling vision-based motion trackers with other modalities, such as speech, for improved activity recognitions.

7. ROBOTICS

As in the previous section, robotics must deal with motion in the “real” world—and not just with idealized models of such motion in a virtual world. As a consequence, besides the usual criteria such as efficiency and simplicity, the performance of motion algorithms in robotics

must be judged based on how accurately they model the real world (for a class of tasks) and how robust they are to the uncertainties.

7.1. Modeling Motion

Some of the key difficulties of modeling motion in robotics include model selection and parameter estimation. Despite extensive work in mechanics over the last few centuries, there are numerous everyday objects that robots need to reason about and manipulate, which are not easy to model using known laws. The fundamental problem lies at the fact that the parameters needed for the models are usually unknown and difficult to acquire. For instance, our ability to model the motion of an object rolling on the ground is limited by the difficulty of knowing the surface friction at the interface and even the inertial properties [Kry and Pai 2002]. Or, consider the problem of modeling a deformable object, say a human organ like the liver, in surgical robotics [Debunne et al. 2001; Brown et al. 2001]. How is the deformation of the liver to be modeled? In order to tackle these problems, one has to address the following issues:

- (i) Algorithms for estimating parameters of motion models from observations [Lang et al. 2002; Richmond and Pai 2000]. These can often lead to ill-posed inverse problems, requiring sophisticated numerical techniques.
- (ii) Techniques for representing, propagating, and reasoning about uncertainty in the motion. Kalman and particle filters are two examples of representing and propagating the uncertainties, but we also need algorithms for incorporating these uncertainties in collision detection and impact problems [Sorenson 1985; Arulampalam et al. 2002].
- (iii) New empirical models. For many complex problems, empirical models are the only option to model motion. For example, in order to model the deformation of nonuniform elastic objects, rather than reconstructing the

internal structure of the object so that known models of elasticity can be applied, it may be preferable to construct an empirical model of the deformation behavior of the object. Essentially, all friction and impact laws used in practice are empirical, and they are applicable only for a small class of tasks [Mason 2001].

Many of the difficulties in modeling motion are because of the the *constraints* that the motion must satisfy. The physical laws governing the motion are obvious constraints. These laws are typically continuous and are modeled using differential equations [Mason 2001]. However, the motion algorithms discretize them. Despite the recent progress in discretization of differential equations, many issues remain unsolved, including the ability to extend them to deformable shapes. Whether the kinetic data structures, discussed in Section 2, can be used to discretize differential equations remains an interesting open question. Contact constraints are another type of constraints that are hard to model [Pai et al. 2001]. For instance, simulating the roll-slide motion during smooth contact is a challenge to both collision-detection algorithms (especially those relying on polyhedral approximations and bounding volumes) and to traditional simulation algorithms (especially those relying on constraint stabilization). How to represent local deformations at contact is another challenging problem—localized (“point”) contacts produce very large stresses, which locally deform apparently rigid objects and can produce important elastic effects such as stick-slip oscillations.

7.2. Collision Detection

A central problem in robotics is detecting collision between many moving objects. It has been extensively studied in both robotics and computational geometry [Lin and Gottschalk 1998]. The work in computational geometry has focused on theoretical approaches, which are, at times, complex to implement, while the

work in robotics has lacked a general approach that works well in all cases. Commonly used in practice are methods that track the closest pair of features between two moving objects and update the closest pair by a local computation after each time step [Mirtich 1998; Lin and Gottschalk 1998]; an underlying assumption is that, if the time step is chosen sufficiently small, then the closest pair can be updated incrementally and efficiently. Feature trackers work well for simple convex objects, or combinations of a small number of such. Hierarchical representations of objects have been incorporated to this approach to handle more complex objects (see, e.g., [Larson et al. 2000] and references therein), and kinetic data structures (Section 2) have been used to alleviate some of the problems of the fixed-time-sampling approach. Alternatively, hierarchies of simple bounding volumes (spheres or boxes) have been used to enclose shapes; these hierarchies are refined only to the coarsest level necessary to establish that the two objects do not intersect.

The above approaches, however, are not suitable for detecting collision between many moving objects, especially if each of them is moving with different speed, because one has to perform collision testing for every pair of objects. Instead a global approach is needed. Recently, a few global approaches are proposed for detecting collision between many moving polygons in the plane, which maintain a tiling of the common exterior of the polygons into flexible cells, so that the polygons are known to be disjoint until one of the cells self-collides [Agarwal et al. 2000; Kirkpatrick et al. 2000; Kirkpatrick and Speckmann 2002]. They maintain additional information that makes it easy to predict when a cell self-collides, and to update the tiling when that occurs. Although this approach looks promising in the plane, there are many stumbling blocks in extending it to 3-space. Even less is known on detecting collision between deformable objects or for detecting self-intersection in a deformable object [Bridson et al. 2002; Guibas et al. 2002].

A different problem is how to respond when a collision occurs [Mirtich 2000]. For example, in virtual-reality applications such as computer games, digital battlefields, haptics interfaces, etc., it is critical that reaction to collision detection is close to reality—Should an object penetrate the other? Should one of the objects bounce back? Should the objects deform or break? This raises several interesting questions that combine geometry with physical models of the objects involved.

