Computational Topology and Geometry: Supplementary Notes

Chee Yap

December 7, 2006

§1. Introduction and Course Mechanics

THIS is a document in progress, so you must be forgiving of mistakes. It is released under the belief that a flawed but timely document may be better than a perfect document that never sees the light of day. Please let me know of any errors!

I will add to this document throughout the course, so that we have only one document to deal with (there are pluses and minuses). The main purpose is to supplement the references listed below.

The assigned homework will always be listed at the end of these notes.

Programming. This course addresses computational issues in computational topology and geometry. As such we will need to do some programming to better appreciate the computational issues.

You will work in groups of 2 (they call this "extreme programming") and I can testify that it makes programming a fun and social acitivity, on top of its educational/intellectual content. I want you to download our "Core Library" from http://cs.nyu.edu/exact/ for doing your programming assignments.

If you live in a Windows environment, my best advice for using Core Library (and many other things!) is to download CYGWIN, a free Unix-like system that sits on top of Windows. Cygwin will have all the tools you need. For our course, I recommend these:

tar, Makefile,
g++ compiler,
some keyboard-based text editor (VIM or GVIM, emacs, etc)

My webpage http://cs.nyu.edu/ yap/prog/ has basic information on CYGWIN, Make and other programming stuff. All serious programmers must learn a keyboard-based editor – I highly recommend GVIM (or VIM, the non-GUI version).

References. The basic references listed below will be augmented with papers and additional notes as needed.

- Chapters from a forthcoming book, "Effective Computational Geometry for Curves and Surfaces" (Eds., J.-D.Boissonnat and M.Teillaud):
  A. Computational Topology: An Introduction, G.Rote and G.Vegter.
  B. Meshing of Surfaces, J.-D.Boissonnat, D.Cohen-Steiner, B.Mourrain, G.Rote, G.Vegter.

  You will need to read each of these chapters; these files are downloadable from our class page.
- Geometry and Topology for Mesh Generation, by H.Edelsbrunner. Cambridge Press, 2001.

  Mainly reference only.

- Topology for Computing, by A.Zomorodian. Cambridge University Press, 2005.

  Easy introduction to computational topology for computer scientists. Reference only.

- Robust Geometric Computation, by K.Mehlhorn and C.Yap.

  Information about numerical-algebraic computation will be from this book manuscript, available from my homepage.

- Fundamental Problems of Algorithmic Algebra, by C. Yap, Oxford Press 2000.

  Mainly reference for algebraic computation. Available from my homepage.

Munkres [5] is an excellent additional reference for algebraic topology. A modern undergrad text (from Springer) by Christine Kinsey is very accessible.

§2. Problems in Topological and Geometric Computation

A main motivation for our subject is the computational aspect of the geometry and topology of curves and surfaces. Such objects may be defined by its algebraic equation (e.g., $S : F(X, Y, Z) = 0$), or more constructively using some mesh with basis elements (e.g., B-splines), or perhaps by differential equations. We may wish to compute geometric properties of such objects (e.g., determine its location in space, or its curvature at particular point, or its singularities). We may wish to determine topological invariants (e.g., its Betti numbers, or determine if two paths on a surface are homotopic) which are global in nature.

The first step in any of the above tasks is to compute some more explicit, combinatorial, approximation of S from its defining equations or description. For example, from the equation $F(X, Y, Z) = 0$ of a surface, the basic properties of the surface are generally not obvious – where is it located in space, is it bounded, does it have singularities? A more explicit representation such as a piecewise linear approximation of the surface may allow some of these questions to be answered more directly. You could display this approximation for visual exploratoration of the surface. Visualization is an important tool for understanding geometric objects.

Such approximations are more generally called a cell complex. In applications, the cell complexes are piecewise linear and are known as **meshes**. There are two main classes of meshes: **surface mesh** and **volume mesh**, both embedded in 3-D. The computational task of converting a continuous characterization of a geometric object into a discrete approximation is called **meshing**. We regard view meshing as the critical step in the transition from continuous to discrete computation. If this step has error, any further computations may be rendered invalid. We have described the typical transition, from continuous to discrete because the continuous description is typically our starting point. But it is interesting to note that there are situations where we seek to reverse this transition. For instance, given a 2-D mesh, we may want to find a quadric surface that best approximates the mesh. Since the abstract continuous description is more compact, this **inverse meshing** can be seen as a data compression problem.

Meshes are very diverse as they arises in many disciplines, such as engineering, physical simulation, geometric design and architecture. In mathematics, we may want to compute meshes (perhaps in high dimensions) to visualize a complicated geometry, or to use it for computing topological invariants. What we call meshes are also known as "unstructured meshes", as distinguished from "structured meshes" whose vertices come from a fixed grid (e.g., $\mathbb{Z}^3$). In recent years, considerable interest is attached to geometric objects with even less structure than meshes: for instance, if we remove from the 2-dimensional cells from a surface mesh, we are left with a wire frame. If we further remove the 1-dimensional cells, we are left with only the vertices. Such a set is called a **point cloud**. The reason for this interest is the availability of new sensing technology and devices which can easily produce such point cloud models of physical models. We now have another form of inverse meshing – how to construct a surface or volume mesh from the point clouds?

The meshing problem thus encompasses a large variety of problems. We can initially classify them according to the nature of the input geometric description, and on the type of desired output mesh. Meshing computation involves a variety of numerical and algebraic techniques. We will briefly touch on some of these

techniques in this course. As a numerical computation, there are inevitable errors. So the big question in meshing is how to guarantee the geometric and topology correctness despite such errors. Correctness criteria, of course, must be clearly specified. Minimally, the mesh should have the same topology, i.e., homeomorphic to the input object. But as an approximation, we also would like to guarantee metric properties, that the mesh is close to the real geometry. Unfortunately, until recently, most published algorithms for meshing have no guarantees.

We clarify our remark about the ways that meshing algorithms may go wrong. First, a meshing algorithm may be using heuristics that are known to be incomplete (e.g., using Newton methods to search for zeros). A second source of error often escapes notice: even when the algorithm use provably correct methods, they may still be inadequate. Typically, the correctness of such algorithms assumes an ideal computational model where the numerical computation are error-free. But its implementation on a real computer may or may not introduce serious difficulties. We are very interested in this transition from ideal to realistic computational models. In recent years, much progress has been made in this direction. For instance, we now know a large class of problems where the translation from the ideal model to an actual model can be automatically achieved by software.

Next, suppose we have obtained a correct mesh representation. On the geometric side, there is the **mesh refinement problem**, i.e., to compute better approximations. This is usually an easier problem than computing the very correct mesh. In the zero-dimensional case, meshing can be viewed as finding zeroes of a real function. Finding the correct initial mesh corresponds to the root isolation problem; mesh refinement amounts to the root refinement (which could be solved by a simple binary search). Concerning mesh refinement, there is an interesting representation of curves and surfaces based **subdivision schemes**. As the name suggests, the refinement strategy for such surfaces is built into the representation, i.e., a predefined subdivision method is used. However, the global aspects of such representations may be difficult to recover.

On the topological side, we usually have no need for further mesh refinement. We just need the appropriate tools from algebraic topology to compute the topological invariants from the mesh. In this course, we will learn about some of these tools: homology, homotopy and Morse theory.

§3. Historical Perspective

The field of Computational Geometry started just over 30 years ago. For the most part, discrete computational problems on linear objects (points, lines, hypersurfaces, polytopes, line arrangements, etc) were the focus. In such a setting, the combinatorial aspects of computing dominates. An impressive set of algorithmic and analysis techniques have been developed over the last 30 years. In this course, we address the more recent interest of Computational Geometry involving nonlinear geometry (curves and surfaces) where the difficulties of continuous computation dominates. We will address the computational history of this topic in three phases:

1. Traditionally, computational scientists and engineers use numerical approximations to compute with curves and surfaces. The problem is that such methods are usually heuristic in nature. Such a numerical approach is still the dominant practice.

2. There was a growing movement in academic circles in the last 20 years to counteract this practice. The idea is to replace numerical computation by algebraic (or symbolic) techniques. The advantage is clear – algebraic techniques are precise and error-free. But the problems of the algebraic approach also begin to manifest themselves: such computations are too general (computes more than we really need) and are too slow for many practical problems. It is not just a matter of trying to find faster algorithms – in many cases, the slowness is intrinsic. As an example, we can solve polynomial equations by reduction to computing with the underlying ideals. This algebraic approach, however, captures more than just the geometry which is embedded in the radicals of the ideals. In a suitably general setting, computing with ideals requires double-exponential time, while the radical ideals is single-exponential time. Of course, even single-exponential time is not practical and so the search for special algorithms continues to be important in purely algebraic algorithms. But this is still not sufficient.

3. In recent years, another trend may be seen. That is the interest in combining algebraic with numerical techniques. This acknowledges the considerable merits of numerical techniques (namely it is fast), while rightly pointing out the need to make computations infallible, through a combination with algebraic techniques. For instance, the introduction of algebraic zero bounds with numerical approximation of roots (e.g., Core Library). This phase of development is still emerging. Although there are not many examples, we plan to look at some of these algorithms in this course.

Successful numerical-algebraic algorithms exhibit "adaptive" complexity. That means that the algorithms performs well for most inputs but not all; its complexity grows in proportion with its distance from singularities in the problem space. The challenge is how to quantify adaptivity.

§4. Review of Abelian Groups

The first part of our lectures is based on the chapter "Computational Topology: An Introduction" by Vegter and Rote (part of a book [3] to appear). We suplement their chapter with details or extensions.

The chapter deals with two topological tools: homology and Morse theory. In particular, since homology is defined through homomorphisms on Abelian groups, we must have some basic knowledge about Abelian groups. We shall write groups additively, e.g., a group G may written more explicitly as $(G, +, 0)$. All groups will be Abelian. Note that an Abelian group G can be viewed as a $\mathbb{Z}$-module where $(n, g) \mapsto ng$.

If G, H are groups, a **homomorphism** is $h : G \to H$ such that $h(x + y) = h(x) + h(y)$. If there is a bijective homomorphism between G and H, then we say they are **isomorphic** and write $G \simeq H$.

Let $S \subseteq G$. Then the subgroup **generated** by S, denoted $\langle S \rangle$ is the set of all finite sums,

$$x = \sum_{g_i \in S} n_i g_i \tag{1}$$

where $n_i \in \mathbb{Z}$. We call S a **generator** of G if $\langle S \rangle = G$. If S is a finite set, then we say G is **finitely generated**. Our main goal is to give a constructive proof of the **Fundamental theorem** of finitely-generated Abelian groups.

If S generates G with the additional property that each $x \in G$ has a unique expression of the form (1), then S is a **basis** of G. If G has a basis, then it is called a **free group**. The **rank** of a free group G is the number of elements in a basis of G.

Note that if G is free, then for all $x \in G$ and $n \in \mathbb{Z}$, the value $nx \neq 0$ for $n \neq 0$. But when $nx = 0$, then we say x is of **finite order**. The smallest $n > 0$ such that $nx = 0$ is called the **order** of x. The set $T = \{x \in G : nx = 0, (\exists) n \in \mathbb{Z}\}$ is called the **torsion subgroup** of G. If $T = \{0\}$, then we say G is **torsion free**.

Let $G_1, \ldots, G_n$ are groups. Then their **direct sum** is the group denoted $G = G_1 \oplus \cdots \oplus G_n$ where the underlying set is $G = G_1 \times \cdots \times G_n$ (Cartesian product) and the group operation is componentwise-operation. If the G_i's are Abelian, then so is G.

Lemma 1. *If H_i is a subgroup of an Abelian group G_i, then*

$$\frac{G_1 \oplus G_2}{H_1 \oplus H_2} \simeq \left(\frac{G_1}{H_1}\right) \oplus \left(\frac{G_2}{H_2}\right).$$

We leave this proof for an exercise. This lemma clearly extends to direct products of n groups, $G_1 \oplus \cdots \oplus G_n$.

Corollary 2. *If $G = G_1 \oplus G_2$ then $G/G_1 \simeq G_2$.*

Exercises

Exercise 4.1: Prove Lemma 1. ♢

Exercise 4.2: Show that if an Abelian group G is finitely generated and torsion-free, then it is free. Show that $\mathbb{Q}$ is torsion-free but not free. Conclude that $\mathbb{Q}$ is not finitely generated. ♢

End Exercises

§5. Smith Normal Form

A key computational tool in finitely generated Abelian groups is the Smith Normal form. See [7, chap. 10] for more details about Smith Normal form. You can download this from my webpage.

Let $A \in \mathbb{Z}^{m\times n}$ be an integer matrix. We say A is in **Smith Normal Form** (SNF) if A is diagonal, and these diagonal elements are $a_1, a_2, \ldots, a_{\min(m,n)}$, with the property that each $a_i \geq 0$ and

$$a_i | a_{i+1}$$

for all i.

Note that $n|0$ for all $n \in \mathbb{Z}$, and if $0|n$ then $n = 0$. This means that in SNF, any zero diagonal element $a_i = 0$ must appear later than any non-zero diagonal element $a_j > 0$, i.e., $i < j$. Similar, any a_i that is equal to 1 must appear before any other $a_j \neq 1$, i.e., $i < j$.

The non-zero diagonal elements are called **Smith invariants** or **invariant factors** of A.

For example, $A = \begin{bmatrix} 2 & 6 & 4 \\ 6 & 1 & 0 \end{bmatrix}$ is ...

A matrix $U \in \mathbb{Z}^{m\times m}$ is said to be **unimodular** if $\det U = \pm 1$. **Elementary row operations** are one of the following:
(C_i) multiply the ith row by -1,
(P_{ij}) permute the ith and jth rows,
($R_{ij}(c)$) replace ith row by c times the jth row (where $c \in \mathbb{Z}$ and $j \neq i$).

Each of these operations on A can be represented by a unimodular matrix $U \in \mathbb{Z}^{m\times m}$ and the corresponding transformation of A is given by matrix multiplication UA. Such a matrix U corresponding to elementary row matrices are called **elementary matrices**. Thus the matrices corresponding to three elementary row operations are:

$$C_i = \begin{bmatrix} 1 & & & & \\ & \ddots & & & \\ & & -1 & & \\ & & & \ddots & \\ & & & & 1 \end{bmatrix}, \quad P_{ij} = \begin{bmatrix} 1 & & & & & \\ & \ddots & & & & \\ & & 0 & & 1 & \\ & & & \ddots & & \\ & & 1 & & 0 & \\ & & & & & \ddots \\ & & & & & & 1 \end{bmatrix}, \quad R_{ij}(c) = \begin{bmatrix} 1 & & & & & \\ & \ddots & & & & \\ & & 1 & & & \\ & & & \ddots & & \\ & & c & & 1 & \\ & & & & & \ddots \\ & & & & & & 1 \end{bmatrix}.$$

The diagonal entries are all 1's unless otherwise indicated. It is also known that every unimodular matrices can be obtained as a product of elementary matrices. Two matrices A and B are row equivalent if $B = UA$.

Similarly, elementary column operations can be obtained by multiplying A on the right by an elementary matrix $V \in \mathbb{Z}^{n\times n}$. By an **elementary operation** we mean an elementary row or elementary column operation.

Let $\delta_k(A)$ be the GCD of all the order k minors of A. In particular, $\delta_1(A)$ is the GCD of all the entries of A. Recall that by convention, GCD is non-negative. It is easy to see from the definition that

$$\delta_k(A) | \delta_{k+1}(A).$$

We have

Theorem 3. *Every matrix A can be reduced to SNF by a sequence of elementary operations.*

Proof. Let A be an $m \times n$ integer matrix. For $i = 1, 2, \ldots, \min\{m, n\}$, we perform PHASE i: inductively assume that there are no non-zeros off the diagonal in the first $i-1$ rows and in the first $i-1$ columns.
1. If the ith row and ith column are all zero, we are done. Go to the next phase. Otherwise, move the smallest non-zero element from ith row or ith column to ith diagonal position (at a_{ii}). Make a_{ii} positive. In the following, a_{ii} will always be non-negative and non-increasing; any operation that strictly decreases a_{ii} is said to be "critical".
2. Using elementary column operations, we make each off-diagonal element in the ith row zero. By inductive hypothesis, $a_{ij} \neq 0$ implies $j > i$. To zero out a_{ij}, we subtract or add a multiple of column i from the column j. We choose the multiple so that the new a_{ij} is non-negative but strictly smaller than a_{ii}. There are two cases: (a) The new a_{ij} is now zero. (b) The new a_{ij} is non-zero. In the latter case, we exchange columns i and j, so that the minimum value of a_{ii} is reduced. Note that (a) or (b) cannot occur indefinitely often. Thus Step 2 must halt, and we go to Step 3.
3. Now, all the off-diagonal elements in the ith row is zero. Again, by elementary row operations, we make every off-diagonal element in the ith column 0. First by a row exchange and possibly multplication by -1, we ensure that the a_{ii}th entry is positive and smallest in magnitude in its column. If this exchange causes the ith row to have two non-zero entries, go back to Step 2. As long as there is some $j > i$ with $a_{ji} \neq 0$, we add or subtract a multiple of the ith row to row j to ensure $0 \leq a_{ji} < a_{ii}$. If $a_{ji} = 0$, we repeat this reduction for other choices of j. If $a_{ji} > 0$, we exchange rows i and j. If this exchange makes row i have more than one non-zero entry, we go back to Step 2. Note that we cannot go from Step 3 to Step 2 indefinitely often since each time this happens a_{11} decreases.
4. Eventually the ith row and ith column has only its diagonal element non-zero. We exit if $i = \min\{m, n\}$, otherwise, proceed to phase $i + 1$.

At the end of these phases, all non-zero entries are found along the diagonal of the matrix. We can assume these non-zero entries are positive. We can make the smallest one divide all the other entries. If we do not succeed, we would have found a strictly smaller smallest non-zero entry. This cannot continue indefinitely. Thus, eventually, the smallest non-zero entry divides all the other entries. We move this entry to position a_{11}, and proceed to ensure that the next smallest entry divides the remaining entries. When this is done, we move it to position a_{22}. This will eventually give us the SNF. **Q.E.D.**

The above algorithm is not meant to be efficient. There are efficient (polynomial time) algorithms for computing the smith invariants (see [Yap]). However, this is still fairly expensive operation.

Lemma 4.
(i) Elementary operations preserve $\delta_k(A)$.
(ii) Elementary operations preserve the rank and set of Smith invariants of A.

It follows that if S is the SNF of A and have invariant factos $a_1, \ldots, a_r$ (r is the rank), then $\delta_1(S) = a_1$, $\delta_2(S) = a_1 a_2$, etc. Since $\delta_i(A)$ are unique, it follows that $a_1, \ldots, a_r$ are unique.

Exercises

Exercise 5.1: Show that the algorithm used to prove the existence of SNF is exponential time. ♢

Exercise 5.2: (a) Implement in Core Library an algorithm for computing the SNF of a matrix A. The algorithm takes as input the matrix A and outputs unitary matrices U and V and S such that $S = UAV$ is SNF.

Do not worry about efficiency. NOTE: there is a Core Extension (COREX) for Linear Algebra that has matrices. Please use this. If necessary, extend the Core Extension

(b) Use your algorithm to compute the SNF of the matrix of ∂_1 and ∂_2 for the triangulation of S^2 above. ♢

End Exercises

§6. Finitely Generated Abelian Groups

The next fact is quite expected:

Lemma 5. *If B is a subgroup of A and A is a free Abelian group then B is free and $rank(B) \leq rank(A)$.*

As corollary, any subgroup B of $\mathbb{Z}$ is isomorphic to $\mathbb{Z}$.

Homomorphisms between free Abelian Groups. Let

$$f : F \to G$$

be a homomorphism between the are free Abelian groups F, G of ranks n and m, respectively.

Let $\mathbf{e} = (e_1, \ldots, e_n)$ be an ordered basis for F and $\mathbf{e}' = (e'_1, \ldots, e'_m)$ be an ordered basis for G. If

$$f(e_j) = \sum_{i=1}^{n} \lambda_{ij} e'_i \tag{2}$$

for $j = 1, \ldots, n$, then the matrix

$$\Lambda_f := (\lambda_{ij}) \in \mathbb{Z}^{m \times n}$$

is called the **matrix of f relative to the bases $\mathbf{e}, \mathbf{e}'$**. Then we have the following basic facts:

Lemma 6.
(i) The rank of $\mathrm{im}f$ *is the column rank of* Λ_f.
(ii) The rank of $\ker f$ *is* $n - rank(\mathrm{im}f)$ *(also called the nullity of* Λ_f*).*

Proof. (i) This follows from the fact that $\mathrm{im}f$ is generated by the columns of Λ_f. Of course, the column rank is the same as the rank of Λ_f.
(ii) This follows from the fact that $rank(\ker f) + rank(\mathrm{im}f) = rank(F)$. This is a generalization of the fact that, for vector spaces, $dim(\ker f) + dim(\mathrm{im}f) = dim(F)$. **Q.E.D.**

We introduce the useful "bar notation": for any $x \in F$, denote by

$$\overline{x} = (x_1, \ldots, x_n)^T \in \mathbb{Z}^n$$

the vector such that $\langle \overline{x}, \mathbf{e} \rangle = \sum_{i=1}^{n} x_i e_i = x$. This bar notation is relative to the choice of an ordered basis $\mathbf{e}$. For instance, relative to $\mathbf{e}$, each $\overline{e_i}$ is the ith elementary vector $(0, \ldots, 0, 1, 0, \ldots, 0)^T$ with a 1 in the i-th position but 0 everywhere else. Similarly, for $y \in G$, let $\overline{y} \in \mathbb{Z}^m$ such that $\langle \overline{y}, \mathbf{e}' \rangle = y$.

Fact 1. *Let Λ_f be the matrix of f (relative to some ordered basis $\mathbf{e}, \mathbf{e}'$). Then we have that for all $x \in F$,*

$$\overline{f(x)} = \Lambda_f \overline{x}.$$

Proof. It is sufficient to note that this lemma holds when $x = e_j$ is a basis element of $\mathbf{e}$: this follows from the definition of Λ_f in (2). **Q.E.D.**

We prove a normal for the matrix of homomorphisms between free Abelian groups:

Theorem 7 (Standard Bases for homomorphisms). *If F, G are free Abelian of ranks n and m, and*

$$f : F \to G$$

is a homomorphishm, then there exists ordered bases $\mathbf{d}, \mathbf{d}'$ for F and G such that the matrix Λ of f is in SNF. We call $\mathbf{d}, \mathbf{d}'$ the "standard bases" for f.

Proof. Let $A \in \mathbb{Z}^{m\times n}$ be the matrix of h relative to some ordered based $\mathbf{e}, \mathbf{e}'$. Let

$$S = UAV$$

be the SNF of A where U, V are unimodular matrices. For $x \in F$, we have

$$\overline{h(x)} = A\overline{x} = U^{-1}SV^{-1}\overline{x}.$$

Note that the vector

$$\overline{d_i} = V\overline{e_i}$$

denote the ith column of V. Let $d_i = \langle \overline{d_i}, \mathbf{e}\rangle$ be the corresponding element of F. Then $\mathbf{d} = (d_1, \ldots, d_n)$ is an ordered basis for F.

