Website. Further information about The HCMR can be found online at the journal's website,

http://hcs.harvard.edu/hcmr (1)

Instructions for Authors. All submissions should include the name(s) of the author(s), institutional affiliations (if any), and both postal and e-mail addresses at which the corresponding author may be reached. General questions should be addressed to Editor-In-Chief Scott Kominers at hcmr@hcs.harvard.edu.

Articles. The Harvard College Mathematics Review invites the submission of quality expository articles from undergraduate students. Articles may highlight any topic in undergraduate mathematics or in related fields, including computer science, physics, applied mathematics, statistics, and mathematical economics.

Authors may submit articles electronically, in .pdf, .ps, or .dvi format, to hcmr@hcs.harvard.edu, or in hard copy to

The Harvard College Mathematics Review
Student Organization Center at Hilles
Box # 360
59 Shepard Street
Cambridge, MA 02138.

Submissions should include an abstract and reference list. Figures, if used, must be of publication quality. If a paper is accepted, high-resolution scans of hand drawn figures and/or scalable digital images (in a format such as .eps) will be required.

Problems. The HCMR welcomes submissions of original problems in all mathematical fields, as well as solutions to previously proposed problems.

Proposers should send problem submissions to Problems Editor Zachary Abel at hcmr-problems@hcs.harvard.edu or to the address above. A complete solution or a detailed sketch of the solution should be included, if known.

Solutions should be sent to hcmr-solutions@hcs.harvard.edu or to the address above. Solutions should include the problem reference number. All correct solutions will be acknowledged in future issues, and the most outstanding solutions received will be published.

Advertising. Print, online, and classified advertisements are available; detailed information regarding rates can be found on The HCMR's website, (1). Advertising inquiries should be directed to hcmr-advertise@hcs.harvard.edu, addressed to Business Manager Charles Nathanson.

Subscriptions. One-year (two issues) subscriptions are available, at rates of $10.00 for students, $15.00 for other individuals, and $30.00 for institutions. Subscribers should mail checks for the appropriate amount to The HCMR's postal address; confirmation e-mails should be directed to Distribution Manager Nike Sun, at hcmr-subscribe@hcs.harvard.edu.

Sponsorship. Sponsoring The HCMR supports the undergraduate mathematics community and provides valuable high-level education to undergraduates in the field. Sponsors will be listed in the print edition of The HCMR and on a special page on the The HCMR's website, (1). Sponsorship is available at the following levels:

Sponsor	$0 - $99
Fellow	$100 - $249
Friend	$250 - $499
Contributor	$500 - $1,999
Donor	$2,000 - $4,999
Patron	$5,000 - $9,999
Benefactor	$10,000 +

Friends · D. E. Shaw & Co. · **Contributors** · QVT Financial LP · **Patrons** · The Harvard University Mathematics Department

Cover Image. The image on the cover was created using circle inversion and is based a method described in "Problems of Circle Tangency," by Gregory Minton (page 47). The image was created in Mathematica™ by Graphic Artist Zachary Abel.

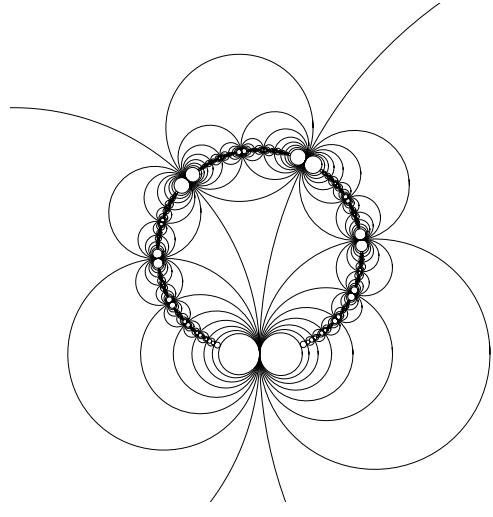

The Harvard College Mathematics Review is produced and edited by a student organization of Harvard College.

-2
Contents

-1 Staff

Editor-In-Chief
Scott Kominers '09

Design Director
Brett Harrison '10

Business Manager
Charles Nathanson '09

Articles Editor
Shrenik Shah '09

Distribution Manager
Nike Sun '09

Features Editor
Sam Lichtenstein '09

Graphic Artist
Zachary Abel '10

Problems Editor
Zachary Abel '10

Cover and Logo Design
Hannah Chung '09

Issue Production Directors
Zachary Abel '10
Menyoung Lee '10
Daniel Litt '10

Webmaster
Brett Harrison '10

Board of Reviewers
Zachary Abel '10
Pablo Azar '09
Connie Chao '08
Justin Chen, Caltech '09
Hannah Chung '09
Kelley Harris '09
Brett Harrison '10
Scott Kominers '09
Menyoung Lee '10
John Lesieutre '09
Sam Lichtenstein '09
Alison Miller '08
Charles Nathanson '09
Shrenik Shah '09
Nike Sun '09
Arnav Tripathy '11

Board of Copy Editors
Zachary Abel '10
Pablo Azar '09
Eleanor Birrell '09
Hannah Chung '09
Grant Dasher '09
Ernest E. Fontes '10
Sherry Gong '11
François Greer '11
Kelley Harris '09
Scott Kominers '09
Menyoung Lee '10
Sam Lichtenstein '09
Daniel Litt '10
Richard Liu '11
Charles Nathanson '09
Shrenik Shah '09
Nike Sun '09
Arnav Tripathy '11
Xiaoqi Zhu '11

Business Board
Zachary Abel '10
Scott Kominers '09
Stella Lee '09
Sam Lichtenstein '09
Daniel Litt '10
Charles Nathanson '09
Shrenik Shah '09
Nike Sun '09

Faculty Advisers
Professor Benedict H. Gross '71, Harvard University
Professor Peter Kronheimer, Harvard University
Dr. Alon Amit, Google
Professor Matthew Steven Carlos, Europäische Universität für Interdisziplinäre Studien

0
From the Editor

Scott Kominers
Harvard University '09
Cambridge, MA 02138
kominers@fas.harvard.edu

> "Man knows how many seeds are in an apple, but only heaven knows how many apples are in a seed."

When we launched *The Harvard College Mathematics Review* (HCMR), we had no idea how many apples would be in the seed.

The response to our first issue has been tremendous. We have received response letters and e-mails from students and faculty at schools spanning not only the United States but the world. Some articles in The HCMR have already been cited professionally; others have been translated into multiple languages.

This summer, at the Mathematical Association of America's *MathFest 2007*, I witnessed first-hand my peers' excitement at viewing the first issue. At *MathFest*, I met math professors who had already shared the first issue's articles with their students—and met students who had already read that first issue and asked their professors to suggest follow-up reading.

Indeed, mathematics is lucky. Young mathematicians are learning and working all over the world—and experienced mathematicians are working hard to teach the new generation.

Just over half of the student articles and problems in this issue came from outside Harvard's walls. Furthermore, the contributors are strikingly diverse: Two of our student authors attend Harvey Mudd College; another is a high school student. Some student problem proposers come from institutions as near to Harvard as the Massachusetts Institute of Technology, and others come from as far away as Romania and Germany. This issue's faculty feature was contributed by Professor Emeritus Elemér E. Rosinger of the University of Pretoria, South Africa, who brought the above aphorism to my attention.

All of these authors have helped us grow a mathematical tree which bears fruit across the world. We are grateful to all of them for their participation and contribution.

As always, we appreciate your commentary and feedback. Please direct your comments and questions to hcmr@hcs.harvard.edu or to me personally at kominers@fas.harvard.edu. We also invite you to submit to future issues. We publish articles, short notes, and problems in any field of pure or applied mathematics at the undergraduate level. Submission guidelines and instructions can be found on the inside front cover.

We especially appreciate the helpful commentary and assistance of our faculty sponsors and advisers, **Professor Benedict H. Gross '71**, **Professor Peter Kronheimer**, **Dr. Alon Amit**, and **Professor Matthew Steven Carlos**. We also owe special thanks to **Professor Clifford H. Taubes** for his continued support and encouragement. We are also grateful to **Dean Paul J. McLoughlin II** and **Mr. David R. Friedrich** for their administrative and organizational assistance. Finally, we could never have produced this issue without the continued, generous support of **The Harvard Mathematics Department**.

It is the honor and joy of everyone at The HCMR to continue to plant seeds. We hope that you enjoy our second issue, $\mathbb{QED}$.

Scott Kominers '09
Editor-In-Chief, The HCMR

STUDENT ARTICLE

1 Determining the Genus of a Graph

Andres Perez[†]
Harvey Mudd College '09
Claremont, CA 91711
aperez@hmc.edu

Abstract

This paper investigates an important aspect of topological graph theory: methods for determining the genus of a graph. We discuss the classification of higher-order surfaces and then determine bounds on the genera of graphs embedded in orientable surfaces. After generalizing Euler's Formula to include graphs embedded on these surfaces, we derive upper and lower bounds for the genera of various families of simple graphs. We then examine some formulas for the genera of particular graphs.

1.1 Introduction

A graph G is **planar** if and only if it can be drawn in the plane such that none of its edges cross. Two examples of non-planar graphs are the complete graph on five vertices K_5 and the complete bipartite graph $K_{3,3}$. A **complete graph** consists of a set of completely connected vertices; a **complete bipartite graph** consists of two independent sets of vertices in which all the vertices in one group are connected to all the vertices in the other group. In 1930, Kazimierz Kuratowski [Ku] arrived at a result that is now known as **Kuratowski's Theorem**:

Theorem 1 (Kuratowski's Theorem)**.** *A graph G is planar if and only if G does not contain a subdivision of K_5 or $K_{3,3}$.*

This characterization of planar graphs leads to a natural question: Does there exist a surface upon which these non-planar graphs can be embedded such that there are no edge crossings? From a topological viewpoint, drawing a graph on a flat plane is equivalent to drawing the same graph on a sphere. We can verify this by taking the stereographic projection of a sphere, i.e. we can unravel the surface of a sphere by creating a hole at the north pole and stretching the sphere's surface out onto the plane (the pole of the sphere can be placed inside some region of the graph and this becomes the "outer face" of the graph). Thus, a graph is planar if and only if it can be drawn on the sphere in a way such that no edges cross.