7.3. Planning Motion

A typical motion-planning problem asks for computing a collision-free motion between two given placements of a given robot in a given environment. In its simplest form, one assumes that the environment is fully known and that there are no constraints on the motion of the robot except that it cannot collide an obstacle. The problem is typically solved in the *configuration space*, in which each placement (or configuration) of the robot is mapped as a point; see Latombe [1991]. The *free configuration space* F is the subset of the configuration space at which the robot does not intersect any obstacle. The robot can move from an initial configuration to a final configuration without intersecting an obstacle if these two configurations lie in the same connected component of F . Planning a collision-free path thus reduces to connectivity and other topological questions in F . Numerous general techniques and techniques for low-dimensional configuration spaces have been developed [Barraquand and Latombe 1991; Halperin et al. 1997]. This has also led to several interesting topological questions on configuration spaces, which are typically represented as semi-algebraic sets (finite Boolean combinations of solution sets of polynomial equalities and inequalities).

The dimension of configuration space depends on the degrees of freedom of the robot, and it can be quite high. Computing high-dimensional configuration spaces exactly is impractical [Barraquand and Latombe 1991]. Therefore, techniques to

compute an approximate representation of F are needed. Much of the difficulty in approximating F is in understanding the topology and in simplifying the topology of F . Recently, Monte Carlo algorithms have been developed for representing F by 1-dimensional networks, called *probabilistic road maps* (PRMs) [Kavraki et al. 1995, 1996]. Intuitively, this network is an approximate representation of the *road map*, a 1-dimensional network that captures the connectivity information of F . These methods sample points in F and connect them by an edge if they can be connected by a simple path inside F . Despite several heuristics to sample well the points in F , better sampling strategies are needed to handle narrow corridors and other difficult areas in F , so that the connectivity of F is preserved. Variants of such methods have already been described in Section 5. PRMs can quickly find a simple collision-free path between initial and final configurations in many situations. But there is no absolute guarantee that they will do so—in fact, very few general purpose techniques are known for establishing that paths do not exist.

In the above discussion, we assumed that there were no constraints on the motion of the robots, which in practice is not the case. Various nonholonomic constraints such as bound on the maximum velocity or acceleration may exist on the moving robot. Planning a collision-free in presence of these constraints is considerable harder, and relatively little is known except various ad-hoc approaches. Better techniques that are practical as well as theoretically sound are needed for nonholonomic motion planning.

In some applications, even more challenging problems arise. If the obstacles are also moving, the free configuration space has to be updated dynamically. Also, flexible objects such as elastic bands, rope, or cloth cannot be properly represented with a finite number of degrees of freedom [Baraff and Witkin 1998; Lamiraux and Kavraki 2001; Meyer et al. 2001]. How one can represent configuration spaces of such objects is a big challenge.

8. SPATIO-TEMPORAL DATABASES

Efficient indexing schemes that support various queries are central to any large database system. Most existing database systems assume that the data is constant unless it is explicitly modified. For example, if the value of the age field is set to 40 years, then this age is assumed to hold (i.e., 40 years is returned in response to a query) until explicitly updated. Such systems are not suitable for representing, storing, and querying continuously moving or varying objects; either the database has to be continuously updated or a query output will be obsolete. Motivated by applications in digital battlefields, air-traffic control, and mobile communication systems, methods for indexing dynamic attributes are needed, which guarantee efficient access to moving objects. A basic scenario is one in which the database system receives trajectories for a large population of objects capable of continuous movement; new objects may arrive, existing objects may leave, and that the positions of existing objects may be updated dynamically. The objects, which can assume to have rigid shapes, may be represented as points, for example, people with mobile phones, people in Internet-worked cars, airplanes, etc. The movement of the objects may be constrained by stationary or moving objects, the movement of the objects may be confined to networks (e.g., cars confined to a road network), and the motion of objects may be highly correlated. The goal is to answer various queries based on the locations, trajectories, and topology, and to explore patterns in motion of objects. Since the data is quite large and resides on disk, one cannot scan the whole data to answer a query.

Instead of continuously updating the position of a moving object, a better approach is to represent the position as a function $f(t)$ of time, so that changes in object position do not require any explicit change in the database system. With this representation, the database needs to be updated only when the function $f(t)$ changes. Recently, there has been some work on extending the capabilities of

existing database systems to handle moving-object databases (MOD); see, for example, Sistla and Wolfson [1995], Sistla et al. [1997], and Chomicki and Revesz [1999].

Methods such as KDS, developed for simulating some structures/phenomena over time, are not always suitable for answering queries on the database of moving objects. The queries might relate either to the current configuration of objects or to a configuration in the future—in the latter case, we are asking to predict the behavior based on the current information. In either case, tracking the current configuration of the objects would be unnecessary and expensive. Another approach to handle motion is to regard the time as another spatial axis and reduce the problem to a static problem in one higher dimension. For example, a set of points moving in the plane can be regarded as lines in \mathbb{R}^3 . This approach has the advantage that the data structure need not be updated unless the trajectory of an object changes. However, it suffers from a number of problems. First, the approach expects the trajectory of the points to be known in advance. Second, it reduces the problem to one higher dimension, which makes the algorithms slower.