Now,

$$\begin{array}{rcll} \overline{f(d_i)} & = & A\overline{d_i} & \\ & = & (U^{-1}SV^{-1})(V\overline{e_i}) & \\ & = & U^{-1}S\overline{e_i} & (\overline{e_i} \in \mathbb{Z}^n) \\ & = & U^{-1}\overline{e'_i} & (\overline{e'_i} \in \mathbb{Z}^m) \\ & = & a_i\overline{d'_i} & (a_i \text{ is the } i\text{th Smith invariant}) \end{array}$$

where we define $\overline{d'_i} = U^{-1}\overline{e'_i}$. In the fourth line of this derivation, we use the fact that e_i is the ith elementary vector in $\mathbb{Z}^n$, and this implies $S\overline{e_i}$ is equal to $\overline{e'_i}$, the ith elementary in $\mathbb{Z}^m$.

Hence $\overline{d'_i}$ is the i-th column of U^{-1}. Moreover, the sequence $\mathbf{d}' = (d'_1, \ldots, d'_m)$ is an ordered basis for G. This proves that the matrix of f relative to the defined bases $\mathbf{d}, \mathbf{d}'$ is S. **Q.E.D.**

The next theorem gives a canonical form for the generators of subgroups of a finitely generated free Abelian group:

Theorem 8 (Standard Bases for Subgroups). *Let F be free Abelian of rank n and R a subgroup of F. Then there is an ordered basis $\mathbf{e} = (e_1, \ldots, e_n)$ for F, and positive integers r and $t_1, \ldots, t_k$ such that*

$$t_1|t_2|\cdots|t_k$$

and $\mathbf{d} = (t_1e_1, \ldots, t_ke_k, e_{k+1}, \ldots, e_r)$ is an ordered basis for R. We call $\mathbf{e}, \mathbf{d}$ the "standard bases" for (F, R).

Proof. By a previous lemma, we know that R is free of rank $r \leq n$. Let

$$j : R \to F$$

be the inclusion homomorphism. By our theorem on normal form for homomorphism, there exist ordered bases $\mathbf{e} = (e_1, \ldots, e_r)$ for R and $\mathbf{e}' = (e'_1, \ldots, e'_n)$ for F, such that the matrix of j relative to these bases is a matrix S in SNF.

Since j is 1-1, S has no zero column. Let $S = Diag(a_1, \ldots, a_r) \in \mathbb{Z}^{m\times r}$. Hence $j(e_i) = a_ie'_i$ $(i = 1, \ldots, r)$.

But since j is 1-1, we have $j(a_i) = a_i$, and the set $\{a_1e'_1, \ldots, a_re'_r\}$ is a basis for R. Suppose k of the elements $a_1, \ldots, a_r$ are greater than 1, then we can rename them to be $t_1, \ldots, t_k$; the remaining $r-k$ elements are all 1's. This gives the form desired by our theorem. **Q.E.D.**

Finally, we can prove:

Theorem 9 (Fundamental Theorem of Finitely Generated Abelian Groups). *Let G be a finitely generated Abelian group.*
(i) Then $G = H \oplus T$ where T is the torsion subgroup of G and H is a free.
(ii) There are finite cyclic groups $T_1, \ldots, T_k$ where T_i has order $t_i > 1$ such that $t_1|\cdots|t_k$ and $T = T_1 \oplus \cdots \oplus T_k$
(iii) The rank of H and $t_1, \ldots, t_k$ are determined by G.

Proof. Let $S = \{g_1, \ldots, g_n\}$ be a set of generators for G, and F be the free group with ordered basis $\mathbf{e} = (e_1, \ldots, e_n)$ generated by S. Consider the homomorphism,

$$h : F \to G$$

where $h(e_i) = g_i$. Then h is onto. Let $R = \ker(h)$. Applying the previous theorem, there are bases for R, F such that $F \simeq F_1 \oplus \cdots \oplus F_n$ and

$$R \simeq (t_1 F_1 \oplus \cdots \oplus t_k F_k) \oplus (F_{k+1} \oplus F_r).$$

But $G \simeq F/R$, and from Lemma 1,

$$\begin{aligned} G \simeq F/R &= (F \simeq F_1 \oplus \cdots \oplus F_n)/(t_1 F_1 \oplus \cdots \oplus t_k F_k) \oplus (F_{k+1} \oplus F_r) \\ &= (F_1/t_1 \oplus \cdots \oplus F_k/t_k) \oplus (F_{r+1} \oplus \cdots \oplus F_n). \end{aligned}$$

We can take $T_i = F_i/t_i$ and $H = F_1^{n-r}$. **Q.E.D.**

REMARK: How can we can represent ("present") finitely generated groups? Let G be an Abelian group. If $S \subseteq G$ is a finite set of elements that generates G, and we are told about all relations about S, then we can say that S is a presentation of G. Here is the formal way to describe this presentation. We are given an onto homomorphism

$$h : F \to G$$

where F is the free group generated by S, and we are given a set of relations of S that characterize all relations of $\ker h$. The relations are just linear combinations of S that equal 0. The matrix Λ and its null space can equally be used to represent this information.

§7. Homology of Simplicial Complex

I want to slightly rewrite the definitions in Vegter/Rote.

For $d \geq 0$, let $\{v_0, \ldots, v_d\}$ be a set of affinely independent points in $\mathbb{R}^m$. The convex hull of such a set is called a **simplex** σ. To be precise, it is the closed set

$$\sigma := \{\sum_{i=0}^{d} c_i v_i : c_i \geq 0, \sum_{i=0}^{d} c_i \leq 1\}.$$

Each v_i is a **vertex** of σ, and let $V(\sigma) = \{v_0, \ldots, v_d\}$ denote the set of vertices. The **dimension** of σ is $dim(\sigma) := |V(\sigma)| - 1$; a d-dimensional simplex is also called a d**-simplex**. Sometimes, it is useful to regard as special case the simplex σ_0 defined by the empty set, with no vertices: $V(\sigma_0) = \emptyset$ (empty set) is allowed, and it has dimension -1. Hence, $-1 \leq dim(\sigma) \leq m$.

Each subset of $V' \subseteq V(\sigma)$ defines a simplex τ; we call τ a **face** of σ and denote this relation by "$\tau \preceq \sigma$". If $0 \leq dim(\tau) < \dim(\sigma)$, then we call τ a **proper face** if σ. In the special cases of $\dim(\tau) = 0, 1, 2, 3, \dim(\sigma) - 1$ we call τ a **vertex**, **edge**, **triangle**, **tetrahedron** and **facet** of σ.

This terminology is familiar from elementary geometry. In fact, as we define the remaining concepts, it would be helpful for readers to keep in mind the example of a triangulation of a polygonal subset of the plane. Consider Figure 1(a). The first step in "algebraization" of topology is to introduce a direction to edges, and a clockwise/ counter clockwise order to triangles. In Figure 1(b), we arbitrarily directed each edge from the smaller indexed vertex to the larger indexed vertex. Also, each triangle is arbitrarily given the counterclockwise direction.

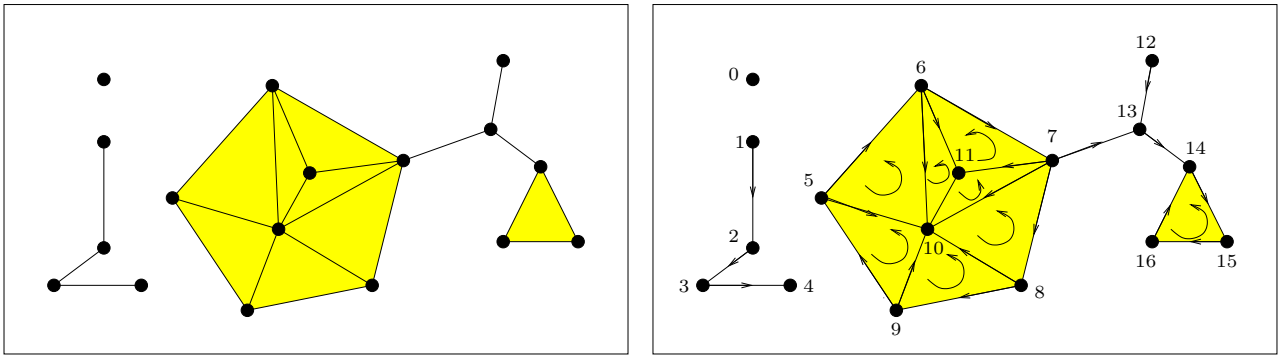

(a) (b)

Figure 1: A Triangulation in the plane

Oriented Simplices. Given a simplex $\sigma = \{v_0, \ldots, v_d\}$, an **ordered simplex** is the sequence of its vertices, $(v_{\pi(0)}, \ldots, v_{\pi(d)})$ where π is a permutation on the index set $\{0, 1, \ldots, d\}$. Two ordered simplices are said to be **equivalent** if their underlying permutations π and π' differ by an even number of transpositions. This is an equivalence relations. Each equivalence classes is known an **oriented simplex** based on σ. Let $[v_{\pi(0)}, \ldots, v_{\pi(d)}]$ denote the oriented simplex correspond to $(v_{\pi(0)}, \ldots, v_{\pi(d)})$.

Note that if $\dim(\sigma) > 0$, then there are two equivalence classes. We arbitrarily choose an orientation class as the **positive orientaion**, and the other (if any) is **negative**. If $\dim(\sigma) = 0$, there is only one equivalence class. If $\dim(\sigma) = -1$, then σ has no equivalence classes.

Let $[K]_d$ denote the set of all oriented d-simplices of K. Also, let $[K] = \bigcup_{d \geq 0} [K]_d$. A subset of $B \subseteq [K]$ is **complete** if, for each simplex $\sigma \in K$, we choose exactly one oriented simplex based on σ. Call K, associated with such a complete set B, an **oriented simplicial complex**. The particular choice of B is not important, but we must be consistent in sticking through with our choice. For instance, in Figure 2(b), we depict a simplicial complex for the 2-sphere, and we choose for each 2-simplex an orientation (counterclockwise). When B is understood, then for $\sigma \in K$, we may write $[\sigma]$ for the oriented simplex based on σ which corresponds to our choice B. It is sometimes convenient to let $[\sigma]^-$ denote the oppositely oriented simplex corresponding to $[\sigma]$. Strictly speaking, the notation $[\sigma]^-$ is undefined when $\dim(\sigma) \leq 0$; but in the context of simplicial chains, we can interpret $[\sigma]^-$ to be $-[\sigma]$.

Compact, connected 2-manifolds. A primary source of examples of simplicial complexes will be 2-dimensional. So you must become familiar with the following basic examples: 2-sphere S^2, torus T^2, Klein bottle, real projective plane $\mathbb{P}^2(\mathbb{R})$. These are all examples of connected, bounded 2-manifolds (S^2 and T^2 are orientable surfaces, but Klein bottle and $\mathbb{P}^2(\mathbb{R})$ are non-orientable).

Every connected, bounded 2-manifolds can be topologically[1] represented as a convex polygon whose directed edges are identified in pairs. Conversely, every convex polygon whose directed edges are identified in pairs represents such a manifold. This is illustrated in Figure 2(a), where we draw a hexagon (a, b, c, d, e, f) where we identify the directed edges ba with bc, the directed edges dc with de, and the directed edges fa with fe. It easy to see that this really describes the boundary of a tetrahedron. After these identifications, we see that the vertices a, c, e are the same vertex. Similarly, Figure 3(a) shows the polygonal representation of a torus, and Figure 3(b) is the polygonal representation of a Klein bottle. Because this is non-orientable, this surface may be a little harder to understand (for instance, if we try to embed this surface in 2-dimensional space, it will self-intersect). Note that the four vertices of the rectangles in Figure 3 are identified: $[a] = [b] = [c] = [d]$.

[1] I.e., up to homeomorphism.

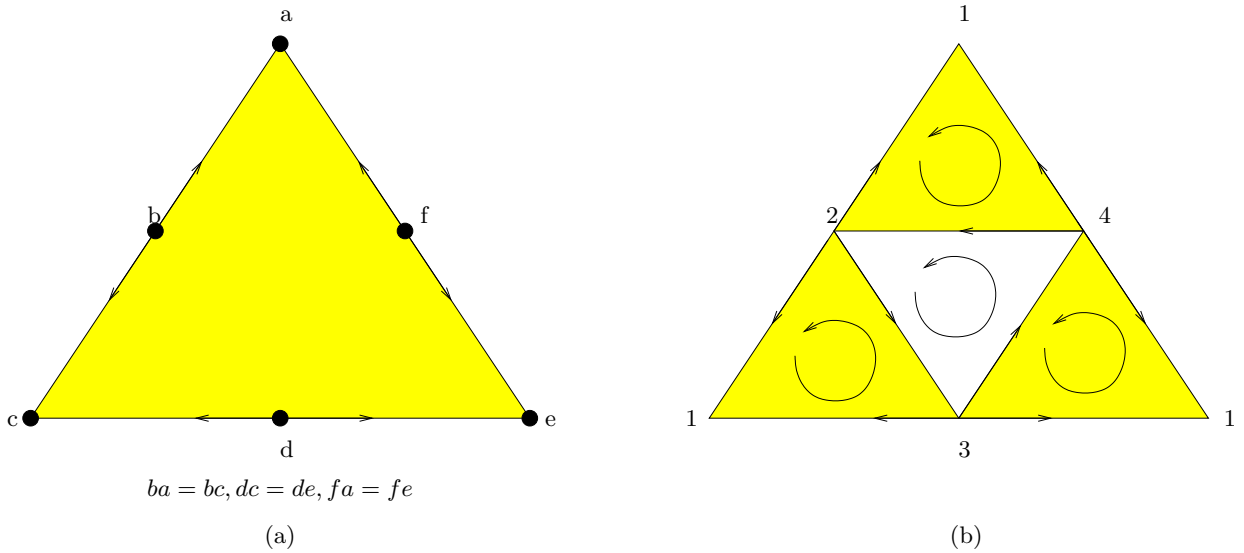

Figure 2: Topological 2-sphere: (a) Boundary of a tetrahedron, (b) Its triangulation

Chain Groups. Let K be a n-dimensional simplicial complex, and G be any commutative ring. Usually, $G = \mathbb{Z}$ but we might also consider $G = \mathbb{Q}, \mathbb{C}, \mathbb{Z}_2$. We can define a **simplicial d-chain** of K to be any function $c : [K]_d \to G$ with the constraint that for any d-simplex σ where $d \geq 2$, we have

$$c[\sigma]^- = -(c[\sigma]).$$

In other words, the coefficients to the two opposite orientations of σ are interdependent.

For any $d = 0, \ldots, n$, let $C_d(K; G)$ denote the set of all simplicial d-chains. We turn $C_d(K; G)$ into a Abelian group by defining the group operation $c + c'$ via $(c + c')[\sigma] = c[\sigma] + c'[\sigma]$. In fact, $C_d(K; G)$ is a free Abelian group generated by its set of (oriented) d-simplices.

We view $C_d(K; G)$ as a G**-module** such that if $g \in G$ (scalar) and $c \in C_d(K; G)$ (vector) then the scalar-vector product $gc \in C_d(K; G)$ is defined by $(gc)[\sigma] = g(c[\sigma])$. Then we may verify that C_d is a G-module:

$$g(c + c') = gc + gc', \quad (g + g')c = gx + g'x, \quad (gg')c = g(g'c).$$

We call C_d the G**-module of simplicial d-chains**.

When convenient, we may write $C_d(K)$ or even C_d for $C_d(K; G)$. The three main examples of G are $G = \mathbb{Z}, \mathbb{Z}_2, \mathbb{Q}$. The main advantage of $G = \mathbb{Q}$ is that $C_d(K; G)$ becomes a vector space. This is the viewpoint of the chapter by Rote/Vegter; it is also our default assumption in these notes.

Boundary Operator. For each $d \geq 1$, let $\partial_d : C_d(K) \to C_{d-1}(K)$ such that

$$\partial_d[v_0, \ldots, v_d] = \sum_{i=0}^{d} [v_0, \ldots, \widehat{v_i}, \ldots, v_d].$$

and ∂_d is extended to C_d via linearity. When $d = 0$, we define $\partial_0 : C_d \to 0$ where $\partial_0[\sigma] = 0$ for all σ.

We may drop the subscript in ∂_d, and simply write ∂ when the context is unambiguous or d is irrelevant.

Note that ∂ is a G-module homomorphism, i.e., for all $g \in G, c \in C_p$, we have

$$\partial(c + c') = \partial(c) + \partial(c'), \quad \partial(gc) = g\partial(c).$$

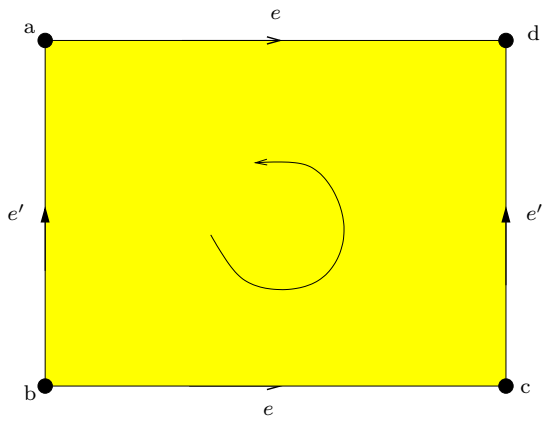

(a) Torus (b) Klein bottle

Figure 3: (a) Torus, (b) Klein bottle

When C_p is a vector space, ∂_p is just a linear transformation between two vector spaces.

We thus have a sequence of module homomorphisms,

$$\cdots \stackrel{\partial_{p+1}}{\to} C_p \stackrel{\partial_p}{\to} C_{p-1} \to \cdots$$

For any d-chain c, we call $\partial(c)$ the **boundary** of c. The main property of this boundary operator is this:

Theorem 10. $\partial \circ \partial = 0$

Proof. The proof amounts to expanding the two applications of the boundary operator, written as a double summation. With care, we show that the double summation vanishes. See Hilton for a short proof. **Q.E.D.**

Given this boundary operator, we can define three interesting groups (or modules): let $d = 0, \ldots, n$.

1. $Z_d(K) = \ker \partial_d$: the d-th dimensional group of **cycles**.
2. $B_d(K) = \mathrm{im}\partial_{d+1}$: the d-th dimensional group of **boundaries**.
3. $H_d(K) = Z_d(K)/B_d(K)$: the d-th dimensional group of **homology cycles**.

The fundamental problem of homology theory is to determine H_d up to group isomorphism. Elements of Z_d and B_d are called **cycles** and **boundaries**, respectively. Since H_d is somewhat abstract, it is useful to develop some language for discussing it: we say two d-chains c, c' are **homologous** to each other, denoted $c \sim c'$, if $c - c' \in B_d = \mathrm{im}\partial_{d+1}$. This is an equivalence relation, and its equivalence classes are called **homology classes**. The elements of H_d are just the homology classes of the d-cycles (i.e., elements of Z_d). Thus, if $c \in \mathrm{im}\partial_{d+1}$, we say c is homologous to 0. In this case, $c = \partial(c')$ for some $(d+1)$-chain c', and we say c **bounds** c' (or simply, c **bounds**).

REMARK: We can be more abstract, by discarding all notions of geometric complexes (simplicial or otherwise), and simply study a "chain complex" as a sequence

$$\cdots \stackrel{\partial_{p+1}}{\to} C_p \stackrel{\partial_p}{\to} C_{p-1} \to \cdots$$

of G-modules C_p and homomorphisms ∂_p's satisfying $\partial_{p+1} \circ \partial_p = 0$.

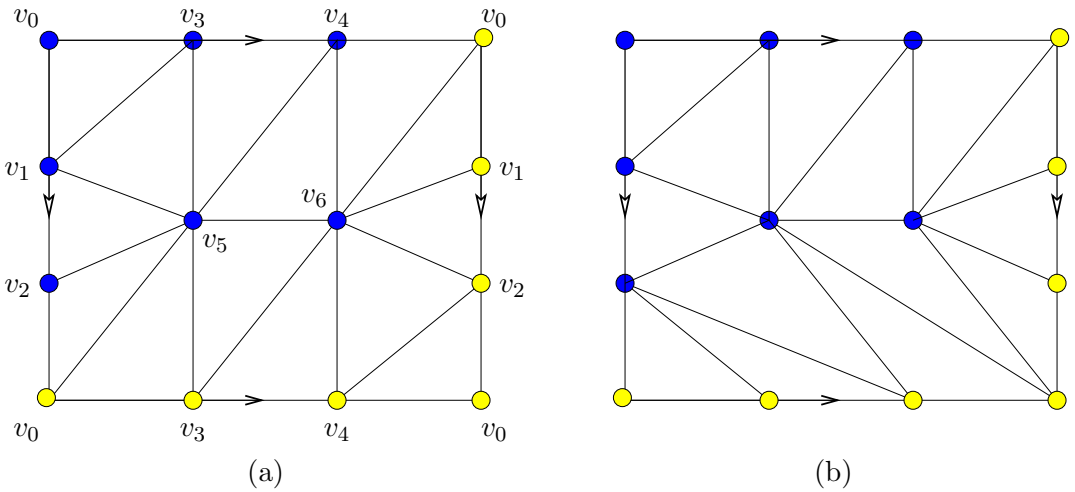

Figure 4: 7-vertex triangulation of the torus. (a) Bad (b) Good

Zero Dimensional Homology. For any simplicial complex K, its 0th homology group $H_0(K)$ is easy to describe: *it is isomorphic to the group* $\mathbb{Z}^m$ where m is the number of connected components of the set $|K|$. Let us prove this. Let $G = (V, E)$ be the undirected graph where $V \subseteq K$ is the set of vertices in K and $E \subseteq K$ is the set of edges in K. The chain group $C_0(K)$ is generated by the set $\{[v_1], \ldots, [v_k]\}$, where $V = \{v_1, \ldots, v_k\}$. Since $\partial[v_i] = 0$ for all i, the cycle group $Z_0(K)$ is equal to $C_0(K)$.

We next determine the boundary group $B_0(K)$. Let $G_1, \ldots, G_m$ be the subgraphs of G, corresponding to the m connected components of G. Choose a spanning tree $T_i \subseteq E$ for each G_i. We may characterize T_i as a set of edges such that for every pair of vertices $u, v \in G_i$, there is a unique sequence path $p(u, v)$ in T_i from u to v. If $e = (u, v)$, then $\partial[e] = \partial[p(u, v)]$. In other words, *for all* $e \in E$, *there is a unique path* $p(e)$ *in some* T_i *such that* $\partial[e] = \partial[p(e)]$. Let $T = \cup_{i=1}^m T_i$. Since $B_0(K)$ is generated by $\{\partial[e] : e \in E\}$, we conclude that it is in fact generated by $\{\partial[e] : e \in T\}$. Pick a representative vertex u_i in each G_i: so every vertex in G_i is homologous to $[u_i]$. Further, $[u_i]$ generates a free group isomorphic to $\mathbb{Z}$. It follows that every 0-chain in G_i is homologous to $n_i[u_i]$ for some $n_i \in \mathbb{Z}$. Every 0-chain $c \in C_0(K) = Z_0(K)$ can be uniquely decomposed into $c = \sum_{i=1}^m c_i$ where c_i is an 0-chain in G_i. But each c_i is homologous to some $n_i[u_i]$. Hence c is homologous to $\sum_{i=1}^m n_i[u_i]$, a chain in the free group generated by $\{[u_1], \ldots, [u_m]\}$. So each element of $H_0(K)$ is an equivalence class of the form $\sum_{i=1}^m n_i[u_i]$. This proves that $H_0(K)$ is isomorphic to the free Abelian group generated by $[u_1], \ldots, [u_m]$, i.e., $H_0(K) \simeq \mathbb{Z}^m$.