Supposing we had one edge crossing in a graph G, we could draw G without any edges crossing by introducing a "handle" to our sphere (which is the topological equivalent to the torus–see Figure 1.1), and drawing one of the edges over the handle so that it no longer crosses the other edge. In this way, we could properly embed the graph $K_{3,3}$ on the torus without any edges crossing (see Figure 1.2).

The torus is an example of a **surface of higher genus**. The sphere is designated to be the surface S_0; the surface formed by adding k handles to the sphere is denoted S_k. The torus is therefore S_1, the double-torus S_2, and so on, where the **genus** of S_k is the number of handles, k.

[†]Andres Perez, Harvey Mudd College '09, is a mathematics major living in Harvey Mudd's West Dorm in Claremont, CA. His family is originally from Peru and he was born in Lake Forest, IL, where he attended Lake Forest Academy. His mathematical interests include graph theory, combinatorics and algebra; he also appreciates theoretical computer science. Beyond math, his interests include poker, photography, and rock climbing.

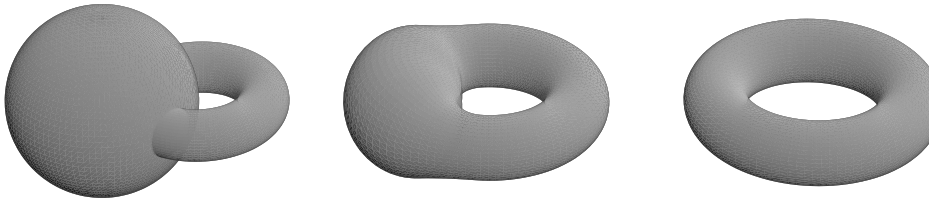

Figure 1.1: Introducing a handle to a sphere is topologically equivalent to a torus.

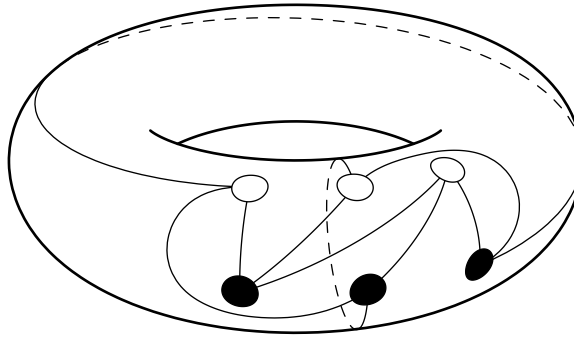

Figure 1.2: The graph $K_{3,3}$ drawn without edge crossings on the torus.

The main purpose of this paper is to determine, for a graph G, the minimal k such that G can be drawn on S_k without edge crossings.

We see that one upper limit to this number is the number of crossings in a drawing of the graph on S_0: we simply introduce a handle for each instance of two edges crossing. But what is the lowest-genus surface required? We define this number to be the **genus** of a graph G (plural, **genera**), denoted $\gamma(G)$. Thus for any planar graph G, $\gamma(G)$ is zero. Both $K_{3,3}$ and K_5 are of genus 1.

Why might we be interested in knowing the genus of a graph? Such a question might arise in circuit design, for example, if we wanted to print electronic circuits on a circuit board to minimize crossings that could result in a short circuit. In this paper, we discuss how graphs embedded on surfaces of higher genus can be represented on the plane, how Kuratowski's Theorem can be extended to higher order surfaces, how Euler's Formula can be modified to account for the genus of the graph, and how we can determine the genus of a graph.

1.2 Surfaces of Higher Order

Since graphs drawn on three-dimensional surfaces are hard to work with, we would like to determine ways to draw graphs on 2-dimensional surfaces. If we consider a torus, we can create a two-dimensional representation of the surface by slicing the handle of the torus to create a cylinder and then slicing the cylinder lengthwise to create a rectangle (see Figure 1.3).

This rectangle has the property that if an edge stretches out to a side it continues back from the other side of the rectangle. Figure 1.4 shows an embedding of K_5 on this rectangle with no crossing edges. Many graphs can be drawn in S_1 with a very high level of symmetry. Figure 1.6 demonstrates two planar embeddings constructed in S_1: $K_{4,4}$ and K_7.

The surfaces that can be created by introducing handles to a sphere are all **orientable**. Intu-

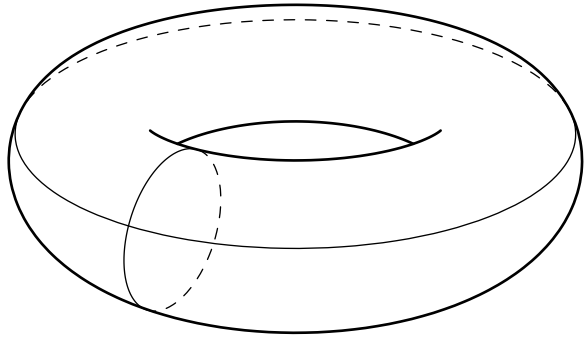

Figure 1.3: A torus representing the cuts required to transform the surface into a rectangle. The vertical cut produces a cylinder, and then with the horizontal cut the surface becomes a rectangle.

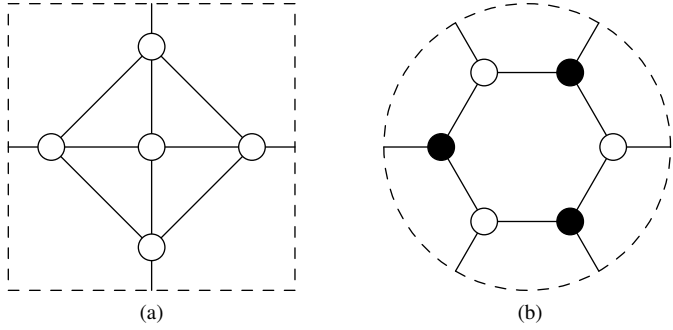

Figure 1.4: (a) A representation of K_5 embedded on the surface S_1, and (b) a representation of $K_{3,3}$ embedded on the projective plane.

itively, a surface is said to be orientable if it has two distinct "sides."[1] There are many interesting properties of **non-orientable surfaces** such as the Möbius strip N_1, and the Klein Bottle N_2, (see Figure 1.5), but the remainder of this paper will only address the orientable surfaces $S_0, S_1, \ldots$.[2] The graphs K_5 and $K_{3,3}$ are called **forbidden minors** in S_0. Formally, a forbidden minor in S_k is

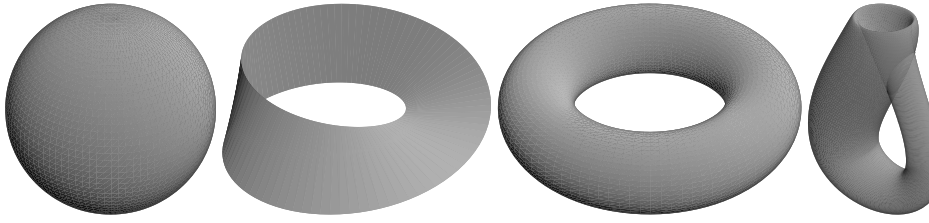

(a) (b) (c) (d)

Figure 1.5: Representations of some common surfaces in 3-space: (a) The sphere, S_0; (b) The Möbius strip, N_1; (c) the torus, S_1; and (d) the Klein Bottle, N_2.

a graph G not embeddable in S_k with the property that, if we were to remove any edge from G, we would be left with a graph that is embeddable in S_k. By Kuratowski's Theorem, K_5 and $K_{3,3}$ are the unique forbidden minors in S_0.[3]

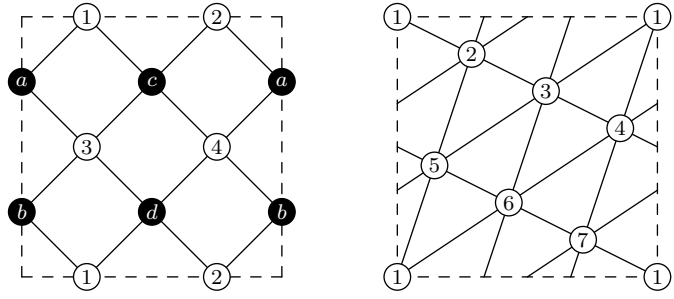

(a) (b)

Figure 1.6: (a) A planar embedding of $K_{4,4}$ in S_1, and (b) a planar embedding of K_7 in S_1.

1.3 Euler's Formula Extended

Euler's Formula states that for a planar embedding of a graph with V vertices, E edges, and F faces (or regions), we have the relation

$$V - E + F = 2.$$

[1] For example, the Möbius strip is not an orientable surface as it only has one "side."

[2] However, we will mention one more noteworthy, non-orientable representation: the **projective plane** is a non-orientable surface that allows for a fairly straightforward embedding of the graph $K_{3,3}$ (see Figure 1.4b).

[3] It is a good exercise to try to find all forbidden minors in the space S_1. The reader may find it surprising to learn that there are in fact over 800 forbidden minors in S_1 alone [We], one example being the graph $2K_4 + K_1$, which is simply two copies of the complete graph on four vertices with an extra vertex adjacent to all the other vertices.