Here are a few different types of queries that one may want to answer on moving objects.

(1) *Location-Based Queries.* These queries include range and proximity queries [Šaltenis et al. 2000; Wolfson et al. 1999]. For example, given a rectangle R and a time t , report all objects that will be inside R at time t ; given a rectangle R and a time interval $[t_1, t_2]$, report all objects that will pass through R in the time interval $[t_1, t_2]$; given a point p and a time interval $[t_1, t_2]$, report the objects that would be nearest to p in that time interval. Despite a flurry of papers on indexing moving objects for answering location-based queries, the problem remains largely open. The theoretical results, based on partition trees and KDSs, are too complex to be practical. The practical indexing schemes

are based on quad-trees, R -trees, and their variants. They do not always work well, and their performance deteriorates with time and the data size. In order to improve the performance of an indexing scheme, some of the researchers have introduced the notion of *horizon* that gives the duration for which the current index performs well. They suggest to rebuild the index periodically or to adjust it when the trajectories are updated (see, e.g., Tayeb et al. [1998]). Although these approaches provide better query performance, it is not always practical to rebuild the index periodically; see Agarwal and Procopiuc [2002] for a recent review of the known indexing techniques.

(2) *Continuous Queries.* In the above queries, the query range was spatially fixed. One could also ask queries when the location of the query range is also moving [Sistla et al. 1997]. For example, keep track of all cars within 5 miles. Of course, such a query can be answered by asking location-based queries repeatedly, but it would be very expensive. One would like to combine KDS with procedures for location-based approaches to answer such queries. The known lower bounds on range-searching data structures imply that it is hard to answer such queries efficiently in the worst case. But can one exploit the constraints on motion to answer them efficiently?

(3) *Trajectory-Based Queries.* These queries involve the topology of trajectories and derived information such as velocity of the objects [Pfoser et al. 2000]. These queries are deemed critical, but also rather expensive. Here is a typical query: Which two objects will be within one mile during the next ten minutes? Which objects are heading *East*, or which objects are moving at speed greater than 65 miles/hour? An even harder query is: report all objects whose speed doubles in the next ten minutes.

(4) *Dynamization.* It is an inherent characteristic that new objects arrive and that the positions of existing objects are updated dynamically. Hence, the index

should support fast insert, delete operations. Although basic spatial indexing schemes, such as *R*-trees or *kd*-trees, support insertion/deletion of objects, the problem is considerably harder for indices handling moving objects because they have to store additional information to predict the motion of objects. This information is expensive to update as objects are inserted or deleted.

The current indexing schemes basically delete and reinsert a point whenever the trajectory of an object changes. In many applications, the trajectory of an object is known in advance. Can this information be used to avoid explicit deletion and reinsertion of the object?

(5) *Uncertainty*. Since in many applications the position of objects is being received through a sensor such as GPS, uncertainty in position and velocity of the objects is inevitable. Indexing schemes that can handle uncertainty are needed. A simple model of uncertainty is to allow an interval of values for each of the parameters defining the dynamic attributes of an object. In this model, the uncertainty in position corresponds to a “ball” in which the object can lie, and the uncertainty in velocity corresponds to a “cone” in which the object lies [Agarwal et al. 2000; Wolfson et al. 1998, 1999]. These simple models have been incorporated in a few indexing schemes, but a more systematic study is needed. More sophisticated parametric models have been developed for modeling uncertainty. Although these methods have been successfully used for tracking moving objects, they have not been used in the context of databases. How to incorporate them in indexing schemes is a challenging open problem.

9. MOBILE WIRELESS NETWORKS

A world of “anytime, anywhere, anyhow” computing is being created by a unique combination of two powerful technological trends: rapid component miniaturization and the emergence of high speed wireless communication. The recent boom in the use of cell phones is merely the be-

ginning of a much larger trend towards broadband wireless networks that carry not only voice, but all kinds of multimedia data. Wireless data networking presents a multitude of challenges that are distinct from their fixed infrastructure networking counterpart:

- (1) [*Heterogeneity*]. There are a variety of wireless technologies (e.g., modulation schemes and spectrum bands), diverse devices (e.g., laptops, PDAs, data-enabled cell phones and pagers), communication protocols (e.g., several 3G standards for wireless and SMS). Finally, the channel conditions under which wireless communication occurs vary widely.
- (2) [*Resource Constraints*]. The nature of wireless technology and devices presents stringent resource constraints, such as the available spectrum (number of channels, codes, slots, etc.) and the power (total and received power), among others.
- (3) [*Mobility*]. The most distinguishing aspect of a wireless network is the mobility of its clients. Dealing with mobility presents new problems (handoffs, paging clients, etc.) as well as it complicates existing ones (such as power allocation schemes where channel conditions vary with mobility patterns). In fact, any structuring mechanism provided by the network that exploits the local topology of the nodes has to be adapted to motion.

9.1. Mobility and Routing

A *mobile ad hoc network* (MANET) is an autonomous system of mobile routers and hosts connected by wireless links. The network is self-organizing and self-configuring, with nodes establishing the necessary routes among themselves without the help of a fixed infrastructure. Since the communicating hosts may be some distance apart, multiple network hops through intermediate nodes may be required to complete the route. The hosts themselves act like “mobile routers,” and cooperatively forward messages not

addressed to them. As these hosts move around, the routing protocol of the network must adapt its routing decisions to maintain communication. The rate of topology change may be quite dramatic in some ad hoc networks.