Homology of Euclidean Balls and Spheres. Besides compact 2-manifolds, the other main set of canonical examples from from Euclidean balls and spheres. Let us compute their homologies.

For $n \geq 1$, let $B^n = \{p \in \mathbb{R}^n : \|p\| < 1\}$ denote the open unit n-ball in $\mathbb{R}^n$. Also let $\overline{B}^n$ denote the closure of B^n, and $\dot{B}^n = \overline{B}^n \setminus B^n$ denote its boundary. We also call $\dot{B}^n$ the **unit** $(n-1)$**-sphere**, denoted S^{n-1}. For instance, B^1 is the open interval $(-1, 1)$ and $S^0 = \{-1, 1\}$.

Let us compute the homology of $\overline{B}^n$ ($n \geq 1$). This amounts to computing $H_k(K_n)$ for all $k = 0, \ldots, n$, where K_n is the simplicial complex comprising the faces of a n-simplex. So K_n has $n+1$ vertices, $\binom{n+1}{2}$ edges, $\binom{n+1}{3}$ triangles, etc. Let its vertices be $v_0, \ldots, v_n$.

Consider the dth homology group $H_d = H_d(K_n)$ for $d = 0, \ldots, n$. Two cases are easily determined:

- (a) $d = 0$: from the previous discussion of the 0th homology group, we know that $H_0 \simeq \mathbb{Z}$ (since $|K_n|$ is connected).
- (b) $d = n$: since there are no $(n+1)$-simplices, $\mathrm{im}\partial_{n+1} = B_n \simeq 0$. There is only one n-simplex $[v_0, \ldots, v_n]$, and $\partial_n[v_0, \ldots, v_n] = [v_1, \ldots, v_n] - [v_0 v_2, \ldots, v_n] + \cdots$ does not vanish. Since the cycles in Z_n must be generated by $\partial_n[v_0, \ldots, v_n]$, we conclude that $Z_n \simeq 0$. Thus $H_n = Z_n/B_n \simeq 0$.

It remains to consider $d = 1, \ldots, n-1$. Let us give a specialized argument for $d = 1$. We show $H_1(K_n) \simeq 0$. Let v_0, v_1, v_2, v_3 be the vertices of K_n. Clearly, $B_1 = \mathrm{im}\partial_2$ is generated by the set of 1-boundaries $b_{ijk} = \partial[v_i v_j v_k]$'s where $0 \leq i < j < k \leq n$. Let e_i ($i = 1, \ldots, N$ where $N = \binom{n}{2}$) be a complete set of oriented edges of K_n, and consider the 1-cycle $c = \sum_{i=1}^N n_i e_i$. It suffices to prove that c can be written as a linear combination of the b_{ijk}'s. We use induction on $t(c) = \sum_{i=1}^N |n_i|$. If $t = 0$, we are done. Otherwise, suppose $n_1 \neq 0$. Wlog, $e_1 = [v_0, v_1]$ and $n_1 > 0$. Wlog, there must be some v_2 such that $e_2 = [v_1, v_2]$ and $n_2 > 0$ (otherwise, $[v_1]$ would have a non-zero coefficient in δc, contradicting $\delta c = 0$). Consider the 1-cycle $c' = c - n_1 b_{012}$. Note that $t(c') \leq t(c) - n_1$ (because the coefficient of $[v_2, v_0]$ in c' can at most increase by n_1 but the coefficients of $[v_0, v_1]$ and $[v_1, v_2]$ each decreased by n_1). So far, we have not specified the coefficient ring for our homology: in case the coefficients come from $\mathbb{Z}$, it is clear that this process must clearly terminate. But even if n_i's come from $\mathbb{Q}$, we see that it must terminate because $t(c')$ belongs to the ideal of $\mathbb{Q}$-combinations generated by $(n_1, \ldots, n_N)$. Unfortunately, this argument does not easily generalize to 2-cycles.

We now show that for all $d = 1, \ldots, n-1$, every d-cycle of K_n is a boundary of some $(d+1)$-cycle. In other words,

$$H_d(K_n) \simeq 0. \tag{3}$$

Our proof uses induction on n. This exploits the observation that the n-simplex K_n is a **cone** $K_n = C(v_n, K_{n-1})$ over the $(n-1)$-simplex K_{n-1}. By definition, the simplices of a cone $C(v_n, K_{n-1})$ is one of two types: σ or $[v_n, \sigma]$, where σ is a simplex of K_{n-1}. Note that $[v_n]$ can be regarded as a special case of the latter type with $\sigma = \emptyset$. More generally, if c is any $(d-1)$-chain of K_{n-1}, let $[v_n, c]$ denote a d-chain of $C(v_n, K_{n-1})$. Moreover, we have

$$\partial[v_n, c] = c - [v_n, \partial c].$$

Hence any d-chain c of K_n can be written as a sum of two chains,

$$c = c' + [v_n, c'']$$

where c', c'' are d- and $(d-1)$-chains of K_{n-1}. To prove our result, assume that c is a cycle, $\partial c = 0$, and we must show that c is a boundary. Hence $0 = \partial c = \partial c' + c'' - [v_n, \partial c'']$. This means

$$\partial c' + c'' = 0, \qquad \partial c'' = 0$$

From the first equation, we see that $c'' = -\partial c'$, and so

$$c = c' - [v_n, \partial c'].$$

But this equation shows that

$$c = \partial[v_n, c'],$$

i.e., c is the boundary of the chain $[v_n, c']$. This completes the general argument for a general d. We leave the case $d = n-1$ as an exercise.

The following slick proof came up from the homework interviews with Gale Morehouse and Michael Burr. Let us use the following elementary fact from the incremental Betti number algorithm: when you add a d-simplex to a complex, you either increase β_d by 1 or decrement β_{d-1} by 1.

Thus, we have shown computed the Betti numbers of the Euclidean n-balls for all $n \geq 1$ and $d = 0, \ldots, n$:

$$\beta_d(\overline{B}^n) = \begin{cases} 1 & \text{if } d = 0 \\ 0 & \text{else.} \end{cases}$$

Building on this result, we next compute the homology of n-spheres S^n for $n \geq 0$.

First, assume $n \geq 1$. Let ∂K^{n+1} denote the triangulation obtained from K^{n+1} by removing its sole $(n+1)$-simplex. Clearly, S^n is homeomorphic to the support $\left|\partial K^{n+1}\right|$. So $Z_d(\partial K^{n+1}) = Z_d(K^{n+1})$ for all $d = 0, \ldots, n$. Also $B_d(\partial K^{n+1}) = B_d(K^{n+1})$ for all $d = 0, \ldots, n-1$, and $B_n(\partial K^{n+1}) = 0$. Thus

$$H_d(\partial K^{n+1}) = H_d(K^{n+1})$$

for all $d = 0, \ldots, n-1$. Also, assuming the coefficient ring is $\mathbb{Q}$, we see that $H_n(\partial K^{n+1}) \simeq \mathbb{Q}$ since $Z_n(\partial K^{n+1}) \simeq \mathbb{Q}$ and $B_n(\partial K^{n+1}) \simeq 0$.

If $n = 0$, then S^0 is just a pair of points, and clearly we have $H_0(S^0) \sim \mathbb{Z}^2$.

Cell Complexes. Although simplicial complexes are easy to understand, their use in computating homology can be tedious (by hand) because we will need many simplices even for simple topological spaces. For instance, the smallest triangulation of the torus requires 7 vertices, 21 edges and 14 triangles. Computing homology using a complex with so many cells is pushing the limits of casual hand computation. It turns out that by generalizing simplicial complexes to "cell complex", one can greatly reduce the number of cells, and bring many simple topological spaces within reach of hand computation.

Before defining the concept, let us see some natural examples of cell complexes. Figure 3 shows the cell complexes of the torus and Klein bottle. In each case, we begin with a rectangle complex: four vertices, four edges and the interior of the rectangle (a 2-cell). Then we identify the opposite edges in pairs: depending on how we do the identification, we get difference surfaces. Here, the torus and Klein bottle are indicated. After identification, we are left with a vertex v, two edges e, e' and a 2-cell c. The set $K = \{v, e, e', c\}$ is our the cell complex.

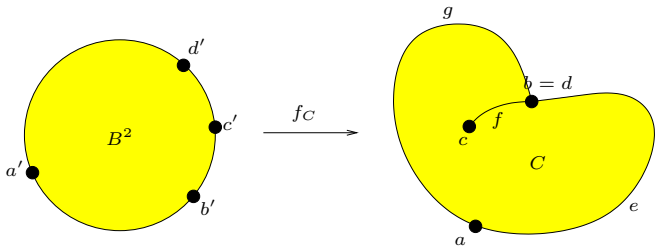

Figure 5: Cell Complex

A **d-cell** is any subset of $\mathbb{R}^n$ that is homeomorphic to of B^d ($d \geq 1$); a **0-cell** is just a singleton set. Let K be any non-empty finite collection of pairwise disjoint cells, where $|K| = \cup K$ is a Hausdorff space. But for our purposes, we may assume $|K| \subseteq \mathbb{R}^n$ for some n. We call K a **cell complex** if, for each d-cell $C \in K$, $d \geq 1$, there is a continuous function $f_C : \overline{B}^d \to |K|$ such that f_C is a homeomorphism from B^d onto C. It can be shown that this implies that $f_C(\dot{B}^d)$ is equal to a union of cells in K [5, p. 215]. NOTE: This is essentially the definition of CW-complex, except that we avoid the complications that arise in CW-complex with infinitely many cells.

In Figure 5, we have $K = \{a, b, c, e, f, g, C\}$ where a, b, c are 0-cells, e, f, g are 1-cells, and C is a 2-cell. The map f_C from $\overline{B}^2$ onto $|K|$ (where $f(a') = a, f(b') = b$, etc) shows that K is a cell complex. Note that

f_C is not a necessarily a homeomorphism of $\overline{B}^d$ as seen in this example. When this extra condition is true for every cell C, we call the cell complex **regular**.

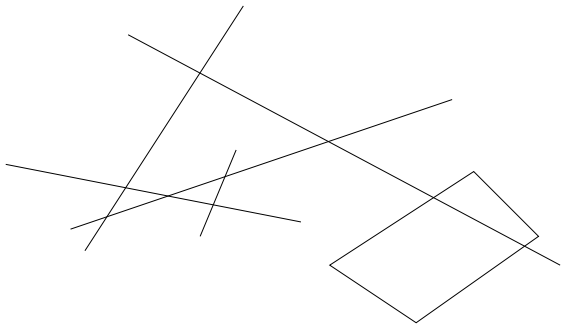

Figure 6: Cell complex from an arrangement of plane line segments

In Computational Geometry, one studies the cell complexes that arise from a set of plane line segments. This is illustrated in Figure 6. In particular, a basic question is the combinatorial of a 2-cell in an arrangement of n line segments.

We can now define cycle groups $C_d(K)$ whose bases are the oriented d-cells of K. The boundary operator ∂ is similarly defined. For instance, in Figure 5, assuming a counterclockwise orientation of C, and with appropriate orientation for edges e and g, we obtain

$$\partial_2 C = [e] + [f] - [f] + [g] = [e] + [g].$$

We again obtain the cycle group $Z_d(K)$, boundary group $B_d(K)$ and homology group $H_d(K)$. The homology groups again depends on the underlying topology of $|K|$. Since a simplicial complex is a special case of a cell complex, and if we accept the fact that the homology is independent of the subdivision of a space into a complex, we conclude that $H_d(K)$ is the same object as that defined using simplicial complexes.

Exercises

Exercise 7.1: Munkres [5, p. 34] noted that in the modern view, obtaining the homology groups of a space is regarded as more important than the classical view of just computing numerical invariants such as Betti numbers or Euler characteristics. Give an example of the information we might want from homology groups that is not available from its Betti numbers. ♢

Exercise 7.2: Figure 4 shows two proposed triangulation of the torus T^2. 7 vertices, 21 edges and 14 triangles.
(a) Why is Figure 4(a) not a triangulation of T^2? Verify that Figure 4(b) is indeed is a triangulation.
(b) Show every triangulation of T^2 satisfies $v \geq 7, e \geq 21, f \geq 21$. Thus Figure 4(b) is a minimal triangulation in a very strong sense. HINTS: In a triangulation, any two vertices determine at most one edge, any three vertices determine at most one face. Also, Euler's characteristic for a torus says that $v - e + f = 0$. You will need another relation involving v, e, f.
(c) What space does Figure 4(a) represent? Compute its homology groups. ♢

Exercise 7.3: Consider the 2-sphere in Figure 7.7 of Vegter/Rote. They have provided canonical bases for C_0, C_1 and C_2 in their notes, and the matrix Λ_i of δ_i $(i = 1, 2)$ relative to these bases were given. We want you to choose bases so that the corresponding matrix is in SNF. HINT: follow the proof of Theorem 7. ♢

Exercise 7.4: Let K be a simplicial complex with n connected components. Then the group $H_0(K)$ is free Abelian with basis given by a set $S = \{[\sigma_i] : i = 1, \ldots, n\}$, where each connected component of $|K|$ is represented by a unique vertex in S. ♢

Exercise 7.5: Most of our examples do not demonstrate torsion (in fact, all subspaces of Euclidean space has no torsion). To see how torsion arise, compute the homology groups of the Klein bottle. ♢

Exercise 7.6: Construct a space whose some homology group contains the torsion group $\mathbb{Z}_3$. HINT: consider how $\mathbb{Z}_2$ arises in the Klein bottle. ♢

Exercise 7.7: (Relative Homology) Let L be a subcomplex of K. Then the chain groups $C_p(L)$ are subgroups of $C_p(K)$. The quotient group $C_p(K)/C_p(L)$ is called the group of **relative chains of** K **modulo** L, denoted $C_p(K, L)$. Relative chain groups are free Abelian. Show that the boundary operator ∂ induces a homomorphism (still denoted ∂,

$$\partial_p : C_p(K, L) \to C_{p-1}(K, L).$$

We can then define the relative p-cycles and p-boundaries $Z_p(K, L), B_p(K, L)$ as before. The relative p-homology group is $Z_p(K, L)/B_p(K, L)$. ♢

Exercise 7.8: Give an algorithm to compute the Betti numbers of a simplicial complex K. ♢

End Exercises

§8. Effective Computation of Homology

It follows from the above considerations that Betti numbers of a simplicial complex K can be reduced to SNF computations.

It is best to illustrate the process with an example. Consider the triangulation of the 2-sphere (by a tetrahedron) in Figure 2.

We see that the maps ∂_1, ∂_2 can be represented by

$$\Lambda(\partial_1) = \begin{bmatrix} 1 \\ 2 \end{bmatrix} etc$$

Since SNF computation is expensive, and involves huge matrices for a large simplicial complex, we seek better methods. There is currently only a limited alterative, which we now present. This is an algorithm from Delfinado and Edelsbrunner [4] for computing Betti numbers of simplicial complexes in S^3.

The basic idea is to use incremental construction of a simplicial complex K. Suppose $\sigma \notin K$ and $K' = K \cup \{\sigma\}$ is also a simplicial complex. How does the Betti numbers of K' differ from that of K? Let us see this in the simple case where K' is 1-dimensional i.e., K' is just a graph (see Figure 7).

Example: Incremental construction of a graph. Let K be the graph in Figure 7(a). We see that $\beta_0(K) = 3$ (the number of connected components) and $\beta_1(K) = 1$ (number of independent cycles). Suppose we augment K with an edge $e = [14]$, as seen in Figure 7(b). How does this affect the first Betti number β_1? Recall that that $\beta_1(K')$ is the number of independent 1-cycles in K' minus the number of independent 1-boundaries. But the number of 1-boundaries is always 0 in a 1-dimensional complex (since there are no 2-simplices). Note that K' now has 3 1-cycles, viz., [12]+[23]+[31], [13]+[34]+[41] and [12]+[23]+[34]+[41]. But it is not hard to see that the number of **independent** 1-cycles in K' has increased by 1. Thus, we conclude that

$$\beta_1(K') = \beta_1(K) + 1.$$

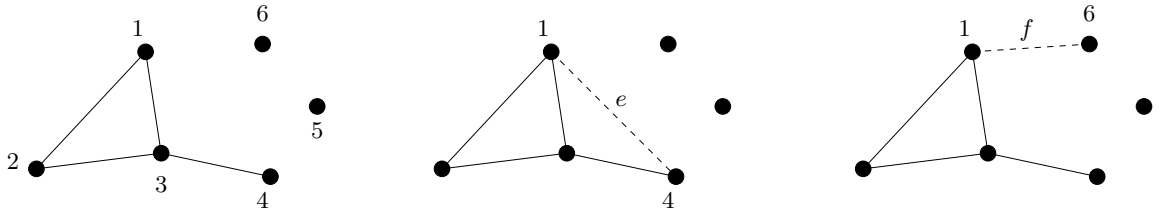

(a) K (b) $K' = K \cup \{e\}$ (c) $K' = K \cup \{f\}$

Figure 7: Augmenting an edge e or f to a graph K

Next, consider the simplex $K' = K \cup \{[16]\}$ shown in Figure 7c. It is clear that in this case, β_1 is unchanged, but β_0 has decreased:

$$\beta_0(K') = \beta_0(K) - 1.$$

It is also easy to see that the only changes to the Betti numbers in both situations are just the ones we described: nothing else changes. From now on, assume this tacit rule, *that Betti numbers of the augmented complex K' are the same as that of K unless we explicitly describe a change.* Of course, in proving correctness of our rules, we must also verify the correctness of this tacit rule.

Let us formalize the preceding discussion into a somewhat more abstract form

RULE$_1$: when augmenting K by a 1-simplex σ, if $\partial\sigma$ bounds in K, then $\beta_1(K') = \beta_1(K) + 1$, and otherwise $\beta_0(K') = \beta_0(K) - 1$.

We can also augment a graph K by adding an isolated vertex v. In this case, it is clear that $\beta_0(K') = \beta_0(K) + 1$. Moreover, $[v]$ bounds in K (since $\partial[v] = 0$, by definition). Hence this case also fits into the pattern of RULE$_1$. In general, for $d \geq 0$, we may formulate the rule:

RULE$_d$: when augmenting K by a d-simplex σ, if $\partial\sigma$ bounds in K, then $\beta_d(K') = \beta_d(K) + 1$, and otherwise $\beta_{d-1}(K') = \beta_{d-1}(K) - 1$.

Let us briefly see that this rule means for $d = 2$: for a triangle $[\sigma] = [u, v, w]$, $\partial[\sigma] = [v, w] - [u, w] + [u, v]$ is the boundary of some 2-chain in K. If so, this means that a new void has been created by adding σ. There is an alternative form of this rule: the criterion that "$\partial\sigma$ bounds in K" is equivalent to "σ is part of a d-cycle in K'".

Lemma 11. *Let K' be the augmentation of K by a d-simplex σ. Then the Betti numbers of K' is obtained from the Betti numbers of K by RULE$_d$ (and also the tacit rule).*

Proof. See Vegter/Rote. **Q.E.D.**

We now address the question of implementing these rules in an algorithm to compute $\beta_d(K)$.

Given a simplicial complex $K = \{\sigma_1, \ldots, \sigma_n\}$, we first fix an ordering of the simplices of K,

$$(\sigma_1, \sigma_2, \ldots, \sigma_n) \tag{4}$$

so that each of the sets $K_i = \{\sigma_1, \ldots, \sigma_i\}$ $(i = 1, \ldots, n)$ is a simplicial complex. Such a sequence (4) is called a **filter** of K. The corresponding sequence $K_1 \subseteq K_2 \subseteq \cdots \subseteq K_n$ is called a filtration. It is actually quite easy to find a filter of K: just list all the $(d-1)$ dimensional simplices before the d dimensional simplices for each $d \geq 1$.

Union-Find Datastructure. The main computational task in implementing our RULE above is to determine, for given K and d-simplex σ, whether $\partial\sigma$ bounds in K.

In case $d = 1$, the simplex σ is a directed edge $[u, v]$. Then $\partial(\sigma) = [v] - [u]$ bounds in K iff u, v belongs to the same connected component. This can be decided very efficiently by using a well-known data structure in Computer Science called the Union-Find data structure. In a certain (amortized) sense, each operation costs $O(\alpha(n))$ where n is the total number of vertices in the eventual complex, and $\alpha(n)$ is a very slow growing function called the inverse Ackermann function.

There is no known computational technique for $d = 2$. However, if the dimension of K is 3, then we can exploit duality: d-simplex in $\mathbb{R}^3$ is the same as a "dual $(3 - d)$-dimensional" simplex. In particular, 2-simplices will be dual 1-simplex. So we can use the Union-Find data structure in the dual setting.

Algorithm for Triangulation of $\mathbb{S}^3$. What is the dual of a complex K in $\mathbb{R}^3$? For simplicity, assume K is a of $\mathbb{S}^3$ (this is just $\mathbb{R}^3$ augmented with a single point at infinity). If K is not already a triangulation of $\mathbb{S}^3$, we can simply extend K into a triangulation L of $\mathbb{S}^3$. Moreover, we can choose a filter (4) for L such that some prefix of this filter is a filter of K. It is then clear that the Betti number computation we perform on L will yield the corresponding information for K.

We observe that in $\mathbb{S}^3$, if σ_i is a tetrahedron, then $\partial\sigma_i$ *bound in* K_i *iff* $i = n$ (the last tetrahedron). So it remains to figure out how to determine the rule when σ_i is a triangle.

Let the $\overline{K}_i := K \setminus K_i$ denote the complement of K_i. Thus we obtain a kind of reverse filtration,

$$\emptyset = \overline{K}_n \subseteq \overline{K}_{n-1} \subseteq \cdots \subseteq \overline{K}_0 = K.$$

Of course, $\overline{K}_i$ is not really a complex, but it can be viewed as an abstract complex G_i whose vertices V_i are the tetrahedrons in $\overline{K}_i$, and whose edges are pairs $\{a, b\} \subseteq V_i$ such that $a \cap b$ is an 2-simplex in $\overline{K}_i$, etc. We could consider triples $\{a, b, c\} \in V_i$ and so one, but in fact, we can ignore them for purposes. Thus, G_i is simply a graph with a vertex set and an edge set.

Let t_i be the transform K_{i-1} to K_i by adding a simplex σ_i. Then the reverse operation $\overline{t}_i$, amounts to adding a simplex from $\overline{K}_i$ to $\overline{K}_{i-1}$. If σ_i is a tet or a triangle, $\overline{t}_i$, this amounts to adding a vertex or an edge to G_i to obtain G_i.

Suppose σ_i is a triangle that belongs to a 2-cycle in K_i. This is equivalent to saying that G_i has one more component than G_{i-1}. Thus the transformation $\overline{t}_i : G_i \mapsto G_{i-1}$ results in the reduction of one component. It we maintain the components of G_i using the Union-Find datastructure, this means we perform a union operation. Whenever we perform such union corresponding to σ_i, we "mark" σ_i.

We are now ready to describe the overall algorithm for $\mathbb{S}^3$:

FORWARD PHASE: We iterate through a filter of K, maintaining the Union-Find data structure to maintain β_0 and β_1. We stop when we reach the 2-simplices.