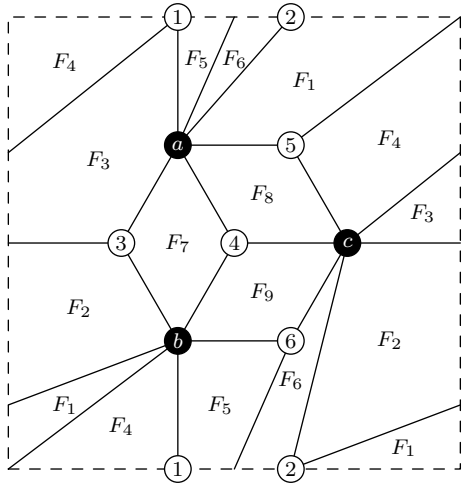

Figure 1.7: A planar embedding of $K_{3,6}$ in S_1. Note that for this embedding, $V - E + F = 0$.

If we examine a planar embedding of the graph $K_{3,6}$ on the surface S_1 (see Figure 1.7), we find that the genus-1 analogue of Euler's Formula is $V - E + F = 0$. This motivates the derivation of one of the most significant theorems in topological graph theory, an extension of Euler's Formula for higher-order surfaces:

$$V - E + F = 2 - 2g,$$

where g is the genus of the surface the graph is embedded upon, and the quantity $2 - 2g$ is defined to be the **Euler characteristic** $\chi(G)$ of the graph G.[4] Before proving this formula, we introduce some basic terminology; cf. [Wh].

Definition 2. A **pseudograph** is a graph with loops and multiple edges allowed.

Definition 3. A region of an embedding of a graph G in a surface M is said to be a **2-cell** if it is homeomorphic to the open disk. If every region of an embedding is a 2-cell, the embedding is said to be a **2-cell embedding**.

Here, when we say that a region is "homeomorphic to the open disk," we mean that it is topologically equivalent to a flat disk. If we were to shrink the the face of such a region to a point, we would find that it does not contain any irregularities such as handles or holes.

Theorem 4 (Euler's Formula)**.** *Let G be a connected pseudograph, with a 2-cell embedding in S_g, with the usual parameters V, E, and F. Then*

$$V - E + F = 2 - 2g. \tag{1.1}$$

Proof. We argue by induction on g. For the base case, we suppose $g = 0$. This reduces to the formula for planar graphs in S_0, $V-E+F = 2$, which we can prove by induction on the number of edges. If there are no edges, then G is an isolated vertex, and therefore $V-E+F = 1-0+1 = 2$. Otherwise, choose any edge e. If e is a loop, then remove it and E and F decrease by 1. If e connects two different vertices, contract e to a point and V and E each decrease by 1. In either case, the result follows by induction.

[4]Note that the Euler characteristic $\chi(G)$ should not be confused with the **chromatic number**, the least number of colors needed to color the vertices of a graph so that no adjacent vertices are of the same color.

Having shown the base case, we now assume the theorem holds true for graphs of genus $g-1$. We wish to show that a connected pseudograph G with a 2-cell embedding in S_g satisfies the formula. Let G be the graph of interest, with parameters V, E, and F.

Since the embedding is a 2-cell, this means that each face must have the property that it can be shrunk down to a point. Therefore if a face were to contain a handle, it would not be a proper 2-cell embedding. This implies that every handle must have at least one edge through it (see Figure 1.8).

Select one handle, and draw two closed curves C_1 and C_2 around the handle (by "around", we mean that if the surface were a coffee mug, you would draw out a curve by wrapping your hand around the handle) such that edges that intersect C_1 intersect C_2, and vice versa; this is always possible.

Suppose edges $e_1, e_2, \ldots, e_n$ run over the handle. Let x_{ij} be the point where curve C_i meets edge e_j, where $1 \leq i \leq 2$ and $1 \leq j \leq n$. Consider the points x_{ij} to be the vertices of a new pseudograph, whose edges consist of the appropriate subdivisions of the original edges, as well as the edges formed along the curves. Call this new graph G'; the graph G' is an extension of the graph G. It also includes the old edges and vertices outside the handle under consideration. We now have the parameters:

$$\begin{aligned} V' &= V + 2n, \\ E' &= E + 4n, \\ F' &= F + 2n. \end{aligned}$$

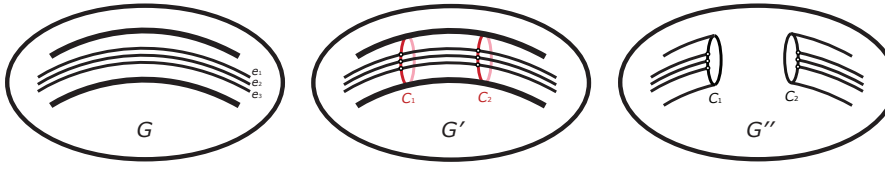

Figure 1.8: The handle under consideration shown in the graphs G, G', and G''.

Now, remove the portion of the handle between C_1 and C_2 and "fill in" the two resulting holes with two disks (the disks can be thought of as "caps"). Note that the edges that formed the curve are still present, and the caps they enclose account for two new faces. The result is the 2-cell embedding shown in Figure 1.8 of a connected pseudograph G'' in S_{g-1}, with the following parameters:

$$\begin{aligned} V'' &= V' = V + 2n, \\ E'' &= E' - n = E + 3n, \\ F'' &= F' - n + 2 = F + n + 2. \end{aligned}$$

By the inductive hypothesis, we now have:

$$\begin{aligned} 2 - 2(g-1) &= V'' - E'' + F'' \\ &= (V + 2n) - (E + 3n) + (F + n + 2) \\ &= V - E + F + (2n - 3n + n) + 2 \\ &= V - E + F + 2, \end{aligned}$$

the desired result. □

1.4 Determining the Maximum Genus of a Graph

Euler's Formula is a useful technique for finding upper and lower bounds for the genus of a graph. An important value used to develop many of the upper bounds for graphs is the **maximum genus**

$\gamma_M(G)$ of a graph.

Definition 5. The **maximum genus** $\gamma_M(G)$ of a connected graph G is the maximum genus among the genera of all surfaces in which G has a 2-cell embedding.

This maximum exists because for a graph to have a proper 2-cell embedding on a surface S_k, it must have at least one edge crossing each of the k handles. This leads to the (very loose) bound of $\gamma_M(G) \leq e + 1$, where e is the number of edges in G. We should also consider the following theorem of Duke [Du] which gives a deeper understanding of a graph's ability to be embedded in the surface S_k:

Theorem 6 (Duke)**.** *A connected graph G has a 2-cell embedding in S_k if and only if $\gamma(G) \leq k \leq \gamma_M(G)$.*

Theorem 6 tells us that if we can embed a graph into both a surface of genus n and into a surface of genus $m > n$, then we can embed the same graph onto any surface of genus g, where $n \leq g \leq m$. In developing upper bounds for the maximum genus, for the sake of algebra, it is helpful to introduce the following construction:

Definition 7. The **Betti number** $\beta(G)$ of a graph G having v vertices, e edges, and m components, is given by: $\beta(G) = e - v + m$.

Hence, $\beta(G) = e-v+1$ for any connected graph G. We can now derive an upper bound for the genus of a connected graph using Euler's Formula: Let G be connected, with a 2-cell embedding in S_k. Then $f \geq 1$, and:

$$\begin{aligned} k &= 1 + \frac{e - v - f}{2} \\ &= \frac{e - v - (f - 2)}{2} \\ &\leq \frac{e - v + 1}{2} \\ &= \frac{\beta(G)}{2}. \end{aligned}$$

We then have the bound,

$$\gamma_M(G) \leq \left\lfloor \frac{\beta(G)}{2} \right\rfloor. \tag{1.2}$$

Graphs for which $\gamma_M(G) = \left\lfloor \frac{\beta(G)}{2} \right\rfloor$ are said to be **upper embeddable**. It has been shown that all complete n-partite graphs (graphs consisting of n completely interconnected, independent sets) are upper embeddable [KRW]. Using bound (1.2) for $\gamma_M(G)$, we can derive upper bounds for the maximum genera of the complete graphs K_n, complete bipartite graphs $K_{m,n}$, and the n-cube Q_n ($n \geq 2$).[5] In particular,

$$\gamma_M(K_n) \leq \left\lfloor \frac{(n-1)(n-2)}{4} \right\rfloor, \tag{1.3}$$

$$\gamma_M(K_{m,n}) \leq \left\lfloor \frac{(n-1)(m-1)}{2} \right\rfloor, \tag{1.4}$$

$$\gamma_M(Q_n) \leq (n-2)2^{n-2}. \tag{1.5}$$

The above formulas can be proven to be equalities for the maximum genera of complete graphs [NSW], complete bipartite graphs [R1], and the n-cube [Za].

[5]The **n-cube** is the simple graph whose vertices are the k-tuples with entries in $\{0, 1\}$ and whose edges are the pairs of k-tuples that differ in exactly one position. For example, Q_2 has the structure of a square, and Q_3 is the cube.

1.5 Lower Bounds for the Genus of a Graph

We now investigate a technique for computing lower bounds for the genera of some simple graphs. Recall that the **degree** of a vertex is the number of adjacent vertices and the **length** of a region is the length of the closed path that bounds the region. Let v_i be the number of vertices of degree i, and let f_i be the number of regions of length i. If we focus only on 2-cell embeddings of graphs with minimum degree 3 ($\delta(G) \geq 3$), also called **polyhedral** graphs, it follows that $v = \sum_{i\geq 3} v_i$ and $f = \sum_{i\geq 3} f_i$. By the well-known **degree-sum formula**, the sum of the degrees of all the vertices is equal to twice the number of edges in a graph. We then have:

$$2e = \sum_{i\geq 3} i \cdot v_i. \tag{1.6}$$

Also, since each edge separates two regions or belongs twice to a single region, summing the sides of each face double-counts the edges, whereby we have:

$$2e = \sum_{i\geq 3} i \cdot f_i. \tag{1.7}$$

The above results hold for all polyhedral graphs, which include the complete graphs on n vertices for $n \geq 3$, the complete bipartite graphs $K_{m,n}$ for $m, n \geq 3$, and the n-cubes Q_n for $n \geq 3$. Thus, for all polyhedral graphs, $2e = \sum_{i\geq 3} i \cdot f_i \geq \sum_{i\geq 3} 3 \cdot f_i = 3f$, and therefore $f \leq \frac{2}{3}e$. Also, for all triangle-free polyhedral graphs, $2e = \sum_{i\geq 4} i \cdot f_i \geq \sum_{i\geq 4} 4 \cdot f_i = 4f$, whence $f \leq \frac{1}{2}e$.