Since the topology of the network is constantly changing, the routing protocols must cope with frequent link failure and disconnections. The distribution of up-to-date information can easily overload the network, while out-of-date information can drive a network into instability [Royer and Toh 1999; Broch et al. 1998]. Thus, the shortest path routing used commonly in fixed infrastructure networks may be the wrong routing algorithm for ad hoc networks [Karp and Kung 2000]. Instead, discovering and routing on “long lived” routes may be more desirable.

In large-scale ad hoc networks, it may be infeasible, and even undesirable, for each host to maintain up-to-date information about the locations of all hosts and the topology of the entire network. Perhaps a more useful approach will be a statistical modeling of the movement of individual hosts. A grand challenge in the field, which may be currently just academic, but which will have tremendous impact on systems and performance, is the following: *Can we model the mobility patterns of users at varying levels of granularity?* The question should be explored in the context of modeling arrival patterns of customers to service queues, web server request sizes, etc. It could be refined in many ways—for example, does an individual user’s mobility pattern fit a small number of short random walks on the plane? Can the number of roaming users in any cell at any moment be predicted without understanding the mobility pattern of individual users? This grand challenge may involve input from geometers on finding a small set of attributes of the users’ trajectories on the globe that are most predictive of the users’ movements. One of the aspects of this grand challenge that is interesting is its “internet” scale of users and resource usages.

9.2. Mobility and Network Layers

Another approach to study the impact of mobility on networking infrastructure is to study its effect on the various architectural layers of the network. We list below some fundamental issues for each layer that are worth exploring.

- (1) [*Physical Layer*]. How does one perform resource scheduling at the physical layer in order to optimize performance for a given mobility pattern. For instance, how does one optimize across issues of power consumption, error correction, coding, rate control, channel assignment, and so on.
- (2) [*MAC Layer*]. How does one discover resources in the presence of *ad-hoc* mobility, such as reachability testing, topology determination, and so on.
- (3) [*Transport Layer*]. How and to what extent must one implement end-to-end TCP properties in presence of host mobility? How does high vs. low mobility affect persistence of TCP connections and fragmentation? Packet loss probabilities are significantly higher in wireless networks than in the wired networks due to widely varying channel conditions that are compounded by users’ mobility. Specifically, the TCP reacts quite violently to packet losses, with its exponential backoff mechanism, because it interprets the loss of packet to mean congestion in the network. However, in the wireless networks, the packet losses are just as likely the result of transmission loss as congestion, so how should the TCP rate control mechanism be modified?
- (4) [*Network Layer*]. How should one balance the cost of “redirection” vs. “paging,” in networks for accessing a user by better understanding of mobility models?
- (5) [*Application Layer*]. How does one support location specific queries on databases of yellow pages or services? How does one support a geographically connected community of users or

“buddies”? These involve testing geographic neighborhoods based on distance constraints where distance updates are obtained by pinging.

Many additional issues are likely to arise as wireless data networking and web evolve to become more widespread. While many of the issues listed above are algorithmic, they are best explored by collaboration with electrical engineers and computer scientists who work on different layers of wireless networking in order to be effective. Such collaborations represent an opportunity for the theoretical computer science community to formulate and solve the fundamental problems of interest in mobility and related issues in wireless networking. Finally, there is clearly a large overlap between the problems discussed here and those in Section 8 on spatio-temporal queries for mobile data.

10. CHALLENGES FOR THE FUTURE

We conclude this article by reviewing a number of challenging issues that cut across many of the previously discussed areas. It is our hope that this article will motivate researchers in these diverse disciplines to seek common ground and enable the cross-fertilization of ideas in motion modeling.

10.1. Motion Representation and Analysis

This covers most of the topics we have already presented in earlier sections:

(1) *Uncertainty*. Whenever we try to model motion in the physical world, we must deal with uncertainty. So far only the computer vision and robotics communities have developed refined tools for the probabilistic treatment of object motion [Arulampalam et al. 2002]. Such ideas could also prove very useful in temporal database and networking domains.

(2) *Robustness*. Many geometric algorithms suffer from robustness problems because of mismatch between fast floating-point computer arithmetic, which is finite precision, and the semantics of

real analysis [Schirra 2000]. Algorithms fail when an approximate computation yields an incorrect or even inconsistent qualitative property (often a topological property). In the context of kinetic algorithms, algebraic representation of trajectories and uncertainty in the input makes the problem even more acute.

(3) *Aggregation*. Related to uncertainty is the issue of statistical models for representing aggregate motion, say the motion of cars on the freeway, or a flock of birds. We may be able to reason reliably about the high-level behavior of the system, even though we have large uncertainty about the motion of each individual object.

(4) *Hierarchical Representations of Complex Motion*. While hierarchical representation methods for shape are highly developed, much less is known for motion. A molecular dynamics simulation, for example, includes all the atomic vibrations caused by thermal noise, which can obscure the large-scale structure of the motion that we want to understand. How to represent the essential dynamics of molecules, which has a physical meaning, remains a challenging problem [Amadei et al. 1992; Teodoro et al. 2000].