BACKWARD PHASE: We then run the dual algorithm starting from $\overline{K}_n$ down to $\overline{K}_1$, but again stopping when we reach the 1-simplicies. In this phase, we just mark the 2-simplices as described above.

FINAL PHASE: Now, we continue from where the FORWARD PHASE got stopped. This time, we use the mark information to update β_2 and β_1.

See the original paper for a direct algorithm to compute the Betti numbers of an arbitrary triangulation of $\mathbb{R}^3$.

Exercises

Exercise 8.1: Carry out a complexity analysis of the above algorithm, paying careful attention to data structures. Conclude that $O(n\alpha(n))$ time and $O(n)$ space suffices if K has n simplices. ♢

End Exercises

§9. Euler Characteristic and Betti Numbers

The Euler Characteristic is an integer that can be associated to topological spaces; in fact, it can be computed as the alternating sum of the Betti numbers. See the interesting account of Imre Lakatos in "Proofs and Refutations", which traces historical development from the initial ideas of Euler to the algebraic view of Poincaré that is our modern viewpoint. But even Poincaré made a mistake and discovered torsion as a result. Lakatos' point (as a historian of science) is that these definitions are subject to various forces akin to negotiations. But I think it is a serious lapse to think that these negotiations are arbitrary and purely power play (as deconstructionists would have us believe). Massey [Chap.VI] also gives a brief historical background of homology theory.

We begin with the original intuitive facts about Euler Characteristic of a space. The initial observation from Euler is that for a planar triangulation of a simply-connected planar region R, the following invariant holds:

$$v - e + f = 2$$

where v, e, f is the number of vertices, edges, faces of the triangulation. It turns out that this number 2 does not depend on the choice of triangulation of R, – so we say the Euler characteristic for R is two, $\chi(R) = 2$. We then generalize this to solid polyhedral objects, and so on. Eventually, we obtain the formula

$$\chi(K) = \sum_{i=0}^{d} (-1)^d rank(C_i(K)).$$

This can be shown inductively. We can further relate this to the Betti numbers,

$$\chi(K) = \sum_{i=0}^{d} (-1)^d \beta_i(K).$$

[See Vegter-Rote].

A basic result of homology theory is that the Betti numbers $\beta_i(K)$ depends only on the topology[2] of the underlying space $|K|$, not on the particular triangulation. We can also show that $\beta_i(K)$ is a **homotopy invariant**: if $|L|$ is homotopic to $|K|$ then $\beta_i(L) = \beta_i(K)$ [See Vegter-Rote].

The interpretation of Betti numbers in $\mathbb{R}^3$ is quite interesting:

β_0 is the number of connected components.

β_1 is the number of **holes**. E.g., a donut has one hole, and an eye-glass frame (typically) has two holes.

β_2 is the number of **voids**. E.g., a soccer ball has one void (which is filled with air). Biological cells can be viewed as a medium filled with some fluid, with numerous voids containing a variety of material.

Here is an application: suppose we are given a model of a very complex molecule, regarded as the union of balls in $\mathbb{R}^3$ of various radii. Each ball corresponds to an atom (e.g., a hydrogen atom has a smaller radius than an oxygen atom). Since the biological functions of molecule often depends on the geometry of the molecules, there is interest in computing the number of holes and voids in such a molecute.

§10. Notes on Homotopy

Another way to get topological invariance is via homotopy. Again we algebraize the concept and discretize it to obtain the group analogue of homology groups, called **fundamental groups**. The relative advantages and disadvantages of using fundamental group invariants will become clear.

§11. Morse Theory

The standard introduction to Morse theory is the example of a torus $M \subseteq \mathbb{R}^3$, perhaps in an unusual position, standing upright as in Figure 8 [cf. Milnor]. We have a "height function" $h : M \to \mathbb{R}$ assigning a real value $h(x, y, z) = z$ to each point $(x, y, z) \in M$. Imagine a horizontal plane sweeping upward in time. Let $M_t := \{p \in M : h(p) \leq t\}$ denote the subset of M swept up to time t. We see that there are four critical moments: $t_0 < t_1 < t_2 < t_3$.

[2]By definition, topological properties of a space is defined up to homeomorphism.

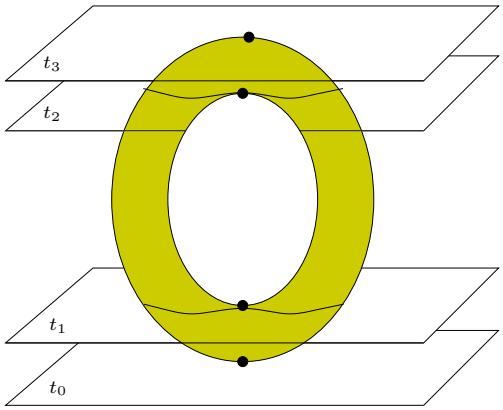

Figure 8: Torus in an upright position

1. $t < t_0$, when M_t is empty.
2. $t = t_0$, when M_t is a single point.
3. $t \in (t_0, t_1)$, when M_t is homeomorphic to a 2-cell.
4. $t = t_1$, when the boundary of the 2-cell is just pinched.
5. $t \in (t_1, t_2)$, when M_t is homeomorphic to a cylinder.
6. $t = t_2$, when the boundary components of the cylinder first meets.
7. $t \in (t_2, t_3)$, when M_t is homeomorphic to a torus with a puncture.
8. $t > t_3$, when M_t is a torus T^2.

In terms of homotopy types,

1. At time t_0, we attach a 0-cell to M_t.
2. At times t_1 and t_2, we attach a 1-cell to M_t,
3. At time t_3, we attach a 2-cell to M_t.

It turns out that the height function h has four critical points p_0, p_1, p_2, p_3, where p_i corresponds to the critical moment t_i. These points are, respectively, a minima, a saddle, a saddle and a maxima. Morse theory assigns a natural number ("index") to these critical points: in fact p_i has index d iff we attach a d-cell to M_t at time t_i. Thus, the study of the critical points and the critical values t_i of a smooth function h could reveal information about topological changes in M_t, and ultimately about M itself.

In general, Morse theory studies topological invariants of smooth manifolds $M \subseteq \mathbb{R}^n$ as revealed by studying critical points of smooth functions on M. Height functions are examples of **Morse functions** on M: such a function is defined to be a smooth function $h : M \to \mathbb{R}$ whose critical points are nondegenerate.

The first part of this lecture is basically concerned with making the concepts in this definition precise. We will need to generalize familiar concepts in Euclidean space to arbitrary manifolds.

The Vegter/Rote treatment has the merit of simplifying definitions by considering only manifolds which are embedded in Euclidean space.

It is helpful to keep two basic principles in mind throughout differential geometry (which provides the foundation for Morse theory): Local Principle [L] and Euclidean Principle [E]. For instance, a manifold is

locally [L] like an Euclidean space [E]. The Euclidean Principle says that we map all concepts in abstract spaces back to Euclidean space. For instance, to define the concept of differentiable function f on manifolds, we first define differentiability at a point [L], and reduce differentiability of f to the differentiability of a transformed function $\overline{f}$ on Euclidean neighborhoods [E]. The Local Principle has a corollary: nonlinear phenomenon when localized becomes a linear phenomenon. Thus the local linear transformations (Jacobians and tangent spaces, etc) becomes the key.

Smooth Manifolds. We say $f : \mathbb{R}^n \to \mathbb{R}$ is **smooth** if, for all $n \geq 0$, the nth derivative $f^{(n)}$ exists. A function $\varphi : \mathbb{R}^n \to \mathbb{R}^m$ is **smooth** if each component f_i of $\varphi(x) = (f_1(x), \ldots, f_n(x))$ is smooth.

We need to introduce the **differential** of φ at $q \in \mathbb{R}^n$: this is the linear map $d\varphi_q : \mathbb{R}^n \to \mathbb{R}^m$ such that for all $v \in \mathbb{R}^n$, and for all curves $\alpha_v : (-\varepsilon, \varepsilon) \to \mathbb{R}^n$ given by $\alpha_v(t) = \varphi(q + tv)$, we have $d\varphi_q(v) = \alpha_v'(0)$ (where α_v' denotes differentiating $\alpha_v(t)$ by t. Concretely, $d\varphi_q$ is a given by the $m \times n$ Jacobian matrix,

$$\begin{bmatrix} \frac{\partial f_1}{\partial x_1}(q) & \cdots & \frac{\partial f_1}{\partial x_n}(q) \\ \vdots & \ddots & \vdots \\ \frac{\partial f_m}{\partial x_1}(q) & \cdots & \frac{\partial f_m}{\partial x_n}(q) \end{bmatrix}$$

Note that $d\varphi_q$ is a constant matrix for each q. Alternatively, we can view $d\varphi$ as a map from $\mathbb{R}^n$ to linear transformations from $\mathbb{R}^n$ to $\mathbb{R}^m$.

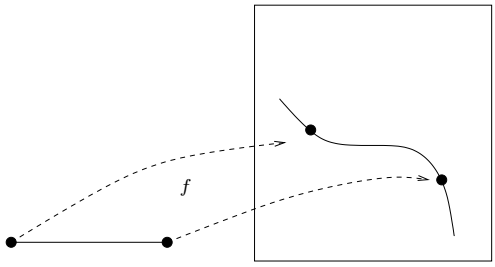

Figure 9: Smooth Curves (picture) and Manifolds

Next, we want to define what it means for $M \subseteq \mathbb{R}^n$ to be a smooth manifold. Sometimes, one says "differentiable" or "C^∞" instead of "smoothness". By a **diffeomorphism** $f : U \to V$ $(U, V \subseteq \mathbb{R}^n)$ we mean a smooth function whose inverse f^{-1} is defined and also a diffeomorphism. For instance $f : \mathbb{R} \to \mathbb{R}$ where $f(x) = x^3$ is a smooth homeomorphism function, but it is not a diffeomorphism because f^{-1} is not smooth at 0.

Let us first consider the special case: what is a smooth curve M embedded in $\mathbb{R}^2$? Applying the local principle [L], we say that the set $M \subseteq \mathbb{R}^2$ is a **smooth curve at** $p \in M$ if there exists an open interval $U \subseteq \mathbb{R}$, an open set $V \subseteq \mathbb{R}^2$ such that $p \in V$ and there exists a smooth function $\varphi : U \to \mathbb{R}^2$ such that (1) φ is a diffeomorphism from U to $V \cap S$, and (2) $d\varphi_q \neq \mathbf{0}$. We say M is a smooth curve if it is smooth at each $p \in M$.

More generally, $M \subseteq \mathbb{R}^n$ is a **smooth m-fold** if for all $p \in M$, there exists an open set $U \subseteq \mathbb{R}^m$, and open set $V \subseteq \mathbb{R}^n$, such that for some smooth onto homeomorphism $\varphi : U \to M \cap V$ such that $d\varphi_q : \mathbb{R}^m \to \mathbb{R}^n$ is injective. We call φ in this definition a **parametrization** or a **local coordinate system** or **chart** at (M, p). Very often, we require $\varphi(\mathbf{0}) = p$ and $\mathbf{0} \in U$.

Tangent space at a point of a manifold A **tangent vector** of M at point p is $\alpha'(0)$ where $\alpha(t)$ is some smooth curve $\alpha : (-\varepsilon, \varepsilon) \to M$ such that $\alpha(0) = p$. The **tangent space** T_qM is the pair (V, p) where V is the set of all tangent vectors of M at p. We call $V + p$ the affine tangent space of M at p, and this space passes through p.

If $\varphi : U \to M$ is a smooth parametrization of M at p, $\mathcal{O} \in U$ and $\varphi(\mathbf{0}) = p$, then $T_pM = d\varphi_{\mathbf{0}}(\mathbb{R}^m) \subseteq \mathbb{R}^n$. Note that T_pM is m-dimensional like M, and it passes through p by definition.

If $\varphi : U \to M$ is a chart of M at p, $0 \in U$, $\varphi(0) = p$, then $T_pM = \varphi_0(\mathbb{R}^m) \subseteq \mathbb{R}^n$.

Topological manifolds. We generalize the above definitions by beginning with a topological space M that is[a] Hausdorff. A **chart** of M is an onto homeomorphism $h : U \to V$ where $U \subseteq \mathbb{R}^n$ and $V \subseteq M$ are open sets. (Note: it clearly does not matter whether we use h or h^{-1} as the definition of a chart, as long as we are consistent.) An **atlas** of M is a collection $\{h_\alpha\}_\alpha$ of charts of M where $h_\alpha : U_\alpha \to V_\alpha$ ($U_\alpha \subseteq \mathbb{R}^n$, $V_\alpha \subseteq M$) such that $\cup_\alpha V_\alpha = M$. For charts h_α, h_β in an atlas, let $U_{\alpha\beta} := U_\alpha \cap U_\beta$. If $U_{\alpha\beta} \neq \emptyset$, then define the **chart transformation**

$$h_{\alpha\beta} : h_\alpha(U_{\alpha\beta}) \to V_\beta$$

where $h_{\alpha\beta} = h_\beta \circ h_\alpha^{-1}$. Since $h_{\alpha\beta}$ is a function on Euclidean subsets, $h_\alpha(U_{\alpha\beta}) \subseteq V_\alpha \subseteq \mathbb{R}^n$ and $h_\beta(U_{\alpha\beta}) \subseteq V_\beta \subseteq \mathbb{R}^n$, we can speak of smoothness of such functions [Principle E]. We say that M is **smooth** if it has an atlas whose chart transformations are smooth functions.
REMARK: The Vegter/Rote treatment affords us to skip chart transformations.
Next, let f be a map between two smooth manifolds, $f : M \to N$. We say that f is smooth at a point $p \in M$ if there are smooth atlases with two charts, $h : (0, U) \to (p, V)$ and $k : (0, U) \to (f(p), V')$ ($U \subseteq \mathbb{R}^n, V \subseteq M, V' \subseteq N$) such that $k \circ f \circ h^{-1} : h^{-1}(U) \to V'$ is smooth.

[a] Hausdorff or T_2 means that for all $p \neq q \in M$, there exists neighborhoods N_p and N_q such that $N_p \cap N_q = \emptyset$.

Morse functions on 2-manifolds are of three types: maxima, minima and saddle points. Let us look at the simplest type of critical point that is degenerate, the **monkey saddle**. Consider the function $h : \mathbb{R}^2 \to \mathbb{R}$. where $h(x, y) = x^3 - 3xy^2$.

Exercises

Exercise 11.1: How would you perturb the Monkey saddle function $h(x, y) = x^3 - 3xy^2$ so that it becomes Morse? ♢

Exercise 11.2: Recall the standard torus T^2 used in Morse theory. One chart for T^2 has been given by Vegter/Rote: it is $\varphi : U \to T^2$ where $U = (0, 2\pi) \times (0, 2\pi)$, and

$$\varphi(u, v) = (r \sin u, (R - r \cos u) \sin v, (R - r \cos u) \cos v).$$

(a) Give other charts so as to cover the rest of T^2. How many additional charts do you need?
(b) Let h be the usual height function on T^2. Compute the gradient field of h.
(c) Consider another different height function $f(\varphi(u, v)) = r \sin u$. Compute the gradient field of f.
(d) Is f a Morse function? ♢

Exercise 11.3: Let $F(X, Y) \in \mathbb{Z}[X, Y]$ and consider the curve $M : F(X, Y) = c$ for some integer $c \in \mathbb{Q}$.
(a) Describe how you can detect whether M is a smooth manifold.
(b) Let M be smooth from (a). For $p_0 \in \mathbb{Q}^2$, define the function $f : M \to \mathbb{R}$ where $f(q) = \|p_0 - q\|$ (Euclidean distance). Describe how to test whether f is Morse.
(c) Let f be Morse from (c). Describe how to compute the critical points of f and to determine the index of each critical point. ♢

End Exercises

§12. Notes on Numerical and Algebraic Methods

We now address the problem of convert continuous data into discrete data. For instance, given a Morse function, how do we determine its critical points? Given an algebraic surface, how do we compute a topologically correct polygonal mesh representation? There are two distinct set of techniques here: numerical and algebraic.

§12.1. UFD and GCD

Let $\mathbb{D}$ be a domain, i.e., a ring with no zero-divisors. The **units** in $\mathbb{D}$ are the invertible elements of $\mathbb{D}$. For $\mathbb{D} = \mathbb{Z}$, there are just two units, ± 1. For a field $\mathbb{D}$ then every non-zero element is a unit. Two elements are **associates** if they are equal to each other up to multiplication by a unit. In $\mathbb{Z}$, the associates comes in pairs, n and $-n$. In a field, every non-zero is an associate of each other. An element x in $\mathbb{D}$ is **irreducible** if x is divisible only by units or its associates. We are exclusively interested in computation over a UFD, where the fundamental theorem of arithmetic holds: every non-unit can be expressed as a power product of irreducible elements, and this is unique up to associates.

In a UFD, the concept of a GCD is well-defined, but up to associates. To make GCD a unique function, we choose a distinguished member of each equivalence class of associates: e.g., in $\mathbb{Z}$, we choose the positive member of each pair $n, -n$ of associates. Then GCD returns the distinguished member. There are well-known algorithms (Euclid's and extensions) for computing GCD in the case $\mathbb{D} = \mathbb{Z}$ and $\mathbb{D} = F[X]$ where F is a field. Gauss's lemma allows us to extend this to the multivariate domain $\mathbb{D}[X_1, \ldots, X_n]$.

Consider GCD in $\mathbb{D}[X]$. The **content** of $A \in \mathbb{D}[X]$ is the GCD of the coefficients of A. We say $A \in \mathbb{D}[X]$ is **primitive** if its content is 1. Then a primitive factorization of A is the factorizatio $A = bB$ where B is primitive, called the **primitive part** of A.

In $\mathbb{D}$, we generally distinguish elements up to non-associates: this means that associates are equal for all practical purposes. We can generalize this: suppose $A, B \in \mathbb{D}[X]$. Then we say A, B are **similar** if $\alpha A = \beta B$ for some non-zero $\alpha, \beta \in \mathbb{D}$. Notice that this is equivalent to saying A and B have the same primitive part. For univariate polynomials, we distinguish them up to non-similarity.

Let us generalize this: if $A, B \in \mathbb{D}[X, Y]$, we say A, B are **similar** if $\alpha A = \beta B$ for some non-zero $\alpha, \beta \in \mathbb{D}[X] \cup \mathbb{D}[Y]$. The 0-content of A is the GCD of its coefficients in $\mathbb{D}$. The X-content of A is the content of A viewed as a polynomial in Y. Similarly for the Y-content of A. Then the **content** of A is the product of its 0-, X- and Y-contents. The **primitive part** of A is given by A divided by its content. We say A is **primitive** its content is 1.

We say $A \in \mathbb{D}[X, Y]$ is **reduced** if $\mathtt{GCD}(A, A_X) = \mathtt{GCD}(A, A_Y) = 1$.

When using polynomials $A \in \mathbb{D}[X, Y]$ to define curves, we are only interested in reduced primitive polynomials.

§12.2. Resultants

Perhaps the most fundamental algebraic tool in this area is the theory of resultants. The multivariate theory of resultants is a current topic of great interest. The basic problem is this: suppose we are given two polynomials $p, q \in \mathbb{D}[X]$ where $\mathbb{D}$ is any UFD. We want to know if they have any common zero. This is equivalent to $\mathtt{GCD}(p, q)$ have positive degree. Generically, we know that $\deg(\mathtt{GCD}(p, q)) = 0$. Hence some very special coincidences have to occur in order that $\deg(\mathtt{GCD}(p, q)) > 0$. This coincidence can be expressed as the vanishing of a polynomial in the coefficients of p, q. A polynomial R in the coefficients of p, q is called the resultant of p and q if the vanishing of R is a necessary and sufficient for $p = q = 0$.

This concept of genericity can be generalized. First of all, a univariate polynomial p has zeros in the generic case (in the real case, this is a nontrivial conclusion), but two univariate polynomials sharing common zero is non-generic. Similarly, suppose $p, q, r \in \mathbb{D}[X, Y]$, then $p = q = 0$ have a solution is a generic fact, but $p = q = r = 0$ having a solution is not a generic fact. We again non-genericity can be expressed as the vanishing of a polynomial R in the coefficients of p, q, r. In general, if the resultant is a set of polynomials (called a **resultant system**).

Fix any UFD $\mathbb{D}$. Assume polynomials in this section have coefficients in $\mathbb{D}$.

Let $A, B \in \mathbb{D}[X]$. If $A = \sum_{i=0}^{m} a_i X^i$ and $B = \sum_{j=0}^{n} b_j X^j$, with $a_m b_n \neq 0$, then the **Sylvester Matrix** of A, B is defined as the following square $(m+n) \times (m+n)$ matrix

$$Syl(A,B) \quad = \quad \begin{bmatrix} a_m & a_{m-1} & \cdots & & a_0 & & & \\ & a_m & a_{m-1} & \cdots & & a_0 & & \\ & & \ddots & & & & \ddots & \\ & & & a_m & a_{m-1} & \cdots & & a_0 \\ b_n & b_{n-1} & \cdots & b_1 & b_0 & & & \\ & b_n & b_{n-1} & \cdots & b_1 & b_0 & & \\ & & \ddots & & & & \ddots & \\ & & & b_n & b_{n-1} & \cdots & & b_0 \end{bmatrix}.$$

Notice the main diagonal elements in this matrix, comprising n copies of a_m and m copies of b_0. The determinant of $Syl(A,B)$ is an element of $\mathbb{D}$. It is denoted $\texttt{res}(A,B)$ and called the **resultant** of A and B.

We also say $\texttt{res}(A,B)$ is the result of **eliminating** the variable X from A, B. Thus, resultants gives us a tool (analogous to Gaussian elimination in the case of linear equations) for eliminating variables. This interpretation will be very important later. To apply this, suppose our polynomials are elements of $\mathbb{Q}[X_1, X_2, \ldots, X_n]$, and suppose we wish to eliminate X_n. We can then take $\mathbb{D}$ to be $\mathbb{Q}[X_1, \ldots, X_{n-1}]$ and X to be X_n.

Theorem 12. *$\texttt{GCD}(A,B)$ is not a constant iff $\texttt{res}(A,B) = 0$.*

Proof. See [Yap,Lemma 6.13(p.156)]. In sketch, assume $\deg A = m, \deg B = n$ with Sylvester matrix $S_0 = Syl(A,B)$. Thus $\texttt{res}(A,B) = \det(S_0)$. Suppose $U = \sum_{i=0}^{n-1} u_i X^i, V = \sum_{i=0}^{m-1} v_i X^i$ are polynomials of degrees $\leq n-1$ and $\leq m-1$ (resp.). Let $\underline{U} = (u_{n-1}, \ldots, u_0)$ be the corresponding row vectors of length n. Similarly for $\underline{V}$. Let $x = (X^{m+n-1}, X^{m+n-2}, \ldots, X, 1)^T$. Then $(\underline{U}, \underline{V}) \cdot S_0 \cdot x = UA + VB$. We see that $UA + UB = 0$ has a solution iff $\det(S_0) = 0$.