Using these two inequalities in conjunction with Euler's Formula, we can obtain lower bounds for all polyhedral graphs and triangle-free polyhedral graphs. First consider the former, where we have $f \leq \frac{2}{3}e$. Using Euler's Formula and solving for g, we find:

$$\begin{aligned} g &= 1 - \frac{v}{2} + \frac{e}{2} - \frac{f}{2} \\ &\geq 1 - \frac{v}{2} + \frac{e}{2} - \frac{1}{2}\left(\frac{2}{3}e\right) \\ &= 1 - \frac{v}{2} + e\left(\frac{1}{2} - \frac{1}{3}\right) \\ &= 1 - \frac{v}{2} + \frac{e}{6}. \end{aligned}$$

Thus, we have the bound

$$\gamma(G) \geq \left\lceil 1 - \frac{v}{2} + \frac{e}{6} \right\rceil. \tag{1.8}$$

We can also develop the bound for triangle-free graphs in the same way to obtain:

$$\gamma(G) \geq \left\lceil 1 - \frac{v}{2} + \frac{e}{4} \right\rceil. \tag{1.9}$$

This gives that

$$\gamma(K_n) \geq \left\lceil 1 - \frac{n}{2} + \frac{n(n-1)}{12} \right\rceil \tag{1.10}$$

$$= \left\lceil \frac{(n-3)(n-4)}{12} \right\rceil. \tag{1.11}$$

In fact [Ha, p. 118-119], we have equality, i.e.,

$$\gamma(K_n) \quad = \quad \left\lceil \frac{(n-3)(n-4)}{12} \right\rceil, \qquad n \geq 3. \tag{1.12}$$

An interesting result to note is that by the inequality (1.3), which is in fact an equality (cf. [NSW]) for large n, $\gamma_M(K_n) \to 3\gamma(K_n)$; this gives a range for the possible surfaces on which a 2-cell embedding of K_n can exist.

Similarly, we have the formula

$$\gamma(K_{m,n}) \quad = \quad \left\lceil \frac{(m-2)(n-2)}{4} \right\rceil, \qquad m, n \geq 2. \tag{1.13}$$

We can determine a lower bound for the genus of the n-cube Q_n in the same fashion, by using the triangle-free inequality (1.9):

$$\gamma(Q_n) \quad \geq \quad 1 + (n-4)2^{n-3}, \qquad n \geq 2. \tag{1.14}$$

It turns out that this is an equality. This was proven to be the genus of the n-cube Q_n in [R2]. The proof involves induction and a more complex technique called "surgery" on the graph [GT] which unfortunately is beyond the scope of this paper.

1.6 Conclusion

We have given several non-trivial bounds on the genera of certain families of graphs, as well as explicit formulas for a few highly symmetric families of graphs (namely the complete and complete bipartite graphs). Determining the genera of more complicated families of graphs usually involves calculating lower bounds using Euler's Formula and then deriving (often using induction) a general construction of an embedding using techniques such as surgery.

One topic we have omitted entirely is the question of finding algorithms to determine the genus of a graph. However, this is a very interesting subject; it has been proven that there exists a linear time algorithm which finds an embedding of G in a surface S, or if this is impossible, finds a subgraph $K \subseteq G$ which is a subdivision of some forbidden minor of S (see [M1]). A general formula for the genus of an arbitrary graph is not known, but using the techniques discussed in this paper, many equalities can be constructed for n-partite graphs.

References

[BL] C. Paul Bonnington and Charles H. C. Little: *The Foundations of Topological Graph Theory*. New York: Springer-Verlag, 1995.

[Du] Richard A. Duke: The genus, regional number and Betti number of a graph, *Canad. J. Math* **18** (1966), 817–822.

[GT] Jonathan L. Gross, Thomas W. Tucker: *Topological Graph Theory*. New York: John Wiley & Sons, 1987.

[Ha] Frank Harary: *Graph Theory*. Redding, Massachusetts: Addison Wesley Publishing Company, 1969.

[KRW] Hudson V. Kronk, Richard D. Ringeisen, and Arthur T. White: On 2-cell imbeddings of complete n-partite graphs, *Colloq. Math.* **36** (1976), 131–140.

[Ku] Kazimierz Kuratowski: Sur le problème des courbes gauches en topologie (in French), *Fund. Math.* **15** (1930), 271–283.

[M1] Bojan Mohar: A linear time algorithm for embedding graphs in an arbitrary surface, *SIAM J. Discrete Math.* **13** (1999), 6–26.

[M2] Bojan Mohar and Carsten Thomassen: *Graphs on Surfaces*. Baltimore: Johns Hopkins University Press, 2001.

[NSW] E. A. Nordhaus, B. M. Stewart, and Arthur T. White: On the maximum genus of a graph, *J. Combinatorial Theory* B **11** (1971), 258–167.

[R1] Richard D. Ringeisen: Determining all compact orientable 2-manifolds upon which $K_{m,n}$ has 2-cell imbeddings, *J. Combinatorial Theory* B **12** (1972), 101–104.

[R2] Gerhard Ringel: Über drei kombinatorische Probleme am n-dimensionalen Würfel und Würfelgitter (in German), *Abh. Math. Sem. Univ. Hamburg* **20** (1955), 10–19.

[We] Douglas B. West: *Introduction to Graph Theory*, 2nd Edition. Upper Saddle River: Prentice-Hall, 2001.

[Wh] Arthur T. White: *Graphs of Groups on Surfaces: Interactions and Models*. New York: Elsevier Science, 2001.

[Za] J. Zaks: The maximum genus of cartesian products of graphs, *Canad. J. Math.* **26** (1974), 1025–1036.

STUDENT ARTICLE

2 The Poincaré Lemma and de Rham Cohomology

Daniel Litt[†]
Harvard University '10
Cambridge, MA 02138
dalitt@fas.harvard.edu

Abstract

The Poincaré Lemma is a staple of rigorous multivariate calculus—however, proofs provided in early undergraduate education are often overly computational and are rarely illuminating. We provide a conceptual proof of the Lemma, making use of some tools from higher mathematics.

The concepts here should be understandable to those familiar with multivariable calculus, linear algebra, and a minimal amount of group theory. Many of the ideas used in the proof are ubiquitous in mathematics, and the Lemma itself has applications in areas ranging from electrodynamics to calculus on manifolds.

2.1 Introduction

Much of calculus and analysis—the path-independence of line- or surface-integrals on certain domains, Cauchy's Theorem (assuming the relevant functions are C^1) on connected complex regions and the more general residue theorem, and various ideas from physics—depends to a large extent on a powerful result known as the Poincaré Lemma. On the way to the statement and proof of this Lemma, we will introduce the concepts of the exterior power and differential forms, as well as de Rham cohomology.

2.2 Linear Algebra and Calculus Preliminaries

2.2.1 The Exterior Power

We begin by defining some useful objects; on the way, we will digress slightly and remark on their interesting properties. We will begin by defining a vector space called the exterior power, in order to extend the notion of a determinant.

Definition 1 (Alternating Multilinear Form, Exterior Power)**.** Let V be a finite-dimensional vector space over a field F. A n**-linear form** is a map $B : \underbrace{V \times \cdots \times V}_{n} \to W$, where W is an arbitrary vector space over F, that is linear in each term, i.e. such that

$$B(a_1, a_2, \ldots, a_n) + B(a_1', a_2, \ldots, a_n) = B(a_1 + a_1', a_2, \ldots, a_n)$$

[†]Daniel Litt, Harvard '10, is a mathematics concentrator living in Leverett House. Originally from Cleveland, Ohio, he has performed research in computational biology, physics, combinatorics, game theory, and category theory. His other academic interests include epistemology and moral philosophy, as well as poetry and music. He is a founding member of The HCMR and currently serves as Issue Production Director.

and similarly in the other variables, and such that

$$s \cdot B(a_1, a_2, \dots, a_n) = B(s \cdot a_1, a_2, \dots, a_n) = B(a_1, s \cdot a_2, \dots, a_n) = \cdots$$

We say that such forms are **multilinear**.

A multilinear form B is **alternating** if it satisfies

$$B(a_1, \dots, a_i, a_{i+1}, \dots, a_n) = -B(a_1, \dots, a_{i+1}, a_i, \dots, a_n)$$

for all $1 \leq i < n$.

Then the **n-th exterior power** of V, denoted $\bigwedge^n(V)$, is a vector space equipped with an alternating multilinear map $\wedge : \underbrace{V \times \cdots \times V}_{n} \to \bigwedge^n(V)$ such that any alternating multilinear map

$$f : \underbrace{V \times \cdots \times V}_{n} \to W$$

factors uniquely through $\wedge$, that is, there exists a unique $f' : \bigwedge^n(V) \to W$ such that the diagram

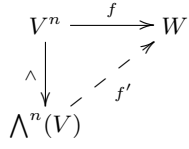

commutes, i.e. $f' \circ \wedge = f$.