(5) *Marrying the Continuous and the Discrete*. Most physical systems evolve following continuous evolution laws, yet their evolution is punctuated by discrete events, such as collisions, that can alter these laws. Historically, the communities that have studied the continuous and discrete aspects of this problem have been distinct (scientific computing vs. discrete algorithms). A much tighter integration between the two can lead to substantial progress.

(6) *Trade-Off between Realism and Efficiency*. Little has been done so far to formally explore trade-offs between accuracy and efficiency [Agarwal and Har-Peled 2001]. Such analyses can benefit all of the areas we have been discussing.

(7) *Decentralization*. Distributed motion algorithms are equally important for large-scale simulations and for low-power

mobile devices in situations such as ad hoc networking or sensor nets [Bonnet et al. 2000, 2001].

(8) *Querying Motion*. Many fundamental questions remain on how best to sense and organize mobile data when the goal is to answer certain queries about the mobile system (and there is no need to know the full state of the system in between).

10.2. Motion Integration

Though each of the areas we discussed has good reasons for the motion representations it has chosen, this diversity hinders the development of end-to-end systems where techniques from several disciplines need to process motion data in an integrated fashion. Currently, big gaps exist between the motion representations best suited for visual motion analysis and those for motion generation, for example. We feel it is important to define a number of fundamental motion representations that can be used across areas and that can form the basis for implementing integrated applications.

REFERENCES

- AGARWAL, P. K., ARGE, L., AND ERICKSON, J. 2000. Indexing moving points. In *Proceedings of the Annual ACM Symposium on the Principles of Database Systems*. 175–186.
- AGARWAL, P. K., BASCH, J., GUIBAS, L. J., HERSHBERGER, J., AND ZHANG, L. 2000. Deformable free space tiling for kinetic collision detection. In *Proc. 4th Workshop Algorithmic Found. Robot.*, to appear.
- AGARWAL, P. K. AND ERICKSON, J. 1999. Geometric range searching and its relatives. In *Advances in Discrete and Computational Geometry*, B. Chazelle, J. E. Goodman, and R. Pollack, eds. *Contemporary Mathematics*, vol. 223. American Mathematical Society, R.I., pp. 1–56.
- AGARWAL, P. K., GUIBAS, L. J., HERSHBERGER, J., AND VEACH, E. 1997. Maintaining the extent of a moving point set. In *Proceedings of the 5th Workshop on Algorithms Data Structure*. Lecture Notes in Computer Science, vol. 1272. Springer-Verlag, New York, pp. 31–44.
- AGARWAL, P. K. AND HAR-PELED, S. 2001. Maintaining approximate extent measures of moving points. In *Proceedings of the 12th ACM-SIAM Symposium on Discrete Algorithms*.
- AGARWAL, P. K. AND PROCOPIUC, C. M. 2002. Advances in indexing mobile objects. *IEEE Bull. Data Eng.* 25, 25–34.
- AGARWAL, P. K. AND SHARIR, M. 2000. Arrangements and their applications. In *Handbook of Computational Geometry*, J.-R. Sack and J. Urrutia, Eds. Elsevier Science Publishers B.V. North-Holland, Amsterdam, The Netherlands, pp. 49–119.
- ALM, E. AND BAKER, D. 1999. Prediction of protein-folding mechanisms from free-energy landscapes derived from native structures. *Proc. Nat. Acad. Sci. USA*, 96, 11305–11310.
- AMADEI, A., LINSSEN, B. M., AND BERENDSEN, H. J. C. 1993. Essential dynamics of protein. *Proteins: Structure, Function, and Genetics* 17, 412–425.
- AMATO, N. M., DILL, K. A., AND SONG, G. 2002. Using motion planning to map protein folding landscapes and analyze folding kinetics of known naive structures. In *Proceedings of the 6th International Conference in Computational Molecular Biology*. pp. 2–11.
- APAYDIN, M., BRUTLAG, D., GUESTRIN, C., HSU, D., AND LATOMBE, J. 2002. Stochastic roadmap simulation: An efficient representation and algorithm for analyzing molecular motion. In *Proceedings of the 6th International Conference in Computational Molecular Biology*. pp. 12–21.
- APAYDIN, M., SINGH, A., BRUTLAG, D., AND LATOMBE, J. 2001. Capturing molecular energy landscapes with probabilistic conformational roadmaps. In *Proceedings of the IEEE Conference on Robotics and Automation*.
- ARULAMPALAM, S., MASKELL, S., GORDON, N., AND CLAPP, T. 2002. A tutorial on particle filters for on-line non-linear/non-gaussian bayesian tracking. *IEEE Trans. Sig. Proc.* 50, 174–188.
- ATALLAH, M. J. 1985. Some dynamic computational geometry problems. *Comput. Math. Appl.* 11, 1171–1181.
- BARAFF, D. 1989. Analytical methods for dynamic simulation of non-penetrating rigid bodies. *Comput. Graph* 23, 223–232.
- BARAFF, D. AND WITKIN, A. P. 1998. Large steps in cloth simulation. In *Proceedings of SIGGRAPH*. pp. 43–54.
- BARRAQUAND, J. AND LATOMBE, J.-C. 1991. Robot motion planning: A distributed representation approach. *Internat. J. Robot. Res.* 10, 628–649.
- BASCH, J., GUIBAS, L. J., AND HERSHBERGER, J. 1997. Data structures for mobile data. In *Proceedings of the 8th ACM-SIAM Symposium Discrete Algorithms*. pp. 747–756.
- BERN, M. AND PLASSMANN, P. 2000. Mesh generation. In *Handbook of Computational Geometry*, J.-R. Sack and J. Urrutia, Eds. Elsevier Science Publishers, B. V. North-Holland, Amsterdam, The Netherlands, 291–332.
- BLACK, M. AND YACOOB, Y. 1997. Recognizing facial expressions in image sequences using local