The theorem holds if we show $UA + UB = 0$ (with $\deg U \leq n-1, \deg V \leq m-1$ iff $\texttt{GCD}(A,B)$ is not a constant. If $UA + VB = 0$ then $A | VB$. Thus $\deg(\texttt{GCD}(A,V)) + \deg(\texttt{GCD}(A,B) = m$. Since $\deg(\texttt{GCD}(A,V)) \leq \deg(V) \leq m-1$, we conclude that $\deg(\texttt{GCD}(A,B)) \geq 1$, as desired. Conversely, if $g = \texttt{GCD}(A,B)$ is non-constant, then we can choose $U = B/g$ and $V = A/g$ to satisfy the equation $UA + VB = 0$. **Q.E.D.**

The following might be called the "fundamental lemma of resultants": let $A, B \in \mathbb{D}[X]$ have degrees m and n. Assume α_i $(i = 1, \ldots, m)$ and β_j $(j = 1, \ldots, n)$ are the zeros of A, B in the algebraic closure of $\mathbb{D}$.

Theorem 13. *For $A, B \in \mathbb{D}[X]$, we have*

$$\texttt{res}_X(A,B) = a^n \prod_{i=1}^{m} B(\alpha_i)$$

where a is the leading coefficient of A.

For instance, if $B(X)$ has degree $n = 2$ then a direct computation shows that $\texttt{res}_X(aX - \alpha, B(X)) = a^2 B(\alpha)$. For the general proof of this theorem, see [7]. It is then easy to deduce the following:

Theorem 14.
(i) The zeros of $\texttt{res}_Y(A(Y), B(X \mp Y))$ are $\alpha_i \pm \beta_j$ (for all $i = 1, \ldots, m, \quad j = 1, \ldots, n$).
(ii) The zeros of $\texttt{res}_Y(A(Y), Y^n B(X/Y))$ are $\alpha_i \beta_j$ (for all $i = 1, \ldots, m, \quad j = 1, \ldots, n$).
(iii) The zeros of $\texttt{res}_Y(A(Y), X^n B(Y/X))$ are α_i / β_j (for all $i = 1, \ldots, m, \quad j = 1, \ldots, n$). It is assumed that each $\beta_j \neq 0$.

Exercises

Exercise 12.1: Prove the special case of Theorem 13: $\texttt{res}_X(A, B) = a^n \prod_{i=1}^m B(\alpha_i)$. ♢

Exercise 12.2: Compute $R(X) = \texttt{res}_Y(A(Y), B(XY))$. In particular, express the leading coefficient of $R(X)$ in terms of the leading and constant coefficients to A, B. ♢

End Exercises

§12.3. Root Separation Bounds

We prove a fundamental result about how close two algebraic numbers can be to each other. The basic result is from Mahler and depends on the discriminant of a polynomial. For $A(X) \in \mathbb{D}[X]$ of degree m, with leading coefficient a, define its discriminant as

$$\text{disc}(A) = a^{-1}\texttt{res}(A, A'). \tag{5}$$

where A' denotes the derivative. Note that $\texttt{res}(A, A')$ is divisible by a since the first column of the Sylvester matrix is a multiple of a. Hence $\text{disc}(A) \in \mathbb{D}$.

Next, if we write $A = a\prod_{i=1}^m (X - \alpha_i)$ where α_i are all the zeros of A in the algebraic closure of $\mathbb{D}$, then it can be shown that

$$\text{disc}(A) = a^{2m-2)} \prod_{1 \le i < j \le m} (\alpha_i - \alpha_j)^2. \tag{6}$$

This proves that if A has a multiple zero, then $\text{disc}(A) = 0$.

Another useful expression for the discriminant is in terms of a Vandermonde matrix:

$$\sqrt{|\text{disc}(A)|} = \pm a^{m-1} \det \begin{bmatrix} 1 & 1 & \cdots & 1 \\ \alpha_1 & \alpha_2 & \cdots & \alpha_m \\ \alpha_1^2 & \alpha_2^2 & \cdots & \alpha_m^2 \\ \vdots & \vdots & \ddots & \vdots \\ \alpha_1^{m-1} & \alpha_2^{m-1} & \cdots & \alpha_m^{m-1} \end{bmatrix}. \tag{7}$$

For any complex polynomial $A(X) = a\prod_{i=1}^m (X - \alpha_i)$ $(a, \alpha_i \in \mathbb{C})$ define its **Mahler measure** to be $M(A) := |a| \prod_{i=1}^m \max\{1, |\alpha_i|\}$.

Theorem 15. *Let $A(X)$ be a complex polynomial of degree m. If α, β are distinct roots of A, then $|\alpha - \beta| \ge \sqrt{|\text{disc}(A)|} M(A)^{-1} m^{-m/2}$.*

Proof. Omitted. See Yap's book. **Q.E.D.**

§12.4. Real Root Isolation

An interval I is called an **isolating interval** of a polynomial P if there is a unique root of P in I. The **real root isolation** problem is this: given P and I, compute a set of pairwise disjoint isolating intervals for P, one interval for each root of P in I.

When P is a real polynomial, Sturm's sequences can be used to determine the the number of distinct real zeros in P. Such sequences can be effectively computed if we have algorithms to perform rational operations on the coefficients of P. Thus, we can solve the real root isolation problem by bisection: keep subdividing the input interval in half until we verify that we have obtained an isolating interval. In case the midpoint of the bisection itself is a root, we can either use root bounds or some other means for handling the problem.

In recent years, it has been observed that it is usually more efficient method to use an alternative method based on **Descartes' rule of sign**. Let $V(P)$ denote the number of sign changes in the sequence $(c_0, \ldots, c_k)$ of coefficients of $P(X) = \sum_{i=0}^k c_i X^i$ (we remove the zero coefficients before counting the sign changes). Descartes' rule says that *the number of positive zeros of $P(X)$ is at most $V(P)$, and moreover, this number*

is less than $V(P)$ be an even number. In particular, if $V(P) = 0$ or $V(P) = 1$, the rule gives us the correct number of positive zeros of P. For instance, if $P(X) = X^{21} - 4X^{15} - X^8 + 5X^3 - 1$ then $V(P) = 3$; and hence P has at most 3 positive roots.

Suppose $\deg P = k$, and let $K = [\frac{a}{c}, \frac{b}{d}]$ be an interval where $a, b, c, d \in \mathbb{R}$, $c, d > 0$. Consider the **linear fractional transformation** (also called **Möbius transformation**) $M_K(X) = \frac{bX+a}{dX+c}$. If we allow arbitrary real values for a, b, c, d, then we can choose $c = d = 1$, and thus have $M_K(X) = \frac{bX+a}{X+1}$. But in general, we like a, b, c, d to be integers. Further, it is best if c, d are powers of 2 (reflecting the fact that the endpoints of K are bigfloats). Clearly, M_K maps reals to reals. But more can be said:

$$M_K(0) = \frac{a}{c}, \qquad M_K(\infty) = \frac{b}{d}.$$

Moreover, for $0 < s < t < \infty$, we have

$$\frac{a}{c} < M_K(s) < M_K(t) < \frac{b}{d}.$$

In other words, M maps the interval $[a/c, b/d]$ continuously into $[0, \infty]$.

For any real polynomial $P(X)$ and any Möbius transformation M, define the transformed polynomial

$$P_M(X) := (cX + d)^k P(M(X)).$$

We may also write $P_K(X)$ for $P_M(X)$ if $M = M_K$. Then we see that $\alpha \in \mathbb{C}$ is a zero of $P(X)$ iff $M^{-1}(\alpha)$ is a zero of $P_M(X)$. In particular, the zeros α of P in K are mapped bijectively into the zeros $M^{-1}(\alpha)$ of P_M in $[0, \infty]$. Thus, Descartes' rule of sign can be applied to $P_M(X)$ to estimate the number of roots of P in K.

In order for this to yield an algorithm for bisection, we need to be assured that when K is small enough, $Var(P_K)$ would be 0 or 1. This is true provided P is square-free. The relevant lemma is called the 1-circle and 2-circle lemmas:

Lemma 16 (One Circle). *Let $P \in \mathbb{R}[X]$ and $K = [a, b] \subseteq \mathbb{R}$. If D_K is the disc with diameter K contains no zero of $P(X)$, then $V(P_K) = 0$.*

Proof. Let M be the Möbius transformation M_K. We check that M^{-1} maps D_K to the half-space to the right of the imaginary (vertical) axis of the complex plane. Hence P_K have factors of the form $X + \alpha$ or $X^2 + 2\alpha + \beta$ where α, β are real positive numbers. Clearly, a product of such factors has no sign variation. **Q.E.D.**

Lemma 17 (Two Circle). *Let $P \in \mathbb{R}[X]$ and $K = [a, b] \subseteq \mathbb{R}$. Let $\underline{T}_K$ and $\overline{T}_K$ be the equilateral triangles, each with K as one of its side. If $\underline{D}_K, \overline{D}_K$ are the two discs whose boundary circumscribe these triangles. If $\underline{D}_K \cup \overline{D}_K$ contain exactly one real zero of $P(X)$ then $V(P_K) = 1$.*

Proof sketch: our proof depends on Ostrowski's theory of normal polynomials: if a_i are the coefficients of a normal polynomial $Q(X) = \sum_i a_i X^i$, then $a_i^2 > a_{i-1}a_{i+1}$ for each i. Thus $\frac{a_i}{a_{i+1}} > \frac{a_{i-1}}{a_i}$ for all i. We note that M_K^{-1} maps $\underline{D}_K \cup \overline{D}_K$ to a cone region (see Figure). Moreover, if $Q(X)$ is normal, then it is easily seen that $V(Q(X)(X - \alpha)) = 1$ where $\alpha > 0$: That is because

$$\begin{aligned} Q(X)(X - \alpha) &= \sum_i (a_{i+1} - \alpha a_i) X^i \\ &= \sum_i a_i (\frac{a_{i-1}}{a_i} - \alpha) X^i \\ &= \sum_i a_i b_i X^i. \end{aligned}$$

And we can see that the sequence of b_i's change sign exactly once. This shows that $V(Q) = 1$.

Exercises

Exercise 12.3: Implement in Core Library Descartes' algorithm to search for real zeros in an interval $[a, b]$. ♢

End Exercises

§12.5. Bernstein Polynomials

Bernstein polynomials gives us a different view of the Descartes' method. Moreover, it has some important advantages, and is useful when generalized to higher dimensions.

Let $n \geq 1$ and $i = 0, \ldots, n$. Also let $a < b$ be real numbers. A Bernstein basis is

$$B_i^n(X; a, b) = \binom{n}{i} \frac{(X-a)^i (b-X)^{n-i}}{(b-a)^n}.$$

When $(a, b) = (0, 1)$, then we simply write $B_i^n(X)$. It is well known that every polynomial $P(X)$ of degree n can be written as

$$P(X) = \sum_{i=0}^{n} c_i B_i^n(X; a, b)$$

where $c_i \in \mathbb{R}$ are called the **Bernstein coefficients** of $P(X)$.

§12.6. Isolating Interval Representation

We can use isolating intervals to represent real algebraic numbers. The **isolating interval representation** of a real algebraic number α is a pair $(A(X), I)$ such that $A(\alpha) = 0$, $A(X)$ is a square-free integer polynomial, and I is an isolating interval of A containing α. We write $\alpha \simeq (A, I)$ in this case. We now discuss arithmetic operations on such representations.

Note that in this representation, we could easily determine if a given (A, I) represents $\alpha = 0$.

Given $\alpha \simeq (A, I)$ and $\beta \simeq (B, J)$, we want to compute $\alpha + \beta$. First compute $C = \mathtt{res}_Y(A(Y), B(X-Y))$. By Theorem 14(i), we know that $\alpha+\beta$ is a root of C. First, we replace C by its square-free part, $C/\mathtt{GCD}(C, C')$ where C' is dC/dX. Next, we compute $K = I + J$ (using interval arithmetic) and check whether the sign variation $V(C_K(X))$ is equal to 1. If so, (C, K) is our desired representation of $\alpha + \beta$. If not, we successively refine I and J, and repeat this test. We will eventually succeed, for the same reason that our root isolation algorithm halts.

Similarly, using the other resultants described in Theorem 14, we can subtract, multiply and divide such representations.

We also need to determine the sign of a given $\alpha \simeq (A, I)$ But this is easy if $0 \notin I$. Otherwise, if $I = [a, b]$, we just check whether $[a, 0]$ or $[0, b]$ is an isolating interval of A. This means that we can compare two real algebraic numbers α, β by determining the sign of $\alpha - \beta$.

Exercises

Exercise 12.4: (a) Implement in Core Library the indicated algorithms for performing the four arithmetic operations.
(b) Implement comparison of real algebraic numbers.
(c) Implement square-root. ♢

Exercise 12.5: Suppose $A(X)$ is a polynomial whose coefficients are real algebraic numbers represented by isolating intervals. How can you isolate all the zeros of $A(X)$. ♢

End Exercises

§12.7. Subresultants

Let S_0 be the Sylvester matrix for $A, B \in \mathbb{D}[X]$ of degrees m and n ($m \geq n \geq 1$). For $i = 0, \ldots, n$, we now define the matrix S_i obtained by deleting the last i rows of coefficients of A, and the last i rows of coefficients of B, and the last i columns of the result. Thus S_i has shape $(m+n-2i) \times (m+n-i)$. Further, let S_i' denote the square matrix obtained by deleting the last i columns of S_i. We call $\det(S_i)$ the **principal subresultant coefficient** of A, B, denoted $\mathrm{psc}_i(A, B)$. Then, $\mathrm{psc}_0(A, B)$ is just the resultant of A, B.

Theorem 12 can be generalized:

Theorem 18. $\deg(\mathtt{GCD}(A, B)) \geq k$ *iff* $\mathrm{psc}_i(A, B) = 0$ *for* $i = 0, \ldots, k$.

Proof. The base case is Theorem 12. Use induction on k. **Q.E.D.**

§12.8. Homogeneous Polynomials

Now assume $A, B \in D'[X_1]$ where $\mathbb{D}' = \mathbb{D}[X_2, \ldots, X_r]$, $r \geq 1$, and $\mathbb{D}$ is a UFD. Thus $A, B \in \mathbb{D}[X_1, \ldots, X_r]$ are multivariate polynomials. Let $\mathtt{res}_{X_1}(A, B) \in \mathbb{D}[X_2, \ldots, X_r]$ denote the resultant where A, B are viewed as polynomials in X_1. To avoid double subscripts, we write $\mathtt{res}_1(A, B)$ instead of $\mathtt{res}_{X_1}(A, B)$.

More generally, define $\mathtt{res}_i(A, B)$ (for $i = 1, \ldots, r$) as the resultant with respect to X_i. Alternatively, the subscript i or X_i says we are "eliminating X_i" from the system of equations $A = B = 0$. In examples, we usually write X for X_1, Y for X_2 and Z for X_3. Now, when we speak of the "degree" of $A \in \mathbb{D}[X_1, \ldots, X_r]$, we mean its **total degree** in $X_1, \ldots, X_r$, still denoted $\deg(A)$. Also, let $\deg_i(A)$ denote the degree of A as a polynomial in X_i. For example, if $A = X^3Y - XY + 2$ then $\deg(A) = 4$, $\deg_1(A) = 3, \deg_2(Y) = 1$ (assuming $X = X_1$ and $Y = X_2$).

A multivariate polynomial can be written as

$$A = \sum_{e \in I} c_e X^e, \qquad c_e \in \mathbb{D}$$

where $I \subseteq \mathbb{N}^r$ is a finite set and $X^e = X_1^{e_1} \cdots X_r^{e_r}$ where $e = (e_1, \ldots, e_r)$. If $c_e \neq 0$, then we say X^e **occurs** in A and call $c_e X^e$ a **term** of A. The polynomial A is X_i**-regular** if $X_i^{\deg(A)}$ occurs in A. We simply say "regular" for X_1-regular.

Let $|e| = e_1 + \cdots + e_r$. So $\deg(A) = \max_{e \in I} |e|$. We say A is **homogeneous** if $|e| = |f|$ for all $e, f \in I$. The zero polynomial is, by definition, homogeneous and has degree $-\infty$.

The following property can be used as the definition of homogeneous polynomials: let $A \in \mathbb{D}[X_1, \ldots, X_r]$ have degree m. Then A is homogeneous iff for all nonzero $t \in \mathbb{D}$,

$$A(tX_1, \ldots, tX_r) = t^m A(X_1, \ldots, X_r). \tag{8}$$

Another property is Euler's identity:

$$mA(X_1, \ldots, X_r) = \sum_{i=1}^{r} X_i \frac{\partial A}{\partial X_i}. \tag{9}$$

If $A \in \mathbb{D}[X_1, \ldots, X_r]$, let $\widehat{A} \in \mathbb{D}[X_1, \ldots, X_r, W]$ denote the standard homogenization of A using a new variable W, with the property that $\deg(A) = \deg(\widehat{A})$ and $\widehat{A}(X_1, \ldots, X_r, 1) = A(X_1, \ldots, X_r)$. Also, for $B \in \mathbb{D}[X_1, \ldots, X_r, W]$, let $B^\vee$ denote the operation of substituting $W = 1$ in B. Thus, $\widehat{A}^\vee = A$.

Lemma 19. *If A, B be arbitrary polynomials (not necessarily homogeneous).*
(i) $\widehat{AB} = \widehat{A}\widehat{B}$
(ii) $\widehat{A+B} = \widehat{A} + \widehat{B}X_0^{\deg(A)-\deg(B)}$ *(where* $\deg(A) \geq \deg(B)$*)*
(iii) $\mathtt{GCD}(\widehat{A}, \widehat{B}) = \widehat{\mathtt{GCD}(A, B)}$.

We are ready to prove the main result of this section.

Theorem 20. *Let A, B be regular homogeneous polynomials in $r \geq 2$ variables, of degrees m and n respectively. If $\mathtt{res}_1(A,B) \neq 0$ then $\mathtt{res}_1(A,B)$ is homogeneous of degree mn.*

Proof. Write $A = \sum_{i=0}^{m} a_i X^i$ and $B = \sum_{j=0}^{n} b_i X^i$ where $X = X_1$ and $a_i, b_j \in \mathbb{D}[V]$ where, for simplicity, we write V for $(X_2, \ldots, X_r)$. By homogeneity, either $a_i = 0$ or $\deg(a_i) = m-i$ for all $i = 0, \ldots, m$. Similarly, either $b_j = 0$ or $\deg(b_j) = n-j$ for $j = 0, \ldots, n$. Let $R(V) = \mathtt{res}_1(A,B)$. For any t, let $tV = (tX_2, \ldots, tX_r)$. From (8), we conclude that

$$R(tV) \quad = \quad \begin{bmatrix} a_m & ta_{m-1} & \cdots & & t^m a_0 & & & \\ & a_m & ta_{m-1} & \cdots & & t^m a_0 & & \\ & & \ddots & & & & \ddots & \\ & & & a_m & ta_{m-1} & \cdots & & t^m a_0 \\ b_n & tb_{n-1} & \cdots & t^{n-1}b_1 & t^n b_0 & & & \\ & b_n & tb_{n-1} & \cdots & t^{n-1}b_1 & t^n b_0 & & \\ & & \ddots & & & & \ddots & \\ & & & b_n & tb_{n-1} & \cdots & & t^n b_0 \end{bmatrix}.$$

If we next multiply the ith row of A's by t^{i-1} and the jth row of the B's by t^{j-1}, we obtain

$$t^p R(tV) = \begin{bmatrix} a_m & ta_{m-1} & \cdots & & t^m a_0 & & & \\ & ta_m & t^2 a_{m-1} & \cdots & & t^{m+1} a_0 & & \\ & & \ddots & & & & \ddots & \\ & & & t^{n-1}a_m & t^n a_{m-1} & \cdots & & t^{m+n-1} a_0 \\ b_n & tb_{n-1} & \cdots & t^{n-1}b_1 & t^n b_0 & & & \\ & tb_n & t^2 b_{n-1} & \cdots & t^n b_1 & t^{n+1} b_0 & & \\ & & \ddots & & & & \ddots & \\ & & & t^{m-1}b_n & t^{m-2} b_{n-1} & \cdots & & t^{m+n-1} b_0 \end{bmatrix}.$$

where $p = \binom{m}{2} + \binom{n}{2}$. But in the righthand side determinant, we can extract a factor of t^{i-1} from the ith column (for $i = 1, \ldots, m+n$). Hence the righthand side determinant is equal to

$$t^{\binom{m+n}{2}} R(V).$$

This proves that

$$t^p R(tV) = t^{\binom{m+n}{2}} R(V).$$

Hence $R(V)$ is homogeneous and its degree is

$$\binom{m+n}{2} - \binom{m}{2} - \binom{n}{2} = mn.$$

Instead of a direct calculation, the reader may instantly see the truth of this last equation in terms of its combinatorial interpretation. **Q.E.D.**

Corollary 21. *Suppose $A, B \in \mathbb{D}[X_1, \ldots, X_r]$ are regular polynomials, not necessarily homogeneous. Then $\mathtt{res}_1(A,B)$, if non-zero, has degree $\leq mn$ in $X_2, \ldots, X_r$.*

Proof. First, we homogenize A and B to $\widehat{A}$ and $\widehat{B}$, using a new variable X_0. Our theorem says that $R(X_0, X_1, \ldots, X_r) = \mathtt{res}_1(\widehat{A}, \widehat{B})$ is homogeneous of degree mn. Since $\mathtt{res}_1(A, B)$ is equal to $R(1, X_1, \ldots, X_r)$, the degree of $\mathtt{res}_1(A, B)$ is at most $\deg(R(X_0, X_1, \ldots, X_r)) = mn$. **Q.E.D.**

Suppose the polynomial $A = \sum_{i=0}^{m} a_i X^i$ ($a_i \in \mathbb{D}[X_2, \ldots, X_r]$) is not regular, i.e., $a_0 \notin \mathbb{D}$. What can we do? One possibility is make divide A by a_0. The coefficients of A/a_0 can now be viewed as elements of a meromorphic series (i.e., power series with finitely many terms with negative powers). In essense, this is Newton's approach. See [1, Lect. 9] for this development. A simpler approach is to consider the following transformation:

$$X_i \mapsto Y_i + c_i X, \qquad (i = 2, \ldots, r).$$

The polynomial $A'(X, Y_2, \ldots, Y_r) = A(X, Y_2 + c_2 X, \ldots, Y_r + c_r X)$ will be regular for some choice of the c_i's. In practice, this will turn a sparse polynomial into a very dense one.