It is not immediately clear that such a vector space exists or is unique. For existence, see [DF, p. 446]; the construction is not important for our purposes so we relegate it to a footnote.[1] Uniqueness follows immediately from the fact that the above definition is a **universal property**. However, we provide the following proof to elucidate this notion:

Proposition 2. *The n-th exterior power of a vector space V is unique up to canonical isomorphism.*

Proof. Consider vector spaces $\bigwedge_1^n(V)$ and $\bigwedge_2^n(V)$ with associated maps $\wedge_1$ and $\wedge_2$ satisfying the definition above. As $\wedge_1, \wedge_2$ are both alternating and multilinear, they must factor through one another; that is, we must have that there exist unique $\wedge_1', \wedge_2'$ such that the diagram

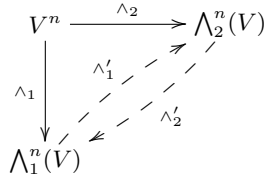

[1] We have that $\bigwedge^0(V) = F$ and $\bigwedge^1(V) = V$. There are three equivalent ways to construct the exterior power for $n \geq 2$. First, the exterior power $\bigwedge^n(V)$ can be viewed as the span of formal strings $v_1 \wedge \cdots \wedge v_n$, where $\wedge$ is a formal symbol satisfying the properties of the wedge product.

Second, for readers familiar with the tensor power, we may, for $n \geq 2$, let $I^2(V) \subset \bigotimes^2(V)$ be the subspace spanned by vectors of the form $v \otimes v$ in $\bigotimes^2(V)$ and for $n > 2$ let $I^n(V) = (V \otimes I^{n-1}(V)) \bigoplus (I^{n-1}(V) \otimes V)$; then $\bigwedge^n(V) \simeq \bigotimes^n(V)/I^n(V)$.

Finally, let $J^2(V) \subset V \otimes V$ be the subspace spanned by vectors of the form $(v \otimes w - w \otimes v)$. Then for $n \geq 2$, $\bigwedge^n(V) \simeq \bigcap_{i=0}^{n-2} \left(\bigotimes^i(V) \otimes J^2(V) \otimes \bigotimes^{n-i-2}(V) \right) \subset \bigotimes^n(V)$.

We leave checking that these constructions satisfy the definition of the exterior power (and are thus isomorphic) as an exercise; the reader may look at the given reference [DF] for the solution.

commutes. Note however that $\wedge_1, \wedge_2$ must be surjective, so $\wedge_1', \wedge_2'$ must be mutually inverse by the commutativity of the diagram above. But then $\bigwedge_1^n(V) \simeq \bigwedge_2^n(V)$ as desired, and we have uniqueness. □

In accordance with convention, for $v_1, \ldots, v_n \in V$, we define

$$v_1 \wedge \cdots \wedge v_n := \wedge(v_1, \ldots, v_n).$$

This is referred to as the **wedge product** of $v_1, \ldots, v_n$. By the fact that $\wedge$ is multilinear and alternating, note that

1. $v \wedge w = -w \wedge v$ (this immediately implies that $v \wedge v = 0$),
2. $a \cdot (v \wedge w) = (a \cdot v) \wedge w = v \wedge (a \cdot w)$,
3. $v \wedge w + v \wedge w' = v \wedge (w + w')$,

with the appropriate generalizations to higher-order exterior powers. We now calculate, for an n-dimensional vector space V, the dimension of $\bigwedge^s(V)$.

Proposition 3. *We have that*

$$\dim \bigwedge^s(V) = \binom{n}{s}.$$

In particular, given a basis $\{e_1, \ldots, e_n\}$ of V, the vectors

$$e_{i_1} \wedge \cdots \wedge e_{i_s} \text{ for } 1 \leq i_1 < i_2 < \cdots < i_s \leq n$$

form a basis of $\bigwedge^s(V)$.

Proof. We begin with the second claim. We first show that

$$\bigwedge^s(V) = \text{span}(e_{i_1} \wedge \cdots \wedge e_{i_s} \mid 1 \leq i_1 < i_2 < \cdots < i_s \leq n).$$

As $\{e_1, \ldots, e_n\}$ is a basis of V, we may write any $v_1 \wedge \cdots \wedge v_s \in \bigwedge^s(V)$ as

$$\left(\sum_{i=1}^n a_i^1 \cdot e_i\right) \wedge \cdots \wedge \left(\sum_{i=1}^n a_i^s \cdot e_i\right)$$

for some $(a_i^j) \in F$. We may distribute the wedge product over this summation by multilinearity—(3) above—and rearrange the terms appropriately by (1) above, so that we have a linear combination of vectors in the desired form.

To see that these vectors are linearly independent, we produce linear maps $B_{i_1,\ldots,i_k} : \bigwedge^s(V) \to F$ such that $B_{i_1,\ldots,i_s}(e_{i_1} \wedge \cdots \wedge e_{i_s}) = 1$ and for $\{i_1', \ldots, i_s'\} \neq \{i_1, \ldots, i_s\}$, $B_{i_1,\ldots,i_s}(e_{i_1'} \wedge \cdots \wedge e_{i_s'}) = 0$. This is sufficient because if $e_{i_1} \wedge \cdots \wedge e_{i_s}$ were to equal

$$\sum_{\{j_1,\ldots,j_s\} \neq \{i_1,\ldots i_s\}} a_{j_1,\ldots,j_s} \cdot e_{j_1} \wedge \cdots \wedge e_{j_s}$$

with some $a_{j_1,\ldots,j_s} \in F$ nonzero, we would have $B_{j_1,\ldots,j_s}(e_{i_1} \wedge \cdots \wedge e_{i_s}) = a_{j_1,\ldots,j_s}$, a contradiction.

Given an ordered n-tuple of basis elements $(e_{i_1}, \dots, e_{i_s}) \in V^n$, with no two indices equal, let $\sigma_{i_1,\dots,i_s}$ be the unique permutation that orders the indices of the basis elements. Consider the map $\Sigma_{i_1,\dots,i_s} : \underbrace{V \times \cdots \times V}_{s} \to F$, $1 \leq i_1 < \cdots < i_s \leq n$ defined by

$$\Sigma_{i_1,\dots,i_s}(e_{j_1}, \dots, e_{j_s}) = \begin{cases} 0, & \text{if } \{i_1, \dots, i_s\} \neq \{j_1, \dots, j_s\} \\ \operatorname{sign}(\sigma_{j_1,\dots,j_s}) & \text{if } \{i_1, \dots, i_s\} = \{j_1, \dots, j_s\} \end{cases}$$

and extended by imposing multilinearity. It is easy to check that this map is multilinear and alternating, so it must factor uniquely through $\wedge$; that the resulting map $\bigwedge^s(V) \to F$ is $B_{i_1,\dots,i_s}$ is also easy to check.

But the number of sets of strictly ordered indices $\{i_1, \dots, i_s\}$ is $\binom{n}{s}$, as claimed, which completes the proof. □

Corollary 4. *Let V be an n-dimensional vector space. Then* $\dim \bigwedge^n(V) = 1$.

Proof. By Proposition 2, we have $\dim \bigwedge^n(V) = \binom{n}{n} = 1$. □

We now extend the standard notion of the determinant. Given an endomorphism $T : V \to V$, we define $\bigwedge^s(T) : \bigwedge^s(V) \to \bigwedge^s(V)$ to be the map

$$v_1 \wedge \cdots \wedge v_s \mapsto T(v_1) \wedge \cdots \wedge T(v_s).$$

This map is linear as T is linear and as $\wedge$ is multilinear.

Definition 5 (Determinant)**.** The **determinant** $\det(T)$ of an endomorphism T of an n-dimensional vector space V is the map

$$\det(T) := \bigwedge^n(T).$$

In particular, note that $\bigwedge^n(T)$ is a map on a one-dimensional vector space (by Corollary 1), and is thus simply multiplication by a scalar. We claim that, having chosen a basis for $\bigwedge^n(V)$, this scalar is exactly the standard notion of the determinant; proving this is an exercise in algebra, which we recommend the reader pursue. Furthermore, this definition allows one to prove easily that the determinant of T is nonzero if and only if T is invertible; the proof follows below.

Proposition 6. *A linear map $T : V \to V$ is invertible if and only if* $\det(T) \neq 0$.

Proof. Note that $\bigwedge^n(\mathrm{id}_V) = \mathrm{id}_{\bigwedge^n(V)}$ and that, given two endomorphisms $T, S : V \to V$, $\bigwedge^n(T \circ S) = \bigwedge^n(T) \circ \bigwedge^n(S)$; that is, $\bigwedge$ respects identity and composition.[2]

To see necessity, note that we have

$$\mathrm{id}_{\bigwedge^n(V)} = \bigwedge^n(\mathrm{id}_V) = \bigwedge^n(T \circ T^{-1}) = \bigwedge^n(T) \circ \bigwedge^n(T^{-1}).$$

But then $\bigwedge^n(T)$ is non-zero, as it is invertible ($\bigwedge^n(T^{-1})$ is its inverse).

To see sufficiency, we show the contrapositive; that is, for non-invertible T, $\det(T) = 0$. Assume that $\dim T(V) < n$. Let $m = \dim T(V)$, and let $\{e_1, \dots, e_m\}$ be a basis for $T(V)$. But then given any $v_1 \wedge \cdots \wedge v_n \in \bigwedge^n(V)$, we have, distributing, that $\bigwedge^n(T)(v_1 \wedge \cdots \wedge v_n) = \sum a_{i_1,\dots,i_n} \cdot e_{i_1} \wedge \cdots \wedge e_{i_n}$; as $m < n$, we have by the pigeonhole principle that each term contains a repeated index. But then by (1) above, the determinant is zero as claimed. □

[2] In fact, $\bigwedge^n(-)$ is a functor.

2.2.2 Homotopies

The motivation here is to classify maps and domains by the existence of continuous transformations between them; we give some definitions that will be useful later. In particular, we wish to characterize the types of domains on which the the Poincaré Lemma will hold.