- parameterized models of image motion. *Int. J. Comput. Vis.* 25, 23–48.
- BLACK, M. J. AND JEPSON, A. D. 1998. Eigentracking: Robust matching and tracking of articulated objects using a view-based representation. *Int. J. Comput. Vis.* 26, 63–84.
- BLAKE, A., NORTH, B., AND ISARD, M. 1999. Learning multi-class dynamics. In *Advances in Neural Information Processing Systems*. vol. 11. MIT Press, pp. 389–395.
- BONNET, P., GEHRKE, J. E., AND SESHADRI, P. 2000. Querying the physical world. *IEEE Pers. Commun.* 7, 10–15.
- BONNET, P., GEHRKE, J. E., AND SESHADRI, P. 2001. Towards sensor database systems. In *Proceedings of the 2nd International Conference on Mobile Data Management*.
- BRIDSON, R., FEDKIW, R., AND ANDERSON, J. 2002. Robust treatment of collisions, contact and friction for cloth animation. In *Proceedings of SIGGRAPH 2002*. ACM Press, pp. 594–603.
- BROCH, J., MALTZ, D. A., JOHNSON, D. B., HU, Y.-C., AND JETCHEVA, J. 1998. A performance comparison of multi-hop wireless ad hoc network routing protocols. In *ACM Mob. Comput. Netw. (MobiCom)*. 85–97.
- BROWN, J., SORKIN, S., BRUYNS, C., LATOMBE, J.-C., MONTGOMERY, K., AND STEPHANIDES, M. 2001. Real-time simulation of deformable objects: Tools and application. *Proc. Comput. Animat.*
- BRYANT, R., EDELSBRUNNER, H., KOEHL, P., AND LEVITT, M. 2000. Inclusion-exclusion formulas for the area derivative of a space-filling diagram. manuscript.
- CHEN, T. AND METAXAS, D. 2000. Image segmentation based on the integration of Markov random fields and deformable models. In *Proceedings of MICAI 2000*. pp. 11–14.
- CHENG, H.-L., DEY, T. K., EDELSBRUNNER, H., AND SULLIVAN, J. 2001. Dynamic skin triangulation. *Disc. Comput. Geom.* 25, 525–568.
- CHOMICIKI, J. AND REVESZ, P. Z. 1999. A geometric framework for specifying spatiotemporal objects. In *Proceedings of the 6th International Workshop on Time Representation and Reasoning*. pp. 41–46.
- COCKBURN, B., KARNIADAKIS, G., AND SHU, C. W. 2000. *Discontinuous Galerkin Methods: Theory, Computation and Applications*. Springer-Verlag, New York.
- DEBUNNE, G., DESBRUN, M., CANI, M.-P., AND BARR, A. H. 2001. Dynamic real-time deformations using space & time adaptive sampling. In *Proceedings of ACM SIGGRAPH*. pp. 31–36.
- DECARLO, D. AND METAXAS, D. 2000. Optical flow constraints on deformable models with applications to face tracking. *Int. J. Comput. Vis.* 32, 99–127.
- DILL, K. A., BROMBERG, S., YUE, K. Z., AND FIEBIG, K. M. 1995. Principles of protein folding. a perspective from simple exact models. *Protein Sci.* 4, 561–602.
- EDELSBRUNNER, H., LETSCHER, D., AND ZOMORODIAN, A. 2000. Topological persistence and simplification. In *Proceedings of the 41st IEEE Symposium on Foundations on Computer Science*, pp. 454–463.
- ERICKSON, J., GUOY, D., SULLIVAN, J. M., AND ÜNGÖR, A. 2002. Building space-time meshies over arbitrary spatial domains. In *Proceedings of the 11th International Meshing Roundtable*, 391–402.
- GAO, J., GUIBAS, L., HERSHBERGER, J., ZHANG, L., AND ZHU, A. 2001. Discrete mobile centers. Manuscript.
- GUIBAS, L., NGUYEN, A., RUSSEL, D., AND ZHANG, L. 2002. Collision detection for deforming necklaces. In *Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms*. pp. 33–42.
- HALPERIN, D., KAVRAKI, L. E., AND LATOMBE, J.-C. 1997. Robotics. In *Handbook of Discrete and Computational Geometry*, J. E. Goodman and J. O'Rourke, Eds. CRC Press LLC, Boca Raton, Fla, pp. 755–778.
- HARTLEY, R. I. AND ZISSERMAN, A. 2000. *Multiple View Geometry in Computer Vision*. Cambridge University Press.
- ISARD, M. AND BLAKE, A. 1998. Condensation—Conditional density propagation for visual tracking. *Int. J. Comput. Vis.* 29, 5–28.
- KAHAN, S. 1991. A model for data in motion. In *Proceedings of the 23th Annual ACM Symposium on the Theory of Computing*. pp. 267–277.
- KARP, B. AND KUNG, H. T. 2000. Greedy perimeter stateless routing (GPSR) for wireless networks. *ACM Mob. Comput. Netw. (MobiCom)*. 243–254.
- KAVRAKI, L. E., LATOMBE, J.-C., MOTWANI, R., AND RAGHAVAN, P. 1995. Randomized query processing in robot path planning. In *Proceedings of the 27th Annual ACM Symposium on Theory of Computing*. pp. 353–362.
- KAVRAKI, L. E., ŠVESTKA, P., LATOMBE, J.-C., AND OVERMARS, M. H. 1996. Probabilistic roadmaps for path planning in high dimensional configuration spaces. *IEEE Trans. Robot. Autom.* 12, 566–580.
- KIRKPATRICK, D., SNOEYINK, J., AND SPECKMANN, B. 2000. Kinetic collision detection for simple polygons. In *Proceedings of the 16th Symposium on Computational Geometry*. pp. 322–330.
- KIRKPATRICK, D. AND SPECKMANN, B. 2000. Separation sensitive kinetic collision detection for simple polygons. In *Proceedings of the Japan Conference on Discrete and Computational Geometry*.
- KIRKPATRICK, D. AND SPECKMANN, B. 2002. Kinetic maintenance of context-sensitive hierarchical representations of disjoint simple polygons. In *Proceedings of the 13th Annual ACM Symposium on Computational Geometry*.
- KRY, P. G. AND PAI, D. K. 2002. Continuous contact simulation for smooth surfaces. *ACM Trans. Graphics*, to appear.