This theorem will be generalized in an exercise. The proof also shows the following result:

Lemma 22. *Let $A = \sum_{i=0}^{m} a_i X^i$, and $B = \sum_{j=0}^{n} b_j X^j$ be polynomials where the coefficients a_i, b_j are indeterminates. We define $\deg(a_i) = m - i$ and $\deg(b_j) = n - j$, so that A, B can be regarded as homogeneous polynomials. Then $\mathtt{res}(A, B)$ is a homogeneous polynomial in $\mathbb{D}[\mathbf{a}, \mathbf{b}] = \mathbb{D}[a_0, \ldots, a_m, b_0, \ldots, b_n]$.*

Exercises

Exercise 12.6: Given $A, B \in \mathbb{D}[X, Y]$, consider the polynomial

$$\Delta(X, Y) = \frac{\det \begin{bmatrix} A(X) & B(X) \\ A(Y) & B(Y) \end{bmatrix}}{X - Y}.$$

(a) Every common root α of A and B satisfies $\Delta(\alpha, Y) = 0$.
(b) Conversely, if $\deg(A) = \deg(B)$ and $\Delta(\alpha, Y) = 0$ then α is a common zero of A, B.
(c) Construct a determinant $R(A, B)$ in the coefficients of A, B such that $R = 0$ iff A, B has a common zero. The determinant $R(A, B)$ is known as the **Bezout resultant** of A, B.
(d) If A, B have degree 2, show that $R(A, B)$ is the same as the usual resultant $\mathtt{res}(A, B)$.
(e) If $\deg(A) = 2$ and $\deg(B) = 3$, how is $R(A, B)$ related to $\mathtt{res}(A, B)$? Generalize this observation. ♢

Exercise 12.7: Given $A, B, C \in \mathbb{D}[X, Y, Z]$, consider

$$\Delta(X, Y, \overline{X}, \overline{Y}) = \frac{\det \begin{bmatrix} A(X,Y) & B(X,Y) & C(X,Y) \\ A(X,\overline{Y}) & B(X,\overline{Y}) & C(X,\overline{Y}) \\ A(\overline{X},\overline{Y}) & B(\overline{X},\overline{Y}) & C(\overline{X},\overline{Y}) \end{bmatrix}}{(X - \overline{X})(Y - \overline{Y})}.$$

(a) Show that every common zero (α, β) of A, B, C satisfies $\Delta(\alpha, \beta, \overline{X}, \overline{Y}) = 0$.
(b) Conversely, if $\deg_i(A) = \deg_i(B) = \deg_i(C)$ ($i = X, Y$) and $\Delta(\alpha, \beta, \overline{X}, \overline{Y}) = 0$. then (α, β) is a common zero of A, B, C.
(c) Construct a determinant $R(A, B, C)$ in the coefficients of A, B, C such that the vanishing of $R(A, B, C)$ is equivalent to $R(A, B, C)$ having a common zero. This is known as the Dixon resultant of A, B, C. ♢

Exercise 12.8: Let $A, B \in \mathbb{D}[X_1, \ldots, X_r]$ be homogeneous with $\deg(A) = m, \deg(B) = n$. If $\deg_i(A)$ denotes the degree of A in X_i, let $\deg_1(A) = m', \deg_1(B) = n'$. Write $\mu = m - m'$ and $\nu = n - n'$. Thus A, B are regular iff $\mu = \nu = 0$. We have the following generalization of a theorem in the text: if $\mathtt{res}_1(A, B) \neq 0$ then $\mathtt{res}_1(A, B)$ is homogeneous of degree $mn - \mu\nu$. ♢

End Exercises

§12.9. Weak Bezout Theorem for Curves

Bezout's theorem that that two curves of degress m and n intersect in exact mn points, when property counted with multiplicities in projective space over the algebraic closure of $\mathbb{D}$. The weak form simply says that that are at most mn intersection points.

Let $A(X,Y) = 0$ and $B(X,Y) = 0$ be two real algebraic curves, and write $R(Y) = \mathtt{res}_X(A,B)$. We may assume that $A, B \in \mathbb{Z}[X,Y]$.

Suppose $p_0 = (x_0, y_0) \in \mathbb{R}^2$ is an intersection point of the two curves. Consider the polynomials $A_0(X) = A(X, y_0)$ and $B_0(X) = B(X, y_0)$ in $\mathbb{R}[X]$: they have a common zero since $A_0(x_0) = B_0(x_0) = 0$. Hence $\deg(gcd(A_0, B_0)) > 0$. Hence $\mathtt{res}_X(A_0, B_0) = 0$. Is $\mathtt{res}_X(A_0, B_0)$ the same as $R(y_0)$? The answer is in the next lemma.

Write $A(X,Y) = \sum_{i=0}^m a_i X^i$ and $B(X,Y) = \sum_{j=0}^n b_j X^j$ where $a_i, b_j \in \mathbb{Z}[Y]$. Write $\mathtt{res}_X(A_0, B_0)$ as the determinant of the matrix S_0 where S_0 is the Sylvester matrix for A_0, B_0. Similarly, let $\mathtt{res}_X(A,B) = \det(S)$ for another Sylvester matrix. Write $S(y_0)$ for the matrix obtained from S by setting $Y = y_0$. So the question is whether $S_0 = S(y_0)$. Note that S_0 depends on the degrees of A_0 and B_0. It is now easy to see:

Lemma 23.
(i) If $a_m(y_0)b_n(y_0) \neq 0$ then $R(y_0) = \mathtt{res}_X(A_0, B_0)$. (ii) If $a_m(y_0) \neq 0$ or $b_n(y_0) \neq 0$ then $R(y_0) = 0$ iff $\mathtt{res}_X(A_0, B_0) = 0$.

Proof. Under the hypothesis of (i), the Sylvester matrix for $\mathtt{res}_X(A_0, B_0)$ is just the specialization of the Sylvester matrix for $\mathtt{res}_X(A,B)$. Under the weaker hypothesis of (ii), the Sylvester matrix for $\mathtt{res}_X(A_0, B_0)$ may just miss a few leading rows of A_0 or B_0. To be specific, assume $a_m(y_0) = 0$ but $b_n(y_0) \neq 0$. Then we see that

$$R(y_0) = b_n(y_0)^k \mathtt{res}_X(A_0, B_0) = 0$$

if the degree of A_0 drops by $k \geq 1$. This proves (ii). **Q.E.D.**

Part (ii) in this lemma is more useful to us than part (i). We can rephrase (ii) as follows. Suppose $R(y_0) = 0$. Then provided $a_m(y_0) \neq 0$ or $b_n(y_0) \neq 0$, we conclude that there exists x_0 such that (x_0, y_0) is a common solution of $A(X,Y) = B(X,Y) = 0$.

Elimination as Projection Operator. For $S \subseteq \mathbb{R}^2$, let $\pi(S) = \{y_0 : (\exists x_0)[(x_0, y_0) \in S]\}$ denote the **X-projection operator**. Thus, the algebraic operation of computing resultant which eliminates the X-variable amounts to the geometric operation of "projecting out the X-coordinate". Let $\textsc{Zero}(A,B) \subseteq \mathbb{R}^2$ denote the zero set of A, B. This shows that

$$\pi_1(\textsc{Zero}(A,B)) \subseteq \textsc{Zero}(R).$$

Intuitively, computing $R = \mathtt{res}_X(A,B)$ corresponds to computing the X-projection of the zero set. We could also project out the Y-coordinate, and obtain the analogous result:

$$\pi_2(\textsc{Zero}(A,B)) \subseteq \textsc{Zero}(\mathtt{res}_Y(A,B)).$$

Let us give another application of the above lemma: suppose $A_Y = \partial A/\partial Y$. Similarly for A_X. Define the set of **singularities** of A to be $Sing(A) := \textsc{Zero}(A, A_X, A_Y)$. Note that $Sing(A)$ includes all self-intersections of A, and isolated points of $\textsc{Zero}(A)$.

$$\pi_2(Sing(A)) \subseteq \textsc{Zero}(\mathtt{res}_Y(A, A_Y).$$

Regular Curves. The preceding discussion motivates the next defintion: a curve $A(X,Y) = 0$ is said to be **regular** (or X-regular) in case $\deg_X(A) = m$ and a_m is a constant. In other words,

$$A(X,Y) = \sum_{i=0}^m a_i(Y) X^i$$

where $a_m(Y)$ is a constant and $\deg_Y(a_i) \leq m - i$. It is easy to see that by a shear transformation

$$(X, Y) \mapsto (X, Y + cX)$$

for almost every constant c, the curve can be made regular (Exercise). We can now prove the weak version of Bezout's theorem.

Theorem 24 (Bezout). *Suppose $A(X,Y), B(X,Y)$ are relatively prime and $\deg(A) = m, \deg(B) = n$. Then the curves $A = 0$ and $B = 0$ has at most mn common points of intersection.*

Proof. By a linear translation, we may assume that A and B are regular. Let $R(Y) = \mathtt{res}_X(A, B)$. Since they are relatively prime, $R(Y)$ is non-zero. Hence $\deg(R) \leq mn$. If (x_i, y_i) $(i = 0, 1, \ldots, mn)$ are $mn + 1$ distinct common intersections, then by a rotation of the curves, we further assume the y_i's are distinct. It follows that $R(y_i) = 0$ for each i. This is a contradiction since there the degree of $R(Y) \leq mn$. **Q.E.D.**

Remark: Bezout's theorem generalize to arbitrary dimensions: given a system of n real polynomials in variables $X_1, \ldots, X_n$, if these polynomials have degrees $m_1, \ldots, m_n$, then there are at most $\prod_{i=1}^n m_i$ non-degenerate zeros. A zero $x = (x_1, \ldots, x_n)$ of the system is **degenerate** if the Jacobian vanishes at x. This theorem can be further generalized to polynomials in $X_1, \ldots, X_n$ and also in exponential terms of the form in $e^{X_1}, \ldots, e^{X_n}$ (Hovanskii's Theorem).

§12.10. Interval Arithmetic

To speed up algebraic computation, we would like to reduce algebraic computation to numerical computation. Numerical computation has the connotation of approximation and uncertainty. So to preserve the exactness that is normally associated with algebraic computation, we must to use some form of interval arithmetic in our numerical computation. The basic idea of interval arithmetic is that we replace a real number x by an interval $[a, b]$ that contains x. We interprete $[a, b]$ as the "interval of uncertainty" associated with x. It is easy enough to extend the basic arithmetic operations $(\pm, \times, \div)$ to such intervals. From such simple ideas, we encounter a rich area of research associated with extending this to more complicated numerical computations. Interval arithmetic is an active area of research that has been greatly influenced by the work of R.E. Moore. More generally, looking beyond arithmetic, it is called **validated computing**.

Boxes and Interval Vectors. Let $\square(\mathbb{R}^n)$ denote the set of closed axes-parallel n-dimensional rectangles which we call **n-boxes** (or simply, boxes). In particular, for $n = 1$, $\square\mathbb{R}$ is the set of closed real intervals, $[a, b]$ where $a \leq b$ and $a, b \in \mathbb{R}$. A typical element of $\square\mathbb{R}^n$ has the form $B = I_1 \times \cdots \times I_n$ where $I_1, \ldots, I_n \in \square\mathbb{R}$. Note that a box B is to be distinguished from the corresponding **interval vector** $(I_1, \ldots, I_n)$. Below, we will need interval vectors when we discuss the interval analogue of the gradient of a function $f : \mathbb{R}^n \to \mathbb{R}$. Let $(\square\mathbb{R})^n$ denote the set of interval n-vectors. We will simply write "$\square\mathbb{R}^n$" for both $\square(\mathbb{R}^n)$ and $(\square\mathbb{R})^n$, since the context will make clear which is intended.

Each box B is regarded as a subset of $\mathbb{R}^n$. We shall regard $\mathbb{R}$ as a subset of $\square\mathbb{R}$ where $a \in \mathbb{R}$ is identified with $[a, b] \in \square\mathbb{R}$. Similarly, $\mathbb{R}^n \subseteq \square\mathbb{R}^n$ for all $n \geq 1$.

For an interval x, let $\overline{x}$ and $\underline{x}$ denote its **lower** and **upper bounds**, i.e., $x = [\overline{x}, \underline{x}]$. Let $\mathrm{width}(x) := \underline{x} - \overline{x}$ denote the **width** of an interval x. Also, $\mathrm{mid}(x) = \frac{x - \overline{x}}{2}$ denote the **midpoint** of x. We extend these concepts to boxes: $\mathrm{width}(x_1 \times \cdots \times x_n) := \max_{i=1}^n \mathrm{width}(x_i)$, and $\mathrm{mid}(x_1 \times \cdots \times x_n) := (\mathrm{mid}(x_1), \ldots, \mathrm{mid}(x_n))$.

Interval Arithmetic. Arithmetic $(+, -, \times, \div)$ on real numbers can be extended to intervals in the obvious way. E.g., $x \pm y = [\underline{x} \pm \underline{y}, \overline{x} \pm \overline{y}]$. For multiplication, we need to consider all four combinations of the end points, $xy = [\min\{\underline{x}\underline{y}, \underline{x}\overline{y}, \overline{x}\underline{y}, \overline{x}\overline{y}\}, \max\{\underline{x}\underline{y}, \underline{x}\overline{y}, \overline{x}\underline{y}, \overline{x}\overline{y}\}]$. Similarly for division x/y provided $0 \notin y$.

Comparison of intervals is only defined for certain pairs: we say $I \leq J$ if for all $x \in I, y \in J$, we have $x \leq y$. Equivalently, $\overline{I} \leq \underline{J}$. Similarly, $I < J$ if for all $x \in I, y \in J$, we have $x < y$. Equivalently, $\overline{I} < \underline{J}$. E.g., we may write $I > 0$ for a positive interval.

For any function $f : \mathbb{R}^n \to \mathbb{R}$, let $\Box f : \Box\mathbb{R}^n \to \Box\mathbb{R}$ denote some box-valued version of f, with the property

$$f(B) \subseteq \Box f(B)$$

and satisfying the convergence property: if $B_{i+1} \subseteq B_i$ for all $i \geq 1$ and $\lim_{i\to\infty} \text{width}(B_i) \to 0$, then $\text{width}(\Box f(B_i)) \to 0$. Now, if f is a continuous function then $B_i \to p$ implies $\Box f(B_i) \to f(p)$.

If f is a polynomial, then $\Box f$ is easy to implement using interval arithmetic: we just take any method to evaluate f, but now assume the interval analogues of the evaluation. E.g., if $f(x) = 2x^2 - 3x$, we can evaluate f as the sum of two monomials, $2x^2$ and $-3x$. Let $x = [-1, 2]$. Then $x^2 = [-2, 4]$ and $\Box f(x) = [-4, 8]+[-6, 3] = [-10, 11]$. Next consider evaluating $\Box f(x)$ using Horner's rule, $\Box f(x) = (2x-3)x = ([-2, 4] - 3)[-1, 2] = [-5, 1][-1, 2] = [-10, 5]$.

This example shows that interval evaluations depends on the order of evaluation. For every polynomial, there are two obvious methods of evaluation, as illustrated above: (1) the power method (compute all powers of x first, form all the monomials, and finally add the monomials) and (2) Horner's rule. It is not hard to show that (2) gives an interval that is no larger than (1). This result follows from the subdistributivity law for interval arithmetic:

$$I(J + K) \subseteq IJ + IK.$$

More generally, for any algebraic expression $F(x)$ that evaluates to $f(x)$, there is an associated interval function $\Box F$ that represents $f(x)$. Note that we view expressions as DAGs or equivalently, straightline programs. The investigation of optimal methods for constructing box functions, not necessarily derived from evaluation of expressions, is studied in interval arithmetic (see Ratschek and Rokne [6]).

Suppose f is a multivariate polynomial in $x_1, \ldots, x_n$. We want to generalize Horner's rule for evaluating f at the box $B = I_1 \times \cdots \times I_n$. We can view $f(x)$ as a polynomial in x_n, whose coefficients are polynomials in $x_1, \ldots, x_{n-1}$. Then recursively, we evaluate the coefficients of f at the box $I_1 \times \cdots \times I_{n-1}$. Finally, we use the univariate version of f to evaluate it at I_n. By reordering the variables, we get different algorithms. It is by no means clear which ordering gives the best results.

How to confirm the shape of a curve inside a box (up to isotopy)? Consider a curve $f(X, Y) = 0$ which is non-singular. Consider a box $B \in \Box\mathbb{R}^2$ whose corners are v_0, v_1, v_2, v_3 (in clockwise order, with the northwest corner at v_0). Suppose $f(v_0)f(v_1) < 0$. Thus the curve passes through B. Let f_X, f_Y denote the partial derivatives of f with respect to X and Y. Thus the gradient of f is $\nabla f = (f_X, f_Y)$. Suppose $\Box f_X(C) > 0$.

Exercises

Exercise 12.9: Prove that if $B \subseteq \mathbb{R}^2$ is a square box and for all $p, q \in B$, we have $\langle \nabla f(p), \nabla f(q) \rangle > 0$ then $f = 0$ is parametrizable in the x- or y-direction. ♢

Exercise 12.10: What other shapes besides a square box would permit such a conclusion? ♢

Exercise 12.11: Generalize the above argument to multidimensional cubes. ♢

End Exercises

§13. Algebraic Cell Decomposition

The problem before us is a very general one: given a set of polynomials $P \subseteq \mathbb{Z}[X_1, \ldots, X_n]$, and a open region $R \subseteq \mathbb{R}^n$, to compute a partition K of R into cells such that over each cell, P is sign-invariant. By a cell (or d-cell) we mean a subset of $\mathbb{R}^n$ that is homeomorphic to some open ball B^d ($d = 1, \ldots, n$) or a singleton ($d = 0$). Typically, R is the interior of a box in $\Box\mathbb{R}^n$, and we call R the **region of interest** (ROI).

We say P is **sign-invariant** over a cell $\beta \in K$ if there is a **sign assignment** $\sigma : P \to \{+, 0, -\}$ such that for all $b \subseteq \beta$ and $f \in P$, $f(b)$ has the sign $\sigma(f)$. An example will clarify this concept: in Figure 10, we

see a partition K of the region R into cells of dimensions 0, 1 and 2. The set $P = \{f, g\}$ corresponds to two curves where f has a loop and g is an ellipse. Assume f is positive inside the loop, and g is positive inside the ellipse. Note that K has six 2-cells, seven 1-cells and two 0-cells in K. The sign assignments for each of the 2-cells are indicated. The sign assignments for the 1- and 0-cells are also easily deduced.

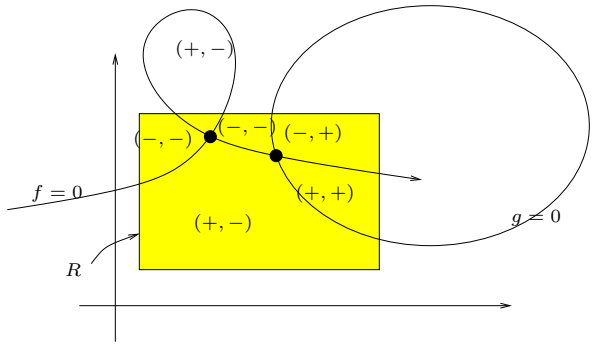

Figure 10: Cell Decomposition

Depending on the applications, we may need additional properties of the cell decomposition K:

- We may want K to be cylindrical (see below) for purposes of quantifier elimination.
- We may want to compute adjacency information in applications such as motion planning. Two cells β, β' are **adjacent** if $\beta \cup \beta'$ is connected. If β, β' are adjacent and $\dim(\beta) = \dim(\beta') + 1$, then we say β' is **incident** on β. Note that incidence is an asymmetric relation, unlike adjacency. We may now define the **incidence graph** of K to be the directed graph whose vertices are the cells of K and $\beta' \to \beta$ is an edge iff β' is incident on β.
- We may need to determine a point p in each cell $\beta \in K$, where p is called a **sample point**. If β has dimension $k \geq 0$, we can ensure that k of the coordinates of p are rational numbers.
- We may want K to form a regular cell complex for the purposes of computing certain topological invariants of the underlying space.

The basic tool for constructing cell decompositions K is the projection operator. Suppose $S \in \mathbb{R}^n$ and $i = 1, \ldots, n$. Then $\Pi_i(S) = \{(x_1, \ldots, \widehat{x_i}, \ldots, x_n) : (x_1, \ldots, x_n) \in S\}$ denotes the i-th projection operator. Also, let $\pi_i(S) = \{x_i : (x_1, \ldots, x_n) \in S\}$.

Suppose $S = \texttt{Zero}(A(X,Y), B(X,Y))$ and $R(X) = \texttt{res}_Y(A, B)$. Also let $a_m(X), b_n(X)$ be the leading coefficients of A, B. For all $x \in \mathbb{R}$, $R(x) = 0$ implies
(i) $a_m(x) = 0$ and $b_n(x) = 0$, or
(ii) there exists a y such that $A(x, y) = B(x, y) = 0$.
Thus

$$\Pi_2(\texttt{Zero}(A, B)) \subseteq \texttt{Zero}(R).$$

We may therefore think of $R(X) = \texttt{res}_Y(A, B)$ as the algebraic analogue of projection along the Y-axis.

Overview: suppose $P \subseteq \mathbb{Z}[X_1, \ldots, X_n]$ and K is a cell decomposition of some open box $R \subseteq \mathbb{R}^n$. We say K is a **P-invariant cell decomposition** if P is sign-invariant over each cell of K. Let $\Pi_n(K)$ denote the set $\{\Pi_n(S) : S \in K\}$. A cell decomposition K is **cylindrical for** $R \subseteq \mathbb{R}$ if K is a partition of R and:

- If $n = 1$ then this simply means that $\texttt{Zero}(f) \cap R \subseteq K$ for each $f \in P$. So assume $n > 1$ in the rest of this.
- $\Pi_n(K)$ is cell decomposition of $\Pi_n(R)$.

- For each cell $\beta' \in \Pi_n(K)$, consider the set of cells $\beta \in K$ such that $\Pi_n(\beta) = \beta'$. This set is called the **stack** over β', denoted $stack(\beta')$.
- There is a total ordering on the elements in $stack(\beta')$ such if $\beta_1 < \beta_2 < \cdots \beta_m$ is the ordering, and if $p_i \in \beta_i$ (for each i) then the n-th component of p_i's would be sorted in the same order.
- If $\dim(\beta') = d$ then $\dim(\beta_i)$ is either d or $d+1$. We call β a **section** if its dimension is d, and a **sector** if its dimension is $d+1$.

When $n = 1$, the cell decomposition of P is easy: it is just real root isolation. For $n > 1$, our goal is to define a set $P' = Proj(P) \subseteq \mathbb{Z}[X_1, \ldots, X_{n-1}]$ such that if K' is a P'-invariant cell decomposition of $Proj_n(R)$ then we can obtain a P-invariant cell decomposition K such that $K' = \Pi_n(K)$. We define $Proj(P)$ as the set comprising the following polynomials: let $f, g \in P$ and $f \neq g$.

- Write $f = \sum_{i=0}^{m} f_i X_m^i$ where $f_i \in \mathbb{Z}[X_1, \ldots, X_{n-1}]$. Then for each $k = 0, \ldots, m$, we add f_k to $Proj(P)$ whenever f_k is non-constant.
- Add the discriminant $psc_k(f, f')$ to $Proj(P)$ for each $k = 0, \ldots, \deg_n(f) - 2$ where $f' = \partial f / \partial X_n$.
- Add $psc_k(f, g)$ to $Proj(P)$ for each $k = 0, \ldots, \min(\deg_n f, \deg_n g)$

The problem could be reduced to single polynomial $p = \prod P$, obtained as a product of all the polynomials in P. In practice, this can be quite inefficient.

We claim that we can "lift" (opposite of project) K' to K as follows: consider the cylinder $\beta' \times \mathbb{R} \subseteq \mathbb{R}^n$ where $\beta' \in K'$. For each $b \in \beta'$ and $f \in P$, we do the real root isolation of $f(b, X_n)$. The number of these real roots do not depend on the choice of b. Let $f_i(b)$ denote the i-th root function defined in this way. Moreover, if $f, g \in P$, then these root function $f_i(b), g_j(b)$ are totally ordered in a way that does not depend on b: either $f_i(b) < g_j(b)$ or $f_i(b) = g_j(b)$ or $f_i(b) < g_j(b)$. In this way, we can define the sectors and sections in the stack.

REMARK: this simple projection does guarantee other nice properties we might like. E.g., K need not be a regular cell complex.

