Issues in Geometric Rounding
Raimund Seidel, Saarland University, Saarbr\"ucken

Loosely speaking, ``geometric rounding'' refers to approximating a geometric object by another one that admits a simple representation using ``simple'' coordinate values. An interesting example concerns approximating a plane straight-edge graph by another one whose vertices have real coordinates that can be expressed as fixed point numbers using few bits. Already this seemingly easy example raises a number of interesting algorithmic and complexity questions. I will address some of them.

Rounding Arrangements of Line Segments: Balancing Priorities in Algorithm Development
John Hershberger, Mentor Graphics Corporation, USA

The problem of rounding an arrangement of line segments to the integer grid is important in many practical applications of geometric computing. In this talk I will use this problem as a case study to discuss the design of algorithms for problems important in practice. Algorithmic considerations include behavioral goals, correctness, simplicity, asymptotic running time and memory use, I/O sensitivity, and implementability. The process of algorithm design continuously balances these priorities against one another.

I will trace the history of segment rounding schemes, leading from early algorithms of [Greene and Yao] and [Milenkovic] to snap rounding [Greene; Hobby]. Snap rounding combines the virtues of simplicity, accuracy, and implementability. I will describe several implementations of snap rounding that successively improved its output-sensitivity, culminating in an edge-erasing algorithm that is simple, efficient, and near-optimal in its output-sensitivity.

Snap rounding is not idempotent (applying snap rounding to a snap rounded arrangement may change the result), and this causes problems in practice, as I will describe based on experience at Mentor Graphics. Previous attempts to address this problem are too complex for efficient implementation. I will describe a new variant of snap rounding, called *stable snap rounding*, that preserves all of ordinary snap rounding's good features---simplicity, accuracy, implementability---and in addition is idempotent: the result of applying stable snap rounding multiple times is the same as applying it once.

New Results in Numerical Subdivision Methods: Complex Roots and Isotopic Surface
Chee K. Yap, Courant Institute, NYU

We consider exact algorithms for curves and surfaces. In contrast to standard approaches based on algebraic (e.g., projection) or geometric (e.g., ray shooting) techniques, we focus on numerical techniques (e.g., evaluation). Numerical techniques tend to be practical with nice properties such as locality (e.g., solving within a box), adaptivity (e.g., fast for most inputs), and generality (e.g., analytic surfaces). The challenge in numerical techniques is how to make them yield exact results.

We look at two classical problems: complex root isolation and computing isotopic approximation of surfaces. We describe two new algorithms based on numerical subdivision approaches:
(i) CEVAL: complex root isolation algorithm [Sagraloff-Y]
(ii) Cxyz: an isotopic surface approximation algorithm based on parametrizability [Lin-Y]
These methods can be easily implemented. We will describe experimental results showing that both algorithms are competive with the current ``best practice'' algorithms. The experimental results for CEVAL is from a Masters thesis by Kamath at Oxford. Time permitting, I will discuss the complexity analysis of such subdivision algorithm. In general, the analysis of adaptive numerical algorithms is an exciting direction for algorithmics (e.g., continuous amortization [Burr-Krahmer-Y]).

Computing Homology Cycles with Certified Geometry
Tamal Dey, Ohio State U. , USA

Computation of cycles representing classes of homology groups is a fundamental problem arising in applications such as parameterization, feature identification, topology simplifications, and data analysis. Variations of the classical Smith Normal Form algorithm and the recently developed persistence algorithm do compute representative cycles of a homology basis for a simplicial complex, but they remain oblivious to the input geometry. Some recent research in computational topology has addressed the problems of computing homology cycles that are optimal with respect to a given metric. In this talk, we concentrate on two such developments: (i) Computing an optimal basis for one dimensional homology of a simplicial complex and using it to approximate such a basis for a smooth manifold from its point data; (ii) Computing an optimal cycle homologous to a given cycle in a simplicial complex. We provide efficient algorithms with their guarantees for (i) and show that classical Linear Programs can solve (ii) for some interesting cases even though the general problem is NP-hard.

Reeb graphs
Vijay Natarajan, IISc Bangalore

Isosurfaces, also called level sets, are extensively used for the visualization of three-dimensional scalar fields. The Reeb graph is an abstract representation of the topology of all level sets of a scalar function. It is obtained by mapping each connected component of the level sets to a point. The Reeb graph serves as an effective user interface for selecting meaningful level sets and for designing transfer functions for volume rendering. It also finds several other applications including topology-based shape matching, topological cleaning of surface models, surface segmentation, and parametrization. In this talk, I will first motivate the application of Reeb graphs for data analysis and visualization. Next, I will present an efficient two-step algorithm for computing the Reeb graph of a piecewise-linear (PL) function. The algorithm is output-sensitive, near-optimal, simple, and works in all dimensions. I will also present a method for effective presentation of Reeb graphs that enables it application to spatially-aware transfer function design.