- LAMIRAUX, F. AND KAVRAKI, L. 2001. Planning paths for elastic objects under manipulation constraints. *Int. J. Robot. Res.* 20, 188–208.
- LANG, J., PAI, D. K., AND WOODHAM, R. J. 2002. Robotic acquisition of deformable models. In *Proceedings of the IEEE International Conference on Robotics and Automation*. pp. 933–938.
- LARSON, E., GOTTSCHALK, S., LIN, M., AND MANOCHA, D. 2000. Fast distance queries using rectangular swept sphere volume. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 3719–3726.
- LASZLO, J., VAN DE PANNE, M., AND FIUME, E. L. 2000. Interactive control for physically-based animation. In *Proceedings of ACM SIGGRAPH*, pp. 201–208.
- LATOMBE, J.-C. 1991. *Robot Motion Planning*. Kluwer Academic Publishers, Boston, Mass.
- LEACH, A. R. 1996. *Molecular Modeling: Principles and Applications*. Addison-Wesley Longman, Rssex, England.
- LEVITT, M., GERSTEIN, M., HUANG, E., SUBBIAH, S., AND TSAI, J. 1997. Protein folding: The endgame. *Ann. Rev. Biochem.* 66, 549–579.
- LIN, M. AND GOTTSCHALK, S. 1998. Collision detection between geometric models: A survey. In *Proceedings of the IMA Conference on Mathematics of Surfaces*.
- LINDBERG, T. 1994. *Scale-Space Theory in Computer Vision*. Kluwer Academic Publishers, Boston, Mass.
- LONGUET-HIGGINS, H. C. 1981. A computer algorithm for reconstructing a scene from two projections. *Nature* 293, 133–135.
- LOTSTEDT, P. 1981. Coulomb friction in two-dimensional rigid body systems. *ZAMM* 61, 605–613.
- LOWRIE, R. B., ROE, P. L., AND VAN LEER, B. 1998. Space-time methods for hyperbolic conservation laws. In *Barriers and Challenges in Computational Fluid Dynamics. ICASE/LaRC Interdisciplinary Series in Science and Engineering*, vol. 6. Kluwer, pp. 79–98.
- MASON, M. T. 2001. *Mechanics of Robotic Manipulation*. MIT Press.
- MATOUŠEK, J. 1994. Geometric range searching. *ACM Comput. Surv.* 26, 421–461.
- MEYER, M., DEBUNNE, G., DESBRUN, M., AND BARR, A. H. 2001. Interactive animation of cloth-like objects in virtual reality. *J. Vis. Comput. Animat.* 12, 1–12.
- MIRTICH, B. V. 1996. Impulse-based dynamic simulation of rigid body systems. Ph.D. dissertation, Univ. California at Berkeley.
- MIRTICH, B. 1998. V-clip: Fast and robust polyhedral collision detection. *ACM Trans. Graph.* 17, 177–208.
- MIRTICH, B. 2000. Timewarp rigid body simulation. In *Proceedings of ACM SIGGRAPH*. pp. 193–200.
- ONUČIĆ, J. N., LUTHEY-SCHULTEN, Z., AND WOLYNES, P. G. 1997. Theory of protein folding: The energy landscape perspective. *Ann. Rev. Phys. Chem.* 48, 545–600.
- PAI, D. K., VAN DEN DOEL, K., JAMES, D. L., LANG, J., LLOYD, J. E., RICHMOND, J. L., AND YAU, S. H. 2001. Scanning physical interaction behavior of 3d objects. In *Proceedings of ACM SIGGRAPH*. pp. 87–96.
- PFOSE, D., JENSEN, C. J., AND THEODORIDIS, Y. 2000. Novel approaches to the indexing of moving object trajectories. In *Proceedings of the 26th International Conference on Very Large Databases*. 395–406.
- RICHMOND, J. L. AND PAI, D. K. 2000. Active measurement and modeling of contact sounds. In *Proceedings of the IEEE International Conference on Robotics and Automation*. pp. 2146–2152.
- RICHTER, G. R. 1994. An explicit finite element method for the wave equation. *Appl. Num. Math.* 16, 65–80.
- ROYER, E. AND TOH, C. 1999. A review of current routing protocols for ad-hoc mobile wireless networks. *IEEE Pers. Commun.* 6, 46–55.
- ŠALTENIS, S., JENSEN, C. S., LEUTENEGGER, S. T., AND LOPEZ, M. A. 2000. Indexing the positions of continuously moving objects. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*. pp. 331–342.
- SCHIRRA, S. 2000. Robustness and precision issues in geometric computation. In *Handbook of Computational Geometry*, J.-R. Sack and J. Urrutia, Eds. Elsevier Science Publishers, B. V. North-Holland, Amsterdam, The Netherlands, pp. 597–632.
- SISTLA, A. P. AND WOLFSON, O. 1995. Temporal conditions and integrity constraints in active database systems. In *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data*. pp. 269–280.
- SISTLA, A. P., WOLFSON, O., CHAMBERLAIN, S., AND DAO, S. 1997. Modeling and querying moving objects. In *Proceedings of the International Conference on Data Engineering*. pp. 422–432.
- SONG, G. AND AMATO, N. M. 2002. Using motion planning to study protein folding pathways. *J. Comput. Biol.* 9, 149–168.
- SORENSEN, H. 1985. *Kalman Filtering: Theory and Applications*. IEEE Press.
- TAYEB, J., ULUSOY, O., AND WOLFSON, O. 1998. A quadtree-based dynamic attribute indexing method. *Comput. J.*, 185–200.
- TEODORO, M., PHILLIPS, G., AND KAVRAKI, L. 2000. Singular value decomposition of protein conformational motions: Application to HIV-1 protease. In *Currents in Computational Molecular Biology*. Universal Academy Press Inc., Tokyo, Japan, pp. 198–199.
- TORR, P. H. S., SZELISKI, R., AND ANANDAN, P. 2001. An integrated Bayesian approach to layer