EXERCISES

Exercise 13.1: Using Core Library, implement a simple cell-decomposition algorithm in 2-D. To keep the problem simple, suppose you are given three polynomial $f, g, h \in \mathbb{Z}[X, Y]$ and an open rectangular region $R = (a, b) \times (c, d)$, we want to compute a sign-invariant decomposition K of $P = \{f, g, h\}$. The output is supposed to be the incidence graph of K. Moreover, we want some to be able to display the curves (so we can verify your decomposition).

REMARK: you may use the tools already implemented in Core. Under ${CORE}/progs/curves/, you will see that we have the ability to roughly trace implicit curves. You should also be able to display the curves (try "make show" in this directory). ♢

Exercise 13.2: Instead of an open rectangular region R, assume that R is defined by a set of polynomial inequalities. Modify the ACD algorithm to deal with this situation. ♢

END EXERCISES

§14. Meshing

The Problem. The **surface meshing problem** is the following: given an surface $S \subseteq \mathbb{R}^3$ defined by the equation $F(X, Y, Z) = 0$, to construct a complex K that is topologically similar to S, and which is geometrically close to S, to within a given $\varepsilon > 0$. Typically, $F(X, Y, Z) \in \mathbb{Z}[X, Y, Z]$.

We may also call K a (surface) **mesh**. In case K is a simplicial complex, it is also called a **triangulation**. Although our main interest is in surface meshing, we sometimes discuss the case of curves $S \subseteq \mathbb{R}^3$. Meshing a curve S amounts to finding a polygonal approximation of S. There is also the problem of **volume meshing** to construct a 3-dimensional complex of the volume bounded by (or exterior to) some given bounded surface.

In practice, we might also be given a given rectangular region $R \subseteq \mathbb{R}^3$ and we only want to mesh the subset $S \cap R$. R is called the region of interest (ROI), typically $R \subseteq \square\mathbb{R}^3$. Furthermore, we can divide the dual requirements of meshing (topological and geometric) into two subproblems: the **topological meshing problem** is to compute a mesh (from scratch) satisfying the topological requirement only. The **meshing refinement problem** is to refine a given topologically correct mesh to satisfy the geometric requirement. Typically, the topological meshing problem is the harder of the two, and we shall focus on this.

There is a precedent for this separation of the topological and geometric requirements: if meshing the surface $f(x, y, z) = 0$ is 2-D meshing, and computing polygonal approximation to a curve $f(x, y) = 0$ is 1-D meshing, then finding roots of a function $f(x)$ may be viewed as 0-D meshing. In finding roots of $f(x)$, we typically subdivide the problem into the root isolation problem and the root refinement problem. These are precisely the topological meshing and mesh refinement subproblems of 0-D meshing.

As noted by Boissonnat et al, the strongest notion of "topologically similar" has evolved in over the last few years. Originally, researchers just require $|K|$ to be homeomorphic to S. Clearly, this is inadequate: if S is a knotted torus embedded in space, we do not regard a standard unknotted torus to be a correct output. In other words, we are also interested in how S is embedded in $\mathbb{R}^3$.

FIGURE

We shall interpret $S, |K|$ to be topologically similar to mean that S and $|K|$ are ambient isotopic to each other. Two surfaces $S, S' \subseteq \mathbb{R}^3$ are **ambient isotopic** if there exists a continuous map

$$\gamma : [0, 1] \times \mathbb{R}^3 \to \mathbb{R}^3$$

such that:
(i) For each $t \in [0, 1]$, the map $\gamma_t : \mathbb{R}^3 \to \mathbb{R}^3$ given by $\gamma_t(p) = \gamma(t, p)$ is a homeomorphism;
(ii) $\gamma_0(p) = p$ for all $p \in \mathbb{R}^3$; and
(iii) $\gamma_1(S) = S'$.

Such a map γ is also called an **ambient isotopy** from S to S'. It is an **ε-ambient isotopy** if, in addition, we have $\|\gamma_0(p) - \gamma_1(p)\| \leq \varepsilon$ for all $p \in S$. Thus, we want $|K|$ to be ε-ambient isotopic to S in the meshing problem.

If we restrict γ to the domain $[0, 1] \times S$, then we obtain an isotopy. We say S is **isotopic** to S' in case there exists such an isotopy. Clearly ambient isotopy implies isotopy, which in turn implies homeomorphism. Thus the difference between ambient isotopy and isotopy is that the former requires a simultaneous transformation of the complementary space $\mathbb{R}^3 \setminus S$. But Hirsch [Differential Topology, Springer 1976] shows that, conversely, a isotopy can be extended to an ambient isotopy in the present setting.

§14.1. Generic Subdivision Algorithm

There are few current algorithms that guarantee topological correctness in the above sense. We shall study two such algorithms: Snyder's algorithm and the Plantinga/Vegter algorithm. Both are based on interval arithmetic and the paradigm of space subdivision.

A simple example of space subdivision algorithm is the **marching cube algorithm**. Suppose we want to compute a simplicial complex for the surface $S : F = 0$ within the rectangular region R. We subdivide R into smaller boxes, using an octtree data structure where each rectangle is subdivided into 8 equal size sub-boxes. For each vertex v of the boxes, we evaluate the function F. If (v, v') is an edge of a box B, and $F(v)F(v') < 0$, then we know that S intersects the box (in fact, intersects an odd number of times). Therefore, this model

requires the ability to evaluate the sign of F at given points. Of course if $F(v)F(v') > 0$, we are not sure if there is any intersection inside this box. But assuming that S does not have "features smaller than than ε", we might be able to conclude that S does not intersect the box. In this way, we can choose the midpoint of (v, v') as a vertex of our mesh K. We then need to know how to connect the vertices of K to form edges, and how to form triangular faces from these edges. These decisions can be done in each box B. As the box has 8 vertices, there are 2^8 possibilities, but by symmetry, we reduce the number of cases to 15 distinct cases. This is the basis of the marching cube algorithm for meshing. The problem with the Marching cube approach is the ambiguity of constructing the triagular faces in a cube B. In fact, the problem already shows up in the case where S is a curve and B is a square: assume that that vertices of B has the **alternating sign** pattern:

[FIGURE: Alternating signs]

Even assuming that we know that the curve S intersect each side of the square B at most once, and S has at most one singular point inside B, we are left with three possibilities as shown in Figure **??**(b,c,d).

The marching cube algorithm and the ones that follow are instances of the following abstract algorithm:

```
Generic Subdivision Algorithm:
    Input: Region of interest R and other data.
    Output: A set of boxes with associated data.
    ▷ INITIALIZE
        Let Q be a queue, initially containing just R
    ▷ MAIN LOOP
        while Q is non-empty
            Remove B from Q
            CASE Test(B)
                TRUE: Output(B)
                    Break
                FALSE: Discard (B)
                    Break
                UNDECIDED: Subdivide B into sub-boxes B_i (i = 1, ..., k)
                    and put each B_i into Q
            END CASE
    ▷ POST-PROCESSING
        Go through the list of outputted boxes and
        to produce the required output data.
```

For instance, $Test(B)$ may simply check whether (1) the surface intersects B, and (2) the diameter of B is smaller than ε. A crude curve plotting algorithm would then place a point in the center of B as a vertex of the polygonal approximation. These vertices are simply connected to neighboring boxes in some heuristic manner.

The fundamental question is how can we guaranteed that the correctness of this algorithm? Correctness here means halting as well as correctness of any output. In general, this requires two additional ingredients.

- (i) We may need additional properties about S. In the following, we shall assume S is regular, meaning that the gradient of $F(X, Y, Z)$ does not vanish on S
- (ii) We may need to strengthen the primitives used by our meshing algorithm. We shall assume the ability to perform interval evaluation of F and its derivatives. In general, $\square f$ denote the interval version of a function f.

Parametrizability. If $S \subseteq \mathbb{R}^2$ is the graph of a function $f : \mathbb{R} \to \mathbb{R}$, then we say[3] S is **x-parametrizable** over $\mathbb{R}$ to mean that for each $x_0 \in \pi_x(B)$, there is at most one y_0 such that $(x_0, y_0) \subseteq S \cap B$.

Note that when S is x-parametrizable over B, we know that S is does not contain any loops inside B and also S has no self-intersection. In case S is an algebraic curve, we know that $S \cap B$ is non-empty iff $S \cap \partial B$ is non-empty. Thus, parametrizability leads to very nice properties for meshing.

In general, let $A = \{k_1, \ldots, k_r\}$ where $1 \leq k_1 < k_2 < \cdots < k_r \leq n$. Recall that $\pi_A((x_1, \ldots, x_n)) = (x_{k_1}, \ldots, x_{k_r})$ and for $S \subseteq \mathbb{R}^n$, we have $\pi_A(S) = \{\pi_A(p) : p \in S\}$. If $S \in \Box\mathbb{R}^n$ and $y = (y_1, \ldots, y_r) \in \mathbb{R}^r$, let

$$S|y := \{(x_1, \ldots, x_n) \in S : x_{k_i} = y_i(\text{ for all } \ i = 1, \ldots, r\}.$$

We say that a set $S \subseteq \mathbb{R}^n$ is **A-parametrizable over** B if there is at most one zero in $B|y$ for each $y \in \pi_A(B)$.

For instance, with $n = 3$ and $A = \{3\}$, we have $J^A(X, Y, Z) = \begin{bmatrix} f_x & f_y \\ g_x & g_y \end{bmatrix}$ where $S = \text{ZERO}(f, g)$. If $f(X, Y, Z) = X$ and $g(x, y, z) = y$, then $J^A(X, Y, Z) = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$.

FIGURE: The unit circle $X^2 + Y^2 = 1$ is Y-parametrizable in the region $B = [\frac{1}{2}, 1] \times [-1, 1]$, and X-parametrizable in the region $B = [-1, 1] \times [\frac{1}{2}, 1]$.

We will be interested in the general situation where $S = \texttt{Zero}(f_1, \ldots, f_{n-r})$ where $f_i : \mathbb{R}^n \to \mathbb{R}$ are differentiable functions. First, we need several related concepts.

Let $\overline{A} = \{1, \ldots, n\} \setminus A$. Assume the elements of $\overline{A}$ are sorted as $1 \leq \ell_1 < \cdots < \ell_{n-r} \leq n$. Consider the Jacobian $J^A = J^A(X)$ which is the $(n-r) \times (n-r)$ matrix $[J_{ij}]_{i,j=1}^{n-r}$ where $J_{ij} = J_{ij}(X) := \frac{\partial f_i}{\partial X_{\ell_j}}(X)$.

Interval Matrices. An **interval matrix** $M = [M_{ij}]_{i,j=1}^n$ is a matrix whose entries are intervals, $M_{ij} \in \Box\mathbb{R}$. Suppose $m = [m_{ij}]_{i,j=1}^n$ where $m_{ij} \in \mathbb{R}$ is an "ordinary matrix". We write $m \in M$ if $m_{ij} \in M_{ij}$ for all i, j. Sometimes, we simply view M as the set of all m such that $m \in M$. We shall be interested in the predicate $0 \notin \det M$, i.e., for all $m \in M$, $\det(m) \neq 0$.

If $B \subseteq \Box\mathbb{R}^n$, then the box version of J^A is given by $\Box J^A(B) = [\Box J_{ij}(B)]_{i,j=1}^{n-r}$ and $\Box J_{ij}(B) \in \Box\mathbb{R}$ is the usual box function of J_{ji}. Thus $\Box J^A(B)$ is an interval matrix. The following is from Snyder [**?**, **?**] can be seen as an interval version of the implicit function theorem:

LEMMA 25. *Suppose $S = \text{ZERO}(f_1, \ldots, f_{n-r})$ and $B \in \Box\mathbb{R}^n$. Fix $A = \{k_1, \ldots, k_r\}$ as above. If $0 \notin \det\Box J^A(B)$ then S is A-parametrizable over B.*

Let us show this for the case where $n = 2$ and $r = 1$. In this case, S is the curve $f(X, Y) = 0$. With $A = \{1\}$, we have $J^A = \frac{\partial f}{\partial y} = f_y$. Then the above lemma simplifies to the claim that $0 \notin \Box f_y(B)$ implies S is x-parametrizable over B. In proof, suppose there are two distinct points (x_0, y_0) and (x_0, y_1) in $S \cap B$. Then by Rolle's theorem, $f_y(x_0, c) = 0$ for some $c \in (y_0, y_1)$. This contradicts the condition $0 \notin \Box f_y(B)$.

This is a kind of implicit function theorem. For a semi-algebraic version of the implicit theorem, see [2, p. 134].

Meshing Using Parametrizability First consider a curve $S : f(X, Y) = 0$ that is a 1-manifold. Thus S is a collection of pairwise disjoint loops or infinite curves. Snyder (1992) shows how we can compute a polynomial approximation:

[3]The indexing system for discussing parametrization can be confusing – should we call this x-parametrizability or y-parametrizability? We suggest the rule that x-parametrizability should mean that x can be used as a parameter for corresponding implicit function h, where $f(x, h(x)) = 0$.

Curve Approximation Algorithm:
Input: Region of interest $R \subseteq \mathbb{R}^2$ and a curve $S : f(x, y) = 0$.
Output: A topologically correct polygonal approximation.
▷ *INITIALIZE*
Let Q be a queue, initially containing just R
▷ *MAIN LOOP*
while Q is non-empty
Remove B from Q
if (S is x- or y-parametrizable over B):
Assume s is x-parametrizable over B
Compute $S \cap \partial B = \{p_1, \ldots, p_m\}$ using root isolation.
Refine the isolating intervals until we order their x-values: $p_1.x < p_2.x < \cdots < p_m.x$
If $m \leq 1$, discard B.
Else for each $i = 1, \ldots, m$,
Assume $a_i < p_i.x < b_i$ with the (a_i, b_i)'s disjoint
If $f((b_i + a_{i+1})/2, y)$ changes sign over the interval $\pi_y(B)$
we "link" p_i and p_{i+1}.
else
Subdivide B into sub-boxes B_i $(i = 1, \ldots, 4)$
Put each B_i into Q
▷ *POST-PROCESSING*
Go through the output boxes, and identify points on the common
edges of adjacent boxes.

Note that linking p_i with p_{i+1} are of three kinds: one of p_i, p_{i+1} is on a vertical edge of B, p_i, p_{i+1} lies on opposite edges of B, and p_i, p_{i+1} lies on the same edge of B,

Does this algorithm halt? If the curve grazes the boundary edge then some root isolation strategies may not halt. See the chapter of Vegter for an example.

§14.2. Vegter Plantinga's Approach

We now present another approach to isotopic meshing. Assume $S : f(X, Y) = 0$ is the curve and f is continuously differentiable. Let $f_x = \frac{\partial f}{\partial X}$ and $f_y = \frac{\partial f}{\partial Y}$. Assume S is non-singular, i.e., for all $p \in S$, the gradient does not vanish:

$$\nabla f(p) = (f_x(p), f_y(p)) \neq (0, 0).$$

Now, we want the interval version of the gradient, where for any box B, we define

$$\square\nabla f(B) := (\square f_x(B), \square f_y(B)).$$

We can form the scalar product of this interval gradient function,

$$\langle \square\nabla f(B), \square\nabla f(B)\rangle := \square f_x(B) \cdot \square f_x(B) + \square f_y(B) \cdot \square f_y(B)).$$

The algorithm is based on two basic predicates or conditions:

- $C_0(B) : 0 \notin \square f(B)$.
- $C_1(B) : 0 \notin \langle \square\nabla f(B), \square\nabla f(B)\rangle$.

Consequences of Small Normal Variation Condition. Clearly, if $C_0(B)$ holds, the curve does not pass through B. The condition $C_1(B)$ tells us that something equality interesting. Intuitively, it means that the gradient of every point in B is roughly pointing in the same general direction. More precisely, for all $p, q \in B$ we have $\langle \nabla f(p), \nabla f(q) \rangle > 0$, i.e., the gradient vectors spans an angle less than 90°. In view of this, $C_1(B)$ is also known as the **small normal variation condition** on B. To use this condition, we need several

Motivation Let us see how useful this condition is: there are two ambiguities in the Marching cube algorithm: the first ambiguity is that the sign of f at the endpoints of an side of a box only tells us the parity (odd or even) of the intersection of the curve with the side. The **epsilon heuristic** is that when the box is small enough, we assume that odd equals one intersection, even equald no intersection. The second ambiguity comes from the alternating sign pattern at the corners of a square box:

FIGURE: (a) Alternating pattern (b) (c) two ways to connect four vertices.

Using the epsilon heuristic, we introduce four vertices, one in the middle of each side of the box. Now, it is unclear how to connect these four vertices within the box in a topologically correct way (fig(b) or (c)). We next see that condition C_1 saves us both ambiguities.

There are several basic arguments that can be applied in is the following 2-dimensional analogue of Rolle's theorem for real functions:

Lemma 26. *Suppose $f(x, y)$ is C^1 and p, q are points such that $\nabla f(p), \nabla f(q)$ are not parallel. Let $u = \alpha \nabla f(p) + (1 - \alpha)\nabla f(q)$ for some $0 \leq \alpha \leq 1$. Let p_t denote the point $(1 - t)p + tq$. Then there exists some $t \in [0, 1]$ such that either $\nabla f(p_t) = 0$ or $\nabla f(p_t)$ is parallel to u.*

Proof. Assume that for all $t \in [0, 1]$, $\nabla f(p_t) \neq 0$. Then the unit vector $d_t = \nabla f(p_t)/\|\nabla f(p_t)\|$ is well-defined for all $t \in [0, 1]$. If $d_t = (0, 0)$ for some t, then the lemma holds trivially. Otherwise, we have d_0 is parallel to $\nabla f(p)$ and d_1 is parallel to $\nabla f(q)$. This means that for some $t \in [0, 1]$, d_t must be parallel to u. **Q.E.D.**

Lemma 27. *Suppose $C_1(B)$ holds at a box B.*
(i) The curve $f = 0$ is parametrizable in the X- or Y-direction in B.
(ii) The sign of f at the four vertices of B cannot form an alternating pattern.

Proof. (i) The condition $C_1(B)$ means that $0 \notin \square f_x(B) \cdot \square f_x(B) + \square f_y(B) \cdot \square f_y(B))$, and this in turn means that $0 \notin \square f_x(B)$ or $0 \notin \square f_y(B)$. The latter two conditions imply parametrizability in either X- or Y-direction.

(ii) Suppose the sign of f is alternating, say $f(NE) > 0, f(NW) < 0, f(SW) > 0, f(SE) < 0$. Then there is a point p on the north side with $f_x(p) < 0$ and a point q on the south side with $f_x(q) > 0$. Thus $0 \in \square f_x(B)$. Similarly, $0 \in \square f_y(B)$. This contradicts the condition $C_1(B)$. **Q.E.D.**

The intuitive idea is that repeated subdivision of a box B will eventually make C_0 or C_1 true. This is the basis of the following algorithm:

```
Isotopic Curve Approximation Algorithm:
    Input: Region of interest R ⊆ ℝ² and a curve S : f(X, Y) = 0.
    Output: A topologically correct polygonal approximation.
        Initialize a quadtree T whose root is B.
        Subdivide each leaf B' of T until either C0(B') or C1(B') holds.
        Balance(T)
        for each box edge of leaf of the quadtree
            if the signs of F at the two endpoints are opposite
            then insert a vertex at the midpoint of the edge
        for each leaf of T
            if the leaf contains two vertices, connect them by a segment
            else if the leaf contains four vertices
            then find the 2 vertices on one side,
                    and connect them to the other two vertices without crossing
```

We first prove termination:

Lemma 28. *The above algorithm terminates.*

Proof. There is a minimum distance $\varepsilon > 0$ between the solutions of $f = 0$ and $\nabla f = 0$. Thus, when a box B has radius less than ε, either $0 \notin f(B)$ or $0 \notin \nabla f(B)$. Then, by the convergence of interval arithmetic, $C_0(B)$ or $C_1(B)$ eventually holds. **Q.E.D.**

Next, let us prove isotopy of the polygonal approximation. We may ignore the case where f is zero at a corner of a box. Actually, it is not hard to see that if f is zero at a corner, we can freely choose either non-zero sign for the corner.

First we do not consider the adaptive grid produced by the algorithm. Instead, assume that the grid is regular, i.e., all leaf boxes have the same size.

FIGURE: semicircle

Lemma 29. *Assuming the regular grid, then V-P algorithm is correct.*

Proof. Note that in the case, each box has either 0 or 2 vertices. Moreover, each edge has 0 or 1 vertex.

We must consider the possibility of the curve cutting the boundary of B more than once. Suppose the input curve cuts a horizontal edge twice at p, q. Then curve must be x-parametrizable in B.

This curve must not exit the semicircle with diameter $[p, q]$. If the curve reaches a point r outside the semicircle, then we get contradiction of $C_1(B)$. We can repeat this argument for every pair (p, q). What if the ends are not paired? **Q.E.D.**

Going back to the adaptive, balanced quadtree, we have:

Lemma 30. *There are zero, two or four vertices on the boundary of each box of the balanced quadtree.*

Proof. The reason for an even number is clear: each vertex is determined by the sign of f at the vertices of B or the midpoints of edges of B. **Q.E.D.**

Lemma 31. *The polygonal approximation for the balanced quadtree is isotopic to the polygonal approximation for the regular grid.*

Proof. **Q.E.D.**

Comparison with Snyder. one big advantage over Snyder is that we do not have to find zeros (i.e., intersect the curve C with edges of the boxes). We just do sign evaluation of f at vertices, and interval evaluation.

Note that because Vegter-Plantinga does not need to determine the exact topology of the curves within a cell, it may terminate faster in many situations.

The extension to 3-D is similar with more case analysis. Presumably this extensions to all dimensions.

Exercises

Exercise 14.1: Provide a convincing proof of correctness of V-P. ♢

Exercise 14.2: Extend the V-P result to all dimensions. ♢

Exercise 14.3: Extend the Vegter-Plantinga result to curves which may be singular. ♢

End Exercises

§15. Notes on Cohomology

Cohomology: in some ways this is more natural than cohomology, but it was much slower in being developed. Wikipedia has a nice history. See Munkres.

The HOM Functor: If A, G are Abelian groups, we obtain another Abelian group

$$Hom(A, G)$$

of all homomorphisms of A into G. This groups is essential for defining cohomology.

Some facts about $Hom(A, G)$. Let $\phi, \psi \in Hom(A, G)$. Then the group operation on $Hom(A, G)$ is defined via $(\phi + \psi)(a) = \phi(a) + \psi(a)$ for all $a \in A$. We must show that $(\phi + \psi)$ is a homomorphism from A to G: for all $a, b \in A$, $(\phi + \psi)(a + b) = \phi(a + b) + \psi(a + b) = (\phi(a) + \phi(b)) + (\psi(a) + \psi(b)) = (\phi(a) + \psi(a)) + (\phi(b) + \psi(b)) = (\phi + \psi)(a) + (\phi + \psi)(b)$. Note that the zero element $Hom(A, G)$ is the function $\phi(a) = 0$ for all $a \in A$. The inverse of ϕ is $-\phi$ where $-\phi(a) = \phi(-a)$.

Example: $Hom(\mathbb{Z}, G)$ is isomorphic to G, where the isomorphism assigns to $\phi \in Hom(\mathbb{Z}, G)$ the element $\phi(1)$. More generally, if A is a free Abelian group of rank n, then $Hom(A, G)$ is isomorphic to $\underbrace{G \oplus \cdots \oplus G}_{n}$.