Seamless Texture Atlas and Meso-textures
Subodh Kumar, IIT Delhi

This talk will describe techniques for displaying textured surfaces. The first part of the talk will focus on surface parameterization and texture charts. I will present an algorithm to produce a texture atlas and to map it seamlessly onto a surface. This algorithm eliminates the problem of seams at chart boundaries -- a problem that ails all textures atlases. The second part of the talk will introduce the notion of meso-textures and will present a technique to apply it as a displacement to surfaces using a fragment shader. I will describe a compressed intermediate representation of meso-textures that enables efficient but accurate ray-surface intersection in the shader.

Memory-Constrained Algorithms: Space-Time Tradeoffs
Tetsuo Asano, JAIST, Japan

Recent progress in computer systems has provided programmers with unlimited amount of work storage for their programs. Nowadays there are plenty of space-inefficient programs which use too much storage and become too slow if sufficiently large memory is not available. There is high demand for space-efficient algorithms or "memory-constrained algorithms." Another motivation for resource-constrained algorithms comes from applications to built-in or embedded software in highly functional hardware. Digital cameras and scanners are good examples. We measure the space efficiency of an algorithm by the number of work storage cells (or the amount of work space) used by the algorithm. Ultimate efficiency is achieved when only a constant number of variables are used in addition to the input storage (e.g., input arrays or lists). Another excellent example is a sensor network; with the drop of prices for flash memory even a large number of cheap sensors can be equipped with huge-volume flash drives. While the data measured by the sensor must be stored onboard for processing, and need to be written, it is preferable to avoid writing to the flash drive, as this is a slow and expensive operation which reduces the flash drive's lifetime. Hence we would like to use only higher level memory (e.g. only CPU registers) that we write into, and perform only read operations on the flash drive.

In some extreme cases an input array is assumed to be read-only, that is, we can read any array element in constant time, but no element can be modified. In addition, we are not allowed to use any array of size dependent on input size. We are allowed to use only a constant number of variables, each of $O(\log n)$ bits, where n is the input size. We call such an algorithm a "constant-work-space algorithm." It is also called ``log-space" algorithm in complexity theory since the amount of work space is $O(\log n)$ bits.

A more general interest is to investigate tradeoff between space and time complexities. It is true that more time is required with less work space. An objective is to derive more concrete dependencies of time complexity on the amount of work space.

In this talk I will introduce some basic algorithmic techniques for memory-constrained algorithms together with some new algorithms especially on computational geometry.

Geometric Transversal Theory in Three Dimensions
Otfried Cheong, KAIST, Korea

In my talk I will give a short introduction into geometric transversal theory, discussing geometric permutations, incompatible pairs of permutations, and Helly-type theorems. After discussing the classic two-dimensional results, I will present recent progress in three dimensions, in particular on geometric permutations, and on Helly-type theorems for isolated line transversals.

Simple Epsilon-Net Constructions
Sathish Govindrajan, IISc Bangalore

Epsilon nets are essentially hitting sets for dense objects and have applications in various geometric problems. In this talk, I will present some simple constructions that give linear-sized epsilon-nets for different geometric objects. I will also briefly discuss some interesting problems on small epsilon-nets.

Geometric problems related to Data Depth
Saurabh Ray, EPFL, Switzerland

The talk will survey the notions of Tukey depth (introduced by John Tukey in 1974) and Simplicial depth (first studied by Boros and Füredi in 1982) of a point set and mention two classical results, the Centerpoint theorem and the First Selection theorem which give optimal bounds for these two measures of depth. The main topic of the talk will be a recent notion of data depth called Ray Shooting depth (introduced by Fox, Gromov, Lafforgue, Naor and Pach in 2010). We will discuss several results and questions including a topological proof of the optimal bounds in the plane, algorithms, relations between the three notions of depth, and further extensions. Several open problems will be presented during the talk.

Stochastic Minimum Spanning Trees in Euclidean Spaces
Subhash Suri, UC Santa Barbara, USA

Consider a master set of n points in d-space, where each point is active (or present) with probability p_i, and inactive otherwise. How easy is to compute the expected length of the minimum spanning tree of this set? We investigate the computational complexity of this and related problems, and present hardness results and approximation algorithms.

Geometric computation on uncertain data
Pankaj K. Agarwal, Duke University & IIT Delhi

Computing geometric structures and answering queries on uncertain data have received much attention in the last few years due to the imprecise nature of many measurement data. This talk describes recent work on range searching and computing maxima in this model.

Estimating small frequency moments in nearly optimal space-time
Sumit Ganguly, IIT Kanpur

A data stream is a sequence of m updates of the form (i,v) where i is an item from $[n]$ and v is an integer in $[-M,M]$ signifying an update. The frequency of i, denoted f_i is the sum of updates v over the (i,v)'s appearing in the stream. For real $p \geq 0$, the pth moment of frequency, Fp is the sum of the pth powers of $|f_i|$. Estimating Fp to within factors of $1 \pm \varepsilon$ is a basic and well-studied problem. For p in [0,2] the space lower bound is known to be of order $(1/\varepsilon^2) \log (mM)$. An open problem in this direction has been to design an algorithm to estimate Fp that matches the space lower bound within small logarithmic factors and can process each stream update in time poly-logarithmic in *m,n* and *N*. We present the first algorithm that achieves both properties.