- extraction from image sequences. *IEEE Trans. Patt. Anal. Mach. Int.* 23, 297–303.
- ÜNGÖR, A. AND SHEFFER, A. 2002. Pitching tents in space-time: Mesh generation for discontinuous Galerkin method. *Int. J. Found. Comput. Sci.* 13, 201–222.
- ÜNGÖR, A., SHEFFER, A., HABER, R. B., AND TENG, S.-H. 2002. Layer based solutions for constrained space-time meshing. *Appl. Numer. Math.*, to appear.
- VAN DEN DOEL, K., KRY, P., AND PAI, D. K. 2001. FOLEYAUTOMATIC: Physically-based sound effects for inter active simulation and animation. *Computer Graphics (ACM SIGGRAPH 2001 Conference Proceedings)*. pp. 537–544.
- VOGLER, C. AND METAXAS, D. 2000. A linguistics-based approach to american sign language recognition. In *Proceedings of the International Conference on Gestures: Meaning and Use*.
- WOLFSON, O., CHAMBERLAIN, S., JIANG, L., AND MENDEZ, G. 1998. Cost and imprecision in modeling the position of moving objects. In *Proceedings of the International Conference Data Engineering*. pp. 588–596.
- WOLFSON, O., JIANG, L., SISTLA, A. P., CHAMBERLAIN, S., AND DENG, M. 1999. Databases for tracking mobile units in real time. In *Proceedings of the International Conference Database Theory*. pp. 169–186.
- WOLFSON, O., SISTLA, A. P., XU, B., ZHOU, S. J., AND CHAMBERLAIN, S. 1999. Domino: Databases for moving objects tracking. In *Proceedings of the SIGMOD International Conference on Management of Data*. pp. 547–559.
- YIN, L., ACHARYA, A., SOBH, N., HABER, R. B., AND TORTORELLI, D. A. 2000. A space-time discontinuous Galerkin method for elastodynamic analysis. In *Discontinuous Galerkin Methods: Theory, Computation and Applications*, B. Cockburn, G. Karniadakis, and C. Shu, Eds. Lecture Notes in Computational Science and Engineering, vol. 11. Springer-Verlag, New York, pp. 459–464.
- ZILLES, C. B. AND SALISBURY, J. K. 1995. A constraint-based god-object method for haptic display. In *Proceedings of the IEEE International Conference on Robotics and Automation*. pp. 146–151.

Received March 2001; revised August 2002; accepted August 2002