Dual Homomorphism. Any homomorphism $f : A \to B$ gives rise to a **dual homomorphism** $\widetilde{f}$ going in the reverse direction,

$$Hom(B, G) \xrightarrow{\widetilde{f}} Hom(A, G)$$

where $\phi : B \to G$ is mapped to

$$\widetilde{f}(\phi) : A \xrightarrow{f} B \xrightarrow{\phi} G.$$

The map $\widetilde{f}$ is a homomorphism because

$$[\widetilde{f}(\phi + \psi)](a) = (\phi + \psi)(f(a)) = \phi(f(a)) + \psi(f(a)) = [\widetilde{f}(\phi)](a) + [\widetilde{f}(\psi)](a).$$

Category. We introduce the useful language of categories. A **category** C consists of 3 things:
(i) A class of **objects** X.
(ii) For every ordered pair (X, Y) of objects, a set $hom(X, Y)$ of **morphisms** f. We write $f : X \to Y$ to denote that $f \in hom(X, Y)$.
(iii) For every triple of objects (X, Y, Z) a function called **composition** of morphisms, $hom(X, Y) \times hom(Y, Z) \to hom(X, Z)$.
If $(f, g) \in hom(X, Y) \times hom(Y, Z)$, then their composition is denoted $g \circ f$. These concepts satisfies two axioms:

Axiom 1 (Associativity). If $f \in hom(W, X)$ and $g \in hom(X, Y)$ and $h \in hom(Y, Z)$ then $h \circ (g \circ f) = (h \circ g) \circ f$.

Axiom 2 (Identities). If X is an object, there is an element $1_X \in hom(X, X)$ such that

$$1_X \circ f = f, \qquad g \circ 1_X = g$$

for all $f \in hom(W, X)$ and $g \in hom(X, Y)$, for all W, Y.

A (covariant) **functor** G from category C to category D is a function assigning each object X of C to an object $G(X)$ of D, and each morphism $f : X \to Y$ of C to a morphism $G(f) : G(Y) \to G(X)$ of D. Moreover, we require $G(1_X) = 1_{G(X)}$ for all X and $G(g \circ f) = G(f) \circ G(g)$.

A **contravariant functor** G from category C to category D is a function assigning each object X of C to an object $G(X)$ of D, and each morphism $f : X \to Y$ of C to a morphism $G(f) : G(X) \to G(Y)$ of D. Moreover, we require $G(1_X) = 1_{G(X)}$ for all X and $G(g \circ f) = G(g) \circ G(f)$.

EXAMPLE: Let us return to our example of Abelian groups. Let C be the category of Abelian groups, where the objects are Abelian groups, and morphisms are homomorphisms of Abelian groups. Let G be fixed, and D be the category whose objects are $Hom(A, G)$ where A is any Abelian group. Then the maps

$$A \to Hom(A, G), \quad \text{and} \quad f \mapsto \widetilde{f}$$

defines a contravariant functor from the category C to categor D.

FACT: If $f : A \to B$ is a homomorphism and is surjective, then its dual $\widetilde{f}$ is injective. That is, the exactness of

$$B \xrightarrow{f} C \to 0$$

implies the exactness of

$$Hom(B, G) \xleftarrow{\widetilde{f}} Hom(C, G) \leftarrow 0$$

More generally (p.247):

Theorem 32. *If the sequence*

$$A \xrightarrow{f} B \xrightarrow{g} C \to 0$$

is exact, then the dual is exact:

$$Hom(A, G) \xleftarrow{\widetilde{f}} Hom(B, G) \xleftarrow{\widetilde{g}} Hom(C, G) \to 0$$

Moreover, if f is injective, and the first sequence splits, then $\widetilde{f}$ is surjective and the second sequence splits.

Simplicial Cohomology Groups. Let K be a simplicial complex. G an Abelian group. The group of **p-dimensional cochains** of K, with coefficients in G, is the group

$$C^p(K; G) = Hom(C_p(K), G)$$

Thus, each cochain in C^p is a homomorphism $h : C_p(K) \to G$ where $h(\sigma) \in G$ for each simplex σ, and we extend this map linearly. NOTE: the notation $C_p(K)$ (from Munkres) is the same as $C_p(K, \mathbb{Z})$ (integral homology coefficients).

REMARK: The basic objects of homology are subspaces of the space K that we wish to study, namely C_p (think of C_p as a subspace of K). In contrast, the basic objects of cohomology are not the subspaces C_p but but functions C^p from the subspaces (to some group G)!

The **coboundary operator** δ is the dual of the boundary operator $\partial : C_{p+1}(K) \to C_p(K)$. Thus,

$$C^{p+1}(K;G) \stackrel{\delta}{\leftarrow} C^p(K;G).$$

NOTE: δ raises (!) the dimension from p to $p+1$. Let us see how δ works: suppose $\phi : C^p(K) \to G$. Then $\delta(\phi) : C^{p+1}(K) \to G$ is defined by

$$[\delta(\phi)](\sigma_{p+1}) = \phi(\partial(\sigma_{p+1}))$$

The kernel of δ $(= \delta_p)$

$$Z^p(K;G) = \ker \delta_p$$

is called the **groups of cocycles**; the image of δ is

$$B^p(K;G) = \mathrm{im}\delta_p$$

is called the **groups of coboundaries**; their quotient

$$H^p(K;G) = Z^p(K;G)/B^p(K;G)$$

is called the **cohomology group**.

Abstract Cohomology. In this case, we begin with any **cochain complex**,

$$\cdots \stackrel{\delta}{\to} C^p \stackrel{\delta}{\to} C^{p+1} \stackrel{\delta}{\to} \cdots$$

which is a sequence of Abelian groups and a boundary operator δ for successive pairs of this sequence. Note the increasing index. We only require $\delta \circ \delta = 0$. Note that such cochain complexes are normally associated with some topological space X.

Then the p**th cohomology group** of the complex is given by

$$H^p = \{\text{Kernel of}\delta : C^p \to C^{p+1}\}/\{\text{Image of}\delta : C^{p-1} \to C^p\}.$$

§16. On Permutations

Permutations is a fundamental concept in computing. It is good to develop some facility for thinking about and for manipulating permutations.

If $\{i_1, \ldots, i_m\} \subseteq \{0, \ldots, d\}$ are $m \geq 0$ distinct elements, then the sequence $[i_1 i_2 \cdots i_m]$ represents the permutation π where $\pi(k) = k$ if $k \notin \{i_1, \ldots, i_m\}$, and otherwise π takes i_1 to i_2, and i_2 to i_3, etc, and finally takes i_m to i_1. We call $[i_1 \cdots i_m]$ a **cyclic permutation** or m**-cycle**. The special cases $m = 0, 1, 2$ are noteworthy: If $m = 0$, we write the permutation as [], and it represents the identity permutation. Similarly, for $m = 1$, the permutation $[i_1]$ is also the identity permutation. For $m = 2$, $[i_1 i_2]$ is called a **transposition**. Clearly, $[i_1 i_2] = [i_2 i_1]$. In general, we have $[i_1 i_2 \cdots i_m] = [i_2 i_4 \cdots i_m i_1]$.

Consider how cyclic permutations interact under function composition: we write composition by juxtaposition: e.g., [13][14] is the composition of the transpositions [13] and [14] (we first apply the permutation [13] followed by the permutation [14]). In this case, we may verify that [13][14] = [134] $\neq$ [143] = [14][13]. Thus composition (or juxtaposition) is not commutative. Also, [13][13] = [1] = [3] = [] (the identity).

If $[i_1 \cdots i_m]$ and $[j_1, \ldots, j_n]$ are two cycles and $\{i_1, \ldots, i_m\} \cap \{j_1, \ldots, j_n\} = \emptyset$, we say the two cycles are **disjoint**. It is clear that $[i_1 \cdots i_m][j_1, \ldots, j_n] = [j_1, \ldots, j_n][i_1 \cdots i_m]$ i.e., *disjoint cycles do commute*. We may verify that in general, if $[i_1 \cdots i_m] = [i_1 i_2][i_1 i_3] \cdots [i_1 i_m]$. That is:

Every m-cycle ($m \geq 2$) can be written as a composition of $m-1$ transpositions; if the m cycle is written as a composition of k compositions then $k \geq m-1$ and k is even. (*)

Given a permutation π, we construct a cycle $[0\pi(0)\pi^2(0)\cdots\pi^k(0)]$ where $k \geq 0$ is the smallest value such that $\pi^{k+1}(0) = 0$. Next, starting from any i_1 from $\{0,\ldots,d\} \setminus \{0,\pi(0),\ldots,\pi^k(0)\}$, we can form another cycle starting from i_1. We can keep repeating this process until we find all the cycles of π. This proves:

Every permutation can be written as composition of disjoint cycles. (**)

For instance, [053][12][4] is a permutation of $\{0,\ldots,5\}$. Usually, we omit the 1-cycles in such a representation, and write [053][12] instead.

Combining (*) and (**), we conclude that every permutation can be written as a product of transpositions. Although the number of transpositions is not unique, the parity of this number is unique. Recall π of $\{0,\ldots,d\}$ into a product of compositions.

Let us now see how two transpositions in juxtaposition $[ij][i'j']$ can be transformed. There are three cases: (1) If the two transpositions are disjoint, we already know that $[ij][i'j'] = [i'j'][ij]$. (2) If they are identical, we also know that they annihilate each other: $[ij][ij] = []$. (3) Finally, assume the two transpositions shares exactly one common element i. In this case, we say the two transpositions are **linked** by the element i. In this case we may assume they have the form $[ij][ik]$ where i,j,k are distinct. In this case, we can rewrite $[ij][ik]$ as $[jk][ij]$. Think of this as "moving $[ij]$ to the right, past $[ik]$". But while moving $[ij]$ to the right, the transposition $[ij]$ survives, but the "passed over" transposition is transformed from $[ik]$ to $[jk]$ (i.e., their link element has changed from i to j). Note that the two transpositions remain linked. Consider the new pair $[jk][ji]$. If we move $[jk]$ to the right, we transform it into $[ki][jk]$, so that k is their link. In other words, we can choose their link to be any of the three elements i,j,k. To ensure canonicalness, we always make the minimum element of i,j,k the link.

In any sequence of transpositions, we can form an undirected graph whose vertices are these transpositions, and whose edges connect pairs of linked transpositions. The connected components of this graph correspond to the cycles of the graph. By repeated local transformations, we can assume that all the transpositions belonging to the same connected component are in contiguous positions. Note that local transformations do not affect these graphs (although the set of transpositions may change). Further, by local transformations, we may assume that ALL the transpositions of the connected component are linked to each other via a single element. We may assume the element is the least element in the set of elements of that component. When this happens, we can replace the component by a single cycle. In this way, we get a set of disjoint cycles. Clearly, this cycle is unique.

Given a juxtaposition of transpositions, we may simplify it as follows: (a) Suppose there are two identical transpositions, $\cdots[ij]\cdots[ij]\cdots$. We can move the leftmost copy $[ij]$ to the right until it is next to another copy of $[ij]$. Then we can annihilate them. So assume there are no identical transpositions. (b) Suppose there is a triple of transpositions, $\cdots[ij]\cdots[jk]\cdots[k\ell]\cdots$, where i,j,k,ℓ are distince, the first two are linked by j and the next two are linked by k. Then...

§17. Assigned Homeworks

1. Sep 12. Compute the homology groups of S^n (the n-spheres).

2. Sep 12. Give the minimum triangulation of the torus T^2. Assume "minimum" means the minimum number of vertices.

   HINTS: Use the standard representation of T by a square rectangular region in which the opposite edges are identified. What to look out for? You must remember that in a triangular, any two vertices determine at most one edge, any three vertices determine at most one face. Also, Euler's characteristic for a torus says that $v - e + f = 0$. You will need another relation involving v, e, f.

3. Sep 19. Show that the naive SNF algorithm (see our notes) can be exponential time. Exponential means that it is exponential at least one of the parameters m, n, L where the matrix is $m \times n$ and L is the maximum bit size of the entries.

4. Sep 19. Implement the "naive" algorithm for computing SNF. Assume that on input matrix A, you output (S, U, V) where $S = UAV$, S is in SNF, and U, V are unimodular. Test your program with any of the examples from Vegter/Rote.

   HOW TO DO THIS HOMEWORK: you must download the Core Library (http://cs.nyu.edu/exact/). Please use Core 1.7x (not the pre-release version Core 2.0). In the Core Library, we have simple extensions (COREX). You need to use the Linear Algebra extension to use matrices and perform basic operations.

   You must do this homework in pairs (this is called Xtreme Programming, which has many benefits). Consult me if you have questions about teams. Also see the introduction of these notes for possible use of CYGWIN. All the software you need can be obtained through CYGWIN.

   How to submit: I suggest you include a simple Makefile (there are many examples in Core Library), and tar everything into one file to send me. For basic information about Makefiles, see my web notes `http://cs.nyu.edu/~yap/prog/make/`. I would like to type "make" in order to compile your program, and "make snf" to see it executing on an example. To ensure that your Makefile can be directly used by me, I want you to put all your homework 1 files inside a subdirectory named $CORE/progs/your-name/hw1/. Similarly for homework 2, etc.

5. Oct 2. Recall the standard torus T^2 used in Morse theory. One chart for T^2 has been given by Vegter/Rote: it is $\varphi : U \to T^2$ where $U = (0, 2\pi) \times (0, 2\pi)$, and

$$\varphi(u, v) = (r \sin u, (R - r \cos u) \sin v, (R - r \cos u) \cos v).$$

   (a) Give other charts so as to cover the rest of T^2. How many additional charts do you need?
   (b) Let h be the usual height function on T^2. Compute the gradient field of h.
   (c) Consider another different height function $f(\varphi(u, v)) = r \sin u$. Compute the gradient field of f.
   (d) Is f a Morse function?

6. Oct 2. Let $F(X, Y) \in \mathbb{Z}[X, Y]$ and consider the curve $M : F(X, Y) = c$ for some rational $c \in \mathbb{Q}$.
   (a) Describe how you can detect whether M is a smooth manifold. REMARKS: this is a computational problem. Although we do not ask for an implementation, you must describe details at the level of implementing this in Core Library. We view M "geometrically", simply as the set of real solutions of $F(X, Y) - c$. Thus we may assume that $F(X, Y)$ is squarefree. Moreover, we shall define smoothness to mean that $F(X, Y) = F_X(X, Y) = F_Y(X, Y) = 0$ has no real solution.
   (b) Let M be smooth from (a). For $p_0 \in \mathbb{Q}^2$, define the function $f : M \to \mathbb{R}$ where $f(q) = \|p_0 - q\|$ (Euclidean distance). Describe how to test whether f is Morse.
   (c) Let f be Morse from (c). Describe how to compute the critical points of f and to determine the index of each critical point.

7. Oct 10. Let $A, B \in \mathbb{D}[X_1, \ldots, X_r]$ be homogeneous with $\deg(A) = m, \deg(B) = n$. If $\deg_i(A)$ denotes the degree of A in X_i, let $\deg_1(A) = m', \deg_1(B) = n'$. Write $\mu = m - m'$ and $\nu = n - n'$. Thus A, B are regular iff $\mu = \nu = 0$. We have the following generalization of a theorem in the text: if $\mathtt{res}_1(A, B) \neq 0$ then $\mathtt{res}_1(A, B)$ is homogeneous of degree $mn - \mu\nu$.

8. Oct 10. Implement the following algorithm: given an interval $I = [a, b]$, and a polynomial $A(X)$, to use Descartes' rule of sign to determine an upper bound on the number of zeros of A in I. Use this algorithm to isolate all the zeros of A in any given interval. NOTE: efficiency is not our primary concern in this exercise.

9. Oct 17. Implement the Descartes method for isolating the roots of a real polynomial. The input to your algorithm should be a square-free polynomial $A(X) \in \mathbb{Z}[X]$ and an interval $I = (a, b)$. If I contains k roots of A, your algorithm return a list of isolating intervals $(J_1, \ldots, J_k)$ for these roots.

CONSIDERATIONS: In Core Library, we have a facility for inputing polynomials. The simplest way to input them is to use the string representation E.g., "3 X^4- 17 X^2 - X +11". Please look at examples under the directory $(CORE)/progs/poly/. Intervals should be a pair of big floats. For bisection, do not forget to check for the midpoint being a root.

10. Oct 24. Implement in Core Library the algorithms for performing the four arithmetic operations, and also comparison of two real algebraic numbers. REMARK: you need not implement resultants of polynomials, which is already available in Core. Note that you should reuse the work in the previous exercise on Descartes method for root isolation.

11. Oct 30. Consider the function $f(x) = \sqrt{x}$ (with the domain of f suitably defined (either in $\mathbb{R}$ or in $\mathbb{C}$). This is an example of an algebraic function. In general, let us say that a partial real function $f : \mathbb{R} \to \mathbb{R}$ is algebraic if it satisfies a polynomial relation of the form

$$F(x, f(x)) = 0$$

for all x in the domain of f, and $F(X, Y) \in \mathbb{Z}[X, Y]$. Note that we are deliberately trying to avoid the machinery of algebraic function theory in this definition. For instance, $F(X, Y) = X - Y^2$ in the case of $f(x) = \sqrt{x}$. Let $g(y)$ be another algebraic function.
(a) Prove that the function composition $h(x) = (g \circ f)(x) = g(f(x))$ is algebraic. HINT: use resultants.
(b) Suppose the domain of $f(x)$ is an interval $[a, b]$ and f is continuous. Call $c \in [a, b]$ a regular point of f if in a sufficiently small neighborhood N of c, the function $f(x)$ is uniquely defined by knowing the data $f(c)$. More precisely, for each c' in the neighborhood of c, there is a unique choice of y such that $F(c', y) = 0$ holds. Note that by this definition, the point $c = 0$ is regular for the function $f(x) = x^{2/3}$. How to you detect if a given point c is regular or not?

CLARIFICATION: Michael Burr points out some bugs in this question. First of all, we must exclude the trivial case where $F(X, Y) \in \mathbb{Z}[X, Y]$ is identically zero. But even so, we have a counter example: let

$$\begin{aligned}
f(x) &= \begin{cases} 1 & \text{if } x = e \\ x & \text{if } x \neq e \end{cases}, \\
g(x) &= \begin{cases} \pi & \text{if } x = 1 \\ x & \text{if } x \neq 1 \end{cases}, \\
g \circ f &= \begin{cases} \pi & \text{if } x = e \\ x \text{ if } x \neq e & \end{cases}.
\end{aligned}$$

Both f and g are algebraic according to my definition, since

$$\begin{aligned}
F(x, f(x)) = 0 \quad &\text{if } F(X, Y) = (Y - 1)(Y - X), \\
G(y, g(y)) = 0 \quad &\text{if } G(Y, Z) = (Y - 1)(Z - Y).
\end{aligned}$$

Furthermore $R(X, Z) = \mathtt{res}_Y(F(X, Y), G(Y, Z))$ is identically 0. This does not necessarily prove that $g \circ f$ is not algebraic. Pf?

One solution to these counter examples is to declare the function f above is algebraic if it is a constant function $f(x) = c$ for some algebraic number c, or there exists $F(X, Y) \in \mathbb{Z}[X, Y]$ such that $F(x, f(x)) = 0$ holds for all x in the domain of f and, in addition, whenever $F(X, Y)$ is **primitive**. This means that for all factorizations $F(X, Y) = F_1(X)F_2(Y)F_3(X, Y)$, we must have F_1 and F_2 to be constants. N.B.: It is possible to check whether a given $F(X, Y)$ is primitive. Then, $f(x)$ and $g(x)$ above are no longer algebraic (pf?).

12. Oct 30. Please refer to (updated) Chapter 12 of my EGC notes on Constructive Zero Bounds. Consider the following expression
$$e = \sqrt{x} + \sqrt{y} - \sqrt{x + y + c\sqrt{x}\sqrt{y}}.$$
View this as a DAG that shares common subexpressions. (a) If x, y, c are L-bit integers, what is the BFMS bound for e?
(b) How good is this bound? In other words, construct asymptotic example that comes as close to this bound as possible.

13. Nov 7. Let $f, g \in \mathbb{Z}[X]$ be polynomials of degrees $\leq m$ and height h. Suppose $f(x) = 0$ and $g(x) \neq 0$.
(a) Give a lower bound on $|g(x)|$ in terms of m, h. HINT: you may use BFMS or Measure bounds.
(b) Let $F(X, Y) \in \mathbb{Z}[X, Y]$. Give a complete procedure to test if $F = 0$ has any solution inside a given box $B = [a, b] \times [c, d]$ where a, b, c, d are bigfloats. Conceptually, you must reduce everything to something you can actually implement using Core Library. HINT: divide into two cases, whether the curve $F = 0$ intersects the boundary of B or not.

14. Nov 7. Using Core Library, implement a simple cell-decomposition algorithm in 2-D. To keep the problem simple, suppose you are given three polynomial $f, g, h \in \mathbb{Z}[X, Y]$ and an open rectangular region $R = (a, b) \times (c, d)$, we want to compute a sign-invariant decomposition K of $P = \{f, g, h\}$. In order to allow different test inputs, we want the inputs to be read from a file. The output is supposed to be the incidence graph of K. Moreover, we want some to be able to display the curves (so we can verify your decomposition).

    NOTES: you may use the tools already implemented in Core. Under ${CORE}/progs/curves/, you will see that we have the ability to roughly trace implicit curves. You should also be able to display the curves (try "make show" in this directory). You should prepare several input files, and we should be able to type "make" to compile your program, "make one" to test your first input file, "make two" to test the second input file, etc.

15. Nov 21. Let $f : \mathbb{R}^2 \to \mathbb{R}$ be a polynomial function, and $\square\nabla f(B)$ denote the interval analogue of the gradient function $\nabla f(x, y)$. Plantinga-Vegter claimed that if $0 \notin \langle \square\nabla f(B), \square\nabla f(B) \rangle$, then f is parametrizable in the X- or Y-direction.

16. Nov 28.

    Vegter-Plantinga Algorithm. The correctness of this algorithm seems to leave a lot to be proved.
    (i) Prove that each leaf box of the quadtree has either 0, 2 or 4 vertices.
    (ii) Prove that the curve is isotopic to the polygonal approximation.

References

[1] S. S. Abhyankar. *Algebraic Geometry for Scientists and Engineers*. Americal Mathematical Society, Providence, Rhode Island, 1990.

[2] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in Real Algebraic Geometry*. Algorithms and Computation in Mathematics. Springer, 2003.

[3] J. Boissonnat and M. Teillaud, editors. *Effective Computational Geometry of Curves and Surfaces*. Number 59 in Mathématiques et Applications. Springer, 2006. To appear.

[4] C. Delfinado and H.Edelsbrunner. An incremental algorithm for Betti numbers of simplicial complexes on the 3-sphere. *Computer Aided Geom. Design*, 12:771–784, 1995.

[5] J. R. Munkres. *Topology: A First Course*. Prentice-Hall, Inc, 1975.

[6] H. Ratschek and J. Rokne. *Computer methods for the range of functions*. Horwood Publishing Limited, Chichester, West Sussex, UK, 1984.

[7] C. K. Yap. *Fundamental Problems of Algorithmic Algebra*. Oxford University Press, 2000.