Definition 7 (Homotopy)**.** Two continuous maps $g_0, g_1 : U \to V$ with $U \subset \mathbb{R}^m, V \subset \mathbb{R}^n$ are said to be **homotopic** if there exists a continuous map $G : [0,1] \times U \to V$ such that for all $x \in U, g_0(x) = G(0,x)$ and $g_1(x) = G(1,x)$.

Intuitively, the notion behind this definition is that $G(t,-)$ interpolates continuously between g_0, g_1. We may use this idea to characterize certain types of domains, which may, speaking imprecisely, be continuously squished to a point.

Definition 8 (Contractibility)**.** We say a domain $U \in \mathbb{R}^m$ is **contractible** if, for some point $c \in U$, the constant map $x \mapsto c$ is homotopic to the identity on U.

Note that all convex and star-shaped domains are contractible.

2.2.3 The Change of Variables Formula

We now begin the calculus preliminaries to the Poincaré Lemma. A well-known theorem from single-variable calculus states that for integrable f defined and continuously differentiable g on $[a,b]$ with integrable derivative, and with f defined on $g([a,b])$, we have

$$\int_{g(a)}^{g(b)} f(x)\,dx = \int_a^b f \circ g(x) \cdot g'(x)\,dx. \tag{2.1}$$

This is proved easily using the chain rule and the fundamental theorem of calculus.

The theorem (2.1) may be generalized to multivariate functions as follows:

Proposition 9 (The Change of Variables Theorem)**.** *Consider open $U, V \subset \mathbb{R}^n$ with $g : U \to V$ an injective differentiable function with continuous partials and an invertible Jacobian for all $x \in U$. Then given continuous f with compact, connected support in $g(U)$, we have*

$$\int_{g(U)} f(x) = \int_U f \circ g(x) \cdot \left|\det(dg|_x)\right|$$

Proof. See [Sp, p. 66-72]. □

While this fact initially seems quite counterintuitive, careful thought gives good reason for the above formula. Consider a small rectangular prism in U with volume v; as long as it is non-degenerate, the vectors parallel to its sides form a basis for $\mathbb{R}^n$. For some x in this prism, we may approximate g at x as $g \approx T + dg|_x$, for some translation T. Then the action of g on this prism is (approximately) to transform the basis vectors parallel to its sides by $dg|_x$, inducing a new parallelepiped, which is non-degenerate if and only if $dg|_x$ is invertible. It is well-known, from computational geometry, that the volume of this new parallelepiped is $\left|\det dg|_x\right| \cdot v$. Considering the definition of the integral from Riemann sums, we have a geometrical motivation for this formula—the volume of each box in the summation is dilated by a factor of $\left|\det dg|_x\right|$.

2.3 Differential Forms

2.3.1 Motivation

The Change of Variables Theorem has an odd implication—that is, that integration is not coordinate-independent. In particular, diffeomorphic distortions of the coordinate system (that is, continuously

differentiable and invertible maps, whose inverse is also continuously differentiable) change the integrals of maps, even though no information is added or lost. This is undesirable because there is no obvious reason why any particular coordinate system is "better" than any other.

Much of mathematics seeks to escape from this type of arbitrary choice—an analogous motivation gives the dot product. The dot product gives a coordinate-free definition of length and angle; similarly, we would like to define a concept of the integral that is invariant under as many diffeomorphisms as possible.

Intuitively, the idea is to construct a class of objects that contain information about how they behave in any given coordinate system. In particular, we wish them to have some notion of "infinitesimal" area at any given point, which can be transformed by diffeomorphisms—we wish to formalize Leibniz's notion of infinitesimals. (The notation we will use will reflect this intention.) The goal is to have such objects encode the change of variables theorem as closely as possible.

It is interesting to note that, as a byproduct of this discussion, we will provide a formal, mathematical motivation for the div, grad, and curl operators, which are usually motivated only physically. We will also provide a generalization of these operators, and justify the physical intuition that they are connected to one another through more than just notation.

2.3.2 Definitions

Let U be a domain in $\mathbb{R}^n$.

Definition 10 (Differential Forms)**.** A **differential k-form** on U is a continuous, infinitely differentiable map $\omega : U \to \bigwedge^k(\mathbb{R}^{n*})$, where $\mathbb{R}^{n*}$ is the dual of $\mathbb{R}^n$ as a vector space. The set of all differential k-forms on U is denoted $\Omega^k(U)$. A k-form ω is also said to be of **degree** k, denoted $\deg \omega = k$.

In particular, we may let $x_1, \dots, x_n$ be a basis for $\mathbb{R}^n$, and let $dx_i \in \mathbb{R}^{n*}$ be the unique linear map $\mathbb{R}^n \to \mathbb{R}$ that satisfies

$$dx_i(x_j) = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j. \end{cases}$$

Then in this basis, we may write any differential k-form ω as

$$\omega(x) = \sum_{1 \leq i_1 < \dots < i_k \leq n} f_{i_1,\dots,i_k}(x) \cdot dx_{i_1} \wedge \cdots \wedge dx_{i_k},$$

for some $f_{i_1,\dots,i_k}$. Intuitively, we may consider $dx_{i_1} \wedge \cdots \wedge dx_{i_k}$ to be an oriented, k-dimensional volume element; this is the notation of "infinitesimal volume" we were looking for above. Note that this notion conforms geometrically with the properties of the exterior power: if one extends one dimension of a parallelepiped or formally sums two parallelepipeds, the volume changes linearly, and the orientation alters when one transposes two neighboring edges.

Differential forms admit a natural multiplication map $\wedge : \Omega^k(U) \times \Omega^l(U) \to \Omega^{k+l}(U)$, which is induced by the wedge product. However, this multiplication is *neither* antisymmetric nor symmetric; in particular, for $\omega \in \Omega^k(U), \alpha \in \Omega^l(U)$, we have

$$\omega \wedge \alpha = (-1)^{kl} \cdot \alpha \wedge \omega,$$

by reordering the dx_i.

Let V be a domain in $\mathbb{R}^k$; let g be a continuous, differentiable map $V \to U$.

Definition 11. The **pullback** of a k-form $\omega \in \Omega^k(U)$ through g, denoted $g^*(\omega) \in \Omega^k(V)$, is, with ω written as above, the map

$$g^*(\omega)(x) = \sum_{1 \leq i_1 < \dots < i_k \leq n} f_{i_1,\dots,i_k} \circ g(x) \cdot \bigwedge^k (dg|_x^*)(dx_{i_1} \wedge \cdots \wedge dx_{i_k})$$

where $dg|_x^*$ is the adjoint of the linear map $dg|_x$.

We claim that the pullback is the desired, (almost) coordinate-free transform we were looking for earlier. To see this, we must first define the integral of a differential form. In fact, for $U \subset \mathbb{R}^n$, we need only worry about n-forms, i.e. $\omega \in \Omega^n(U)$. Writing $\omega(x) = f(x) \cdot dx_1 \wedge \cdots \wedge dx_n$, we define

$$\int_U \omega := \int_U f(x)$$

where the integral on the right is the standard integral on real-valued functions. (Note the illustrative abuse of notation; if we write the term on the left out, we have

$$\int_U f \cdot dx_1 \wedge \cdots \wedge dx_n.$$

Omitting the wedges gives the standard notation from multivariate calculus.)

It is important to note that the sign of the integral is non-canonical; we have chosen an orientation of $\mathbb{R}^n$ by choosing an ordering of its basis vectors.

2.3.3 The Change of Variables Formula Revisited

We claim that the value of the integral is invariant under diffeomorphism, up to a sign. Let U, V be domains in $\mathbb{R}^n$, with $\omega \in \Omega^n(U)$. Consider a diffeomorphism $g : V \to U$. Then we claim

$$\int_U \omega = \pm \int_V g^*(\omega).$$

But writing out the term on the right according to our definitions, and writing $\omega(x)$ as $f(x) \cdot dx_1 \wedge \ldots \wedge dx_n$ gives us exactly

$$\begin{aligned}
\int_V g^*(\omega) &= \int_V f \circ g \cdot \bigwedge^k (dg|_x^*)(dx_{i_1} \wedge \cdots \wedge dx_{i_k}) \\
&= \int_V f \circ g \cdot \det(dg|_x^*) \cdot dx_{i_1} \wedge \cdots \wedge dx_{i_k} \\
&= \int_V f \circ g \cdot \det(dg|_x^*) \\
&= \int_V f \circ g \cdot \det(dg|_x),
\end{aligned}$$

where we use the fact that $\det(dg|_x) = \det(dg|_x^*)$. But this is exactly the change of variables formula without the absolute value sign. As g is a diffeomorphism, $dg|_x$ is always invertible and thus has nonzero derivative; continuity implies then that $\det(dg|_x)$ is everywhere-positive or everywhere-negative. So

$$\int_U \omega = \int_V f \circ g \cdot \left|\det(dg|_x)\right| = \pm \int_V f \circ g \cdot \det(dg|_x) = \pm \int_V g^*(\omega),$$

as claimed. We say that a diffeomorphism g is orientation-preserving if $\det(dg|_x)$ is everywhere-positive; in this case, we have strict equality above.

In fact, we may use this notion to redefine the notions of the line integral, the surface integral, and so on; in general, we may take the k-integral of a k-form ω on $U \subset \mathbb{R}^n$. Consider a domain $V \in \mathbb{R}^k$. Then the k-integral over a curve $g : V \to U$ is just the integral of $g^*(\omega)$ as defined above. Note that this is the integral of a k-form in $\mathbb{R}^k$, and is thus well-defined. This definition immediately gives invariance of the integral under appropriate re-parameterization, as before.

It is worth pausing here to examine what we have achieved. A careful reader might say that we have achieved *nothing*, at least insofar as pursuit of truth is concerned. We have just redefined

some terms: the integral, and coordinate transformations, to be precise. We have replaced them with ideas that conform to normative notions we have about how the objects in question should behave. In some sense, this evaluation would be true from a purely epistemological view, but it would miss the pedagogical point. By restricting our attention to objects which are coordinate-free, we can examine the coordinate-independent properties of the objects they correspond to with much greater ease. The value of this labor will become clear as we develop this machinery in the next few sections.