Orthogonal Range Reporting in 3 and Higher Dimensions
Lars Arge, Univ of Aarhus, Denmark

Orthogonal range reporting is the problem of storing a set of n points in d-dimensional space, such that the k points in an axis-orthogonal query box can be reported efficiently. While the 2-d version of the problem was completely characterized in the pointer machine model more than two decades ago, until recently this was not the case in higher dimensions.

In this talk the we describe techniques that can be used to obtain a space optimal pointer machine data structure for 3-d orthogonal range reporting that answers queries in O(log n + k) time. Thus we settle the complexity of the problem in 3-d. The techniques can also be used to obtain improved structures in higher dimensions, namely structures with a log n/ log log n factor increase in space and query time per dimension. Thus for d ≥ 3 we obtain a structure that both uses optimal $O(n(\log n/ \log\log n)^{d-1})$ space and answers queries in the best known query bound $O(\log n(\log n/ \log\log n)^{d-3} + k)$. At the end of the talk we will mention how the new techniques can also be used to obtain improved or even optimal external memory structures, and how we have also proved lower bounds show that the optimal query bound increases from O(log n+k) to $O((\log n/ \log\log n)^2+k)$ somewhere between three and six dimensions.

Mathematical Software Development, An Industrial Perspective
Shripad Kale, Geometric Ltd.

a) Challenges faced by our industrial OEMs (Automotive, Industrial, Aerospace etc.) and Software OEMs (CAD, CAM, CAE companies)
b) Major areas of work required in mathematical and geometrical areas like visualization, triangulation, meshing, geometric operations, interoperability, specialized algorithms.
c) More specific discussion on key challenges faced for CAD-CAE interoperability and specific issues
d) Work areas, challenges and technology development in medical domain.
e) Demonstrations of some of technologies developed by Geometric. (Volumetric rendering, Geometry reconstruction using medical scan (CT, MRI etc) data, Orthodontic CAD/CAM application)

Approximation Algorithms for Line Segment Coverage in Wireless Sensor Networks
Subhas Nandy, ISI Kolkata

The coverage problem in wireless sensor network deals with the problem of covering a region or parts of it with sensors. We address the problem of covering a set of line segments with minimum number of sensors. A line segment is said to be *covered* if it intersects the sensing region of at least one among the sensors distributed in a bounded rectangular region *R*. We assume that the sensing radius of each sensor is uniform. The problem of finding the minimum number of sensors needed to *cover* each member in a given set of line segments in *R* is NP-hard. We propose two constant factor approximation algorithms and a PTAS for the problem for covering a set of axis-parallel line segments. Finally, we show that a PTAS exists for covering a set of arbitrarily oriented line segments in *R* where the length of the line segments are bounded within a constant factor of the sensing radius of each sensor.

Linear-Time Approximation Schemes for Clustering Problems in any Dimensions
Yogish Sabharwal, IBM IRL Delhi

We present a general approach for designing approximation algorithms for a fundamental class of geometric clustering problems in arbitrary dimensions. More specifically, our approach leads to simple randomized algorithms for the *k*-means, *k*-median and discrete *k*-means problems that yield (1+ε) approximations with probability ≥ 1/2 and running times of $O(2^{(k/\varepsilon)^{O(1)}} dn)$. These are the first algorithms for these problems whose running times are linear in the size of the input (*nd* for *n* points in *d* dimensions) assuming *k* and ε are fixed. Our method is general enough to be applicable to clustering problems satisfying certain simple properties and is likely to have further applications.

The covert set cover problem
Sandeep Sen, IIT Delhi

We address a version of the set-cover problem where we do not know the sets initially (and hence referred to as covert) but we can query an element to find out which sets contain this element as well as query a set to know the elements. We want to find a small set-cover using a minimal number of such queries. We present a Monte Carlo randomized algorithm that approximates an optimal set-cover of size OPT within O(log *N*) factor with high probability using $O(\mathrm{OPT}\log^2 N)$ queries where *N* is the number of element in the universal set.

Hybrid Methods for Improving the Efficiency of Topology and Arrangement Computations
Michael Sagraloff, MPI Informatik, Germany

Abstract: Computing the topology of an algebraic curve/surface or computing an arrangements of algebraic objects constitutes a fundamental operation in geometric computing. So far, only algorithms based on elimination techniques such as resultants or Groebner Bases have proven practical, complete and exact at the same time. However, a major drawback of these approaches is that they suffer from a large amount of costly symbolic operations. We report on some recent results with respect to arrangement computation of planar algebraic curves. The proposed algorithm is designed in a way such that the amount of symbolic operations is reduced to a minimum. More precisely, only resultant computation and square-free factorization have to be considered. Furthermore, our implementation profits from a very efficient and highly parallel method for computing resultants on the graphics card. By outsourcing the resultant computation, we were able to eliminate one of the major bottlenecks of an elimination approach. In addition, we combine information obtained from the resultant computation with that obtained from a numerical solver to deduce our final result. We compared our methods with state-of-the art implementations on challenging benchmark instances. The results show that our hybrid method outperforms the existing approaches by far.