2.3.4 The Exterior Derivative

We wish to define an operator on differential forms that is similar to the derivative; in particular, it should satisfy some analogue of the product rule, and in some sense be invariant under coordinate transformations. Taking our cue from the antisymmetry of the wedge product, we want to find a collection of operators (d^k) that satisfies

- d^k is a linear map $\Omega^k(U) \to \Omega^{k+1}(U)$.
- Given two differential forms $\omega \in \Omega^k(U), \alpha \in \Omega^l(U)$, $d^{k+l}(\omega \wedge \alpha) = d^k(\omega) \wedge \alpha + (-1)^k \omega \wedge d^l(\alpha)$. This is analogous to the product rule. This condition and the prior condition make d a **derivation of degree 1**.
- For f a 0-form, i.e. a function $U \to \mathbb{R}, U \subset \mathbb{R}^n$, d coincides with the derivative in the following sense: $d^0(f) = \sum_{i=1}^n \frac{\partial f}{\partial x_i} \cdot dx_i$. That is, in matrix form, $d^0(f)(x)$ is exactly $df|_x$ (albeit in a different space, which, having chosen the bases $\{x_1, \ldots, x_n\}, \{dx_1, \ldots, dx_n\}$, is non-canonically isomorphic to the usual space).
- $d^{k+1} \circ d^k = 0$ for all k.

In general, we omit the superscript and the parentheses; i.e. $d^k(\omega)$ is written $d\omega$, and we write d^k as simply d for all k; the last condition above would then be written $d \circ d = 0$.

Proposition 12. *The map d is uniquely defined by the above conditions.*

Proof. Consider the function $\chi_i : \mathbb{R}^n \to \mathbb{R}$ given by $(x_1, x_2, \ldots, x_n) \mapsto x_i$; this function coincides with dx_i as defined above, but we use dx_i from here on to denote the constant differential form $x \mapsto 1 \cdot dx_i$, by (confusing) convention, just as we might use the constant c to denote the map $x \mapsto c$. Note that, by the third condition above, $d\chi_i = dx_i$. So by the fourth condition above, $d(dx_i) = 0$. We may now proceed to define d^k inductively, through the second condition above. In particular, it is clear that $\Omega^k(U)$ is spanned by the set of differential forms with a single term, e.g. $\omega = f(x) \cdot dx_{i_1} \wedge \ldots \wedge dx_{i_k}$. But then

$$\begin{aligned} d^k(f(x) \cdot dx_{i_1} \wedge \ldots \wedge dx_{i_{k-1}} \wedge dx_{i_k}) &= d^{k-1}(f(x) \cdot dx_{i_1} \wedge \ldots \wedge dx_{i_{k-1}}) \wedge dx_{i_k} \\ &\quad + (-1)^{k-1} f(x) \cdot dx_{i_1} \wedge \ldots \wedge dx_{i_{k-1}} \wedge d(dx_{i_k}) \\ &= d^{k-1}(f(x) \cdot dx_{i_1} \wedge \ldots \wedge dx_{i_{k-1}}) \wedge dx_{i_k} \end{aligned}$$

and we may extend d linearly to linear combinations of single-term forms. While this proves uniqueness, it is not immediately clear that d is well-defined, as we must check that property (2) holds for all k, l, rather than just for $l = 1$.

To show that d is well-defined, we give an explicit construction that satisfies the inductive construction given above. In particular, for single-term forms $\omega(x) = f(x) \cdot dx_{i_1} \wedge \ldots \wedge dx_{i_k}$, we let

$$d\omega = \sum_{j=1}^n \frac{\partial f}{\partial x_j} \cdot dx_j \wedge dx_{i_1} \wedge \ldots \wedge dx_{i_k}.$$

The interested reader can check that this construction satisfies the first three conditions above; we check the fourth. The proof is inductive. We have that for 0-forms, e.g. $f(x)$, that

$$\begin{aligned} d \circ d(f) &= d\left(\sum_{i=1}^{n} \frac{\partial f}{\partial x_i} \cdot dx_i\right) \\ &= \sum_{j=1}^{n}\sum_{i=1}^{n} \frac{\partial^2 f}{\partial x_j \partial x_i} \cdot dx_j \wedge dx_i \\ &= \sum_{i<j} \left(\frac{\partial^2 f}{\partial x_i \partial x_j} dx_i \wedge dx_j + \frac{\partial^2 f}{\partial x_j \partial x_i} dx_j \wedge dx_i\right) \\ &= \sum_{i<j} \left(\frac{\partial^2 f}{\partial x_i \partial x_j} dx_i \wedge dx_j - \frac{\partial^2 f}{\partial x_j \partial x_i} dx_i \wedge dx_j\right) \\ &= 0, \end{aligned}$$

where the last step uses the equality of mixed partials. We have already noted that $d(dx_i) = 0$ above. Assume that for $i < p$, we have for $\omega \in \Omega^i(U), d \circ d(\omega) = 0$. Then for $\alpha \in \Omega^p(U)$ with one term, we may write $\alpha = \beta \wedge \gamma$ for $\beta \in \Omega^1(U), \gamma \in \Omega^{p-1}(U)$. We have

$$\begin{aligned} d \circ d(\alpha) &= d \circ d(\beta \wedge \gamma) \\ &= d(d\beta \wedge \gamma - \beta \wedge d\gamma) \\ &= d(d\beta \wedge \gamma) - d(\beta \wedge d\gamma) \\ &= d(d\beta) \wedge \gamma + d\beta \wedge d\gamma - d\beta \wedge d\gamma + \beta \wedge d(d\gamma) \\ &= 0 \end{aligned}$$

by the induction hypothesis; differential forms with more than one term satisfy the same claim by linearity. This completes the proof. □

Calculation gives that the exterior derivative commutes with the pullback, e.g. $g^*(d\omega) = d(g^*(\omega))$. That is, in some sense, the exterior derivative "flows" with changes of coordinates; for 0-forms, this is just the chain rule.

For the rest of this subsection, we will restrict our attention to $\mathbb{R}^3$. Note that, from Proposition 2, we have that the dimensions of $\bigwedge^0(\mathbb{R}^3), \bigwedge^1(\mathbb{R}^3), \bigwedge^2(\mathbb{R}^3)$, and $\bigwedge^3(\mathbb{R}^3)$ are $1, 3, 3$, and 1, respectively. In particular, we can identify $\bigwedge^0(\mathbb{R}^3)$ and $\bigwedge^3(\mathbb{R}^3)$ with $\mathbb{R}$ (actually the former is identified as such canonically), and $\bigwedge^1(\mathbb{R}^3), \bigwedge^2(\mathbb{R}^3)$ with $\mathbb{R}^3$. Then d gives maps $\mathbb{R} \to \mathbb{R}^3$, etc., and, by precomposition, maps $\nabla : C^\infty(\mathbb{R}^3) \to (\mathbb{R}^3 \to \mathbb{R}^3), \nabla\times : (\mathbb{R}^3 \to \mathbb{R}^3) \to (\mathbb{R}^3 \to \mathbb{R}^3), \nabla\cdot : (\mathbb{R}^3 \to \mathbb{R}^3) \to C^\infty(\mathbb{R}^3)$. Easy computation gives that these maps correspond, respectively, to the gradient, curl, and divergence. In fact, the first of these three computations follows immediately from the third bullet in the definition of the exterior derivative.

It is important to note that the identifications above are non-canonical; in the most general case, we define a canonical isomorphism called the **Hodge dual**, denoted

$$* : \bigwedge^k(V) \xrightarrow{\sim} \bigwedge^{n-k}(V),$$

where n is the dimension of V. Using this isomorphism, we may extend the ideas of gradient, curl, and divergence given above to vector spaces with arbitrary finite dimension.

2.3.5 The Interior Product and the Lie Derivative

Consider an element $w \in \bigwedge^k(V)$, where V is a finite-dimensional vector space over some field F. Then w can be viewed as an alternating, multilinear mapping $w : V^{*k} \to F$, where V^* is the

vector space dual to V (here we take advantage of the canonical isomorphism $V \simeq V^{**}$). Given $a_1, \dots, a_n \in V^*$, we may define, for each $\alpha \in V^*$, a function ι_α such that $\iota_\alpha(w)(a_1, \dots, a_n) = w(\alpha, a_1, \dots, a_n)$.

In particular, we may uniquely define ι_α as follows [Wa, p. 61]:

- $\iota_\alpha : \bigwedge^k(V) \to \bigwedge^{k-1}(V)$,
- For $v \in \bigwedge^1(V), \iota_\alpha(v) = \alpha(v)$,
- ι_α is a derivation of degree -1, i.e. $\iota_\alpha(a \wedge b) = \iota_\alpha(a) \wedge b + (-1)^{\deg a} a \wedge \iota_\alpha(b)$.

The proof that these properties uniquely define ι_α is analogous to the proof for d above and is left to the reader.

We now restrict our attention to $V = (\mathbb{R}^n)^*$. Consider a vector field $\xi : U \to \mathbb{R}^n$, where U is a domain in $\mathbb{R}^n$. Then for a differential form ω on U, we may let ι_ξ act on U point-wise, e.g. $\iota_\xi(\omega)(x) = \iota_{\xi(x)}(\omega(x))$. We define the **Lie Derivative** of a form ω with respect to a vector field ξ by

$$\text{Lie}_\xi(\omega) := d \circ \iota_\xi(\omega) + \iota_\xi \circ d(\omega).$$

In some sense, this operator takes the derivative of a form with respect to a (possibly time-dependent) vector field. This intuition is clear for constant vector field; computation, which we omit, gives that for a constant vector field $\vec{x}$,

$$\text{Lie}_{\vec{x}}(f \cdot dx_1 \wedge \cdots \wedge dx_n) = \frac{\partial f}{\partial x} \cdot dx_1 \wedge \cdots \wedge dx_n.$$

This fact will be useful later, and to remind ourselves of it, we will denote a constant vector field with respect to some coordinate x_i as $\frac{\partial}{\partial x_i}$.

2.4 Chain Complexes

Above, we noted that the exterior derivative satisfies $d \circ d = 0$. This fact suggests a more general structure, which we abstract as follows:

Definition 13 (Chain Complex)**.** A **chain complex** is a sequence of Abelian groups (or algebraic objects with Abelian structure, e.g. modules, vector spaces) $A_{-1}, A_0, A_1, A_2, \dots$ with connecting homomorphisms $d^k : A_k \to A_{k-1}$ such that for all $k, d^k \circ d^{k+1} = 0$. We denote all of this data as $(A_\bullet, d_\bullet)$.

In a **cochain complex**, the connecting homomorphisms proceed in the opposite direction; i.e. $d^k : A_k \to A_{k+1}$ and $d^k \circ d^{k-1} = 0$. In this case, we denote the entire collection of Abelian groups and connecting homomorphisms as $(A^\bullet, d^\bullet)$.

Note that the chains and cochains are identical, but with different indexing; the terminology stems from convention. The notion of the pullback suggests the following:

Definition 14 (Map of Complexes)**.** A **map of complexes** $\psi^\bullet : (A^\bullet, d^\bullet) \to (B^\bullet, e^\bullet)$, in the case of a cochain, is a collection of maps $\psi^k : A_k \to B_k$ such that $e^k \circ \psi^k = \psi^{k+1} \circ d^k$, i.e. the diagram in Figure 2.1 commutes. The case of chains is analogous.

It should be clear by now that differential forms on some domain $U \subset \mathbb{R}^n$ form a complex

$$\cdots \longrightarrow 0 \longrightarrow \Omega^0(U) \xrightarrow{d} \Omega^1(U) \xrightarrow{d} \cdots \xrightarrow{d} \Omega^n(U) \longrightarrow 0 \longrightarrow \cdots$$

and, from the fact that pullbacks commute with the exterior derivative, that pullbacks are maps of complexes. We call this complex the **de Rham complex** and denote the de Rham complex on U

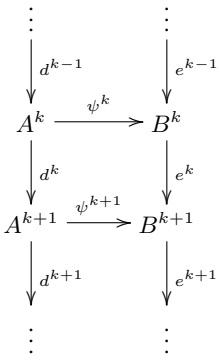

Figure 2.1: A Map of Complexes.

as $(\Omega^\bullet(U), d^\bullet)$. As in the de Rham derivative, we often ignore the superscripts on the connecting homomorphisms, e.g. $d \circ d = 0$, in arbitrary chains or cochains. Also, when we are discussing more than one complex, it is common to use the same symbol for their respective connecting homomorphisms, e.g. $d \circ \psi^k = \psi^{k+1} \circ d$.

Definition 15 (Closed, Exact)**.** Given an element $\omega \in A_k$, we say that ω is **closed** if $d\omega = 0$. We say that ω is **exact** if there exists α such that $d\alpha = \omega$.

Note that, as $d \circ d = 0$, we have that im $d^k \subset \ker d^{k+1}$; that is, all exact elements of a chain or cochain are closed. It is natural to ask when closed elements are exact—in the de Rham complex, the Poincaré Lemma addresses this question to a large extent. Pursuing it in general, we define:

Definition 16 (Homology, Cohomology)**.** The k**-th homology group** of a chain $(A_\bullet, d_\bullet)$ is

$$H_k(A_\bullet) := \frac{\ker d^k}{\text{im } d^{k+1}}.$$

Analogously, in a co-chain $(B^\bullet, d^\bullet)$, the k**-th cohomology group of** $B^\bullet$ is

$$H^k(B^\bullet) := \frac{\ker d^k}{\text{im } d^{k-1}}.$$

Intuitively, this group characterizes those closed forms that are not exact; i.e. elements in any given coset are identical up to an exact form.

Consider two co-chains $A^\bullet, B^\bullet$ and a map of complexes $\phi^\bullet : A^\bullet \to B^\bullet$. We claim that $\phi^\bullet$ induces well-defined maps $H^k(\phi^\bullet) : H^k(A^\bullet) \to H^k(B^\bullet)$. (An analogous claim holds for chains.)

Proof. Consider an element $[a] \in H^k(A^\bullet)$. We claim that the map $H^k(\phi^\bullet) : [a] \to [\phi(a)]$ is well-defined. By definition, any element $a' \in [a]$ differs from a by an exact element ω; as it is exact, we may write $\omega = d\alpha$ for some α. Then $[\phi(a')] = [\phi(a + d\alpha)] = [\phi(a) + \phi(d\alpha)] = [\phi(a) + d(\phi(\alpha))] = [\phi(a)]$, where we use the fact that maps of complexes commute with the complexes' connecting maps. So the map of cohomologies (resp. homologies) is well-defined. □

It is natural to ask when two maps of complexes induce the same map between cohomologies. To this end, we consider the following definition:

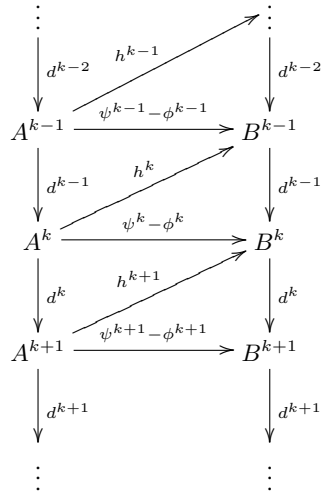

Figure 2.2: A Homotopy of Cochain Complexes.

Definition 17 (Homotopy)**.** In a (justifiable, as we shall see) homonym, we say that two maps $\psi^\bullet, \phi^\bullet : A^\bullet \to B^\bullet$ of complexes are **homotopic** through the **homotopy** (h^k) if there exists a sequence of maps $h^k : A^k \to B^{k-1}$ (in a co-chain, with analogous indexing for chains) such that

$$\psi^k - \phi^k = d \circ h^k + h^{k+1} \circ d.$$

That is, in the diagram in Figure 2.2, the parallelograms commute with the horizontal arrows.

Proposition 18. *If two maps of complexes $\psi^\bullet, \phi^\bullet$ are homotopic through some homotopy (h^k), then $H^k(\psi^\bullet) = H^k(\phi^\bullet)$.*

Proof. Note that, by linearity $H^k(\psi^\bullet) - H^k(\phi^\bullet) = H^k(\psi^\bullet - \phi^\bullet) = H^k(d \circ h^k + h^{k+1} \circ d) = H^k(d \circ h^k) + H^k(h^{k+1} \circ d)$. We claim that $H^k(d \circ h^k) = H^k(h^{k+1} \circ d) = 0$. To see that $H^k(h^{k+1} \circ d) = 0$, note that for $[a] \in H^k(A^\bullet)$, we have that $a \in \ker(d)$, so $H^k(h^{k+1} \circ d)([a]) = [h^{k+1} \circ d(a)] = 0$. Furthermore, $H^k(d \circ h^k)([a]) = [d(h^k(a))]$. But $d(h^k(a)) \in \text{im}(d)$, so $[d(h^k(a))] = [0]$. But then $H^k(\psi^\bullet) - H^k(\phi^\bullet) = 0$, so $H^k(\psi^\bullet) = H^k(\phi^\bullet)$ as claimed. □

2.5 The Poincaré Lemma

We finally are able to state and prove the Poincaré Lemma. We wish to characterize situations in which closed forms are also exact.

Theorem 19 (The Poincaré Lemma)**.** *Let U be a contractible domain in $\mathbb{R}^n$, and let k be a positive integer. Then for $\omega \in \Omega^k(U)$ such that $d\omega = 0$, there exists $\alpha \in \Omega^{k-1}(U)$ such that $\omega = d\alpha$. In other words all closed differential k-forms on contractible domains are exact.*

Proof. We first prove a general lemma—that is, that the pullbacks through homotopic maps are homotopic as maps of complexes, as is suggested by the terminology.

Lemma 20. *Let V and W be domains, $V \subset \mathbb{R}^n, W \subset \mathbb{R}^m$. Consider maps $g_0, g_1 : V \to W$ that are homotopic, i.e. there is a map $G : I \times V \to W$, where $I = [0,1]$ such that $G(0,x) = g_0(x), G(1,x) = g_1(x)$. Then the maps of complexes $g_0^*, g_1^* : \Omega^k(W) \to \Omega^k(V)$ are homotopic.*

Proof of Lemma 20. Let $G_t : V \to W$ be the map $x \mapsto G(t,x)$. For $\omega \in \Omega^k(W)$, define

$$h^k(\omega)(x) = \int_{t=0}^{t=1} \iota_{\frac{\partial}{\partial t}}(G_t^*(\omega))(x).$$

We claim that this is the desired homotopy of complexes. In particular, we have that

$$\begin{aligned}
(d^{k-1} \circ h^k + h^{k+1} \circ d^k)(\omega) &= d\left(\int_{t=0}^{t=1} \iota_{\frac{\partial}{\partial t}}(G_t^*(\omega))\right) + \int_{t=0}^{t=1} \iota_{\frac{\partial}{\partial t}}(G_t^*(d\omega)) \\
&= \int_{t=0}^{t=1} (d \circ \iota_{\frac{\partial}{\partial t}} + \iota_{\frac{\partial}{\partial t}} \circ d)(G_t^*(\omega)) \\
&= \int_{t=0}^{t=1} \mathrm{Lie}_{\frac{\partial}{\partial t}}(G_t^*(\omega)) \\
&= \int_{t=0}^{t=1} \frac{\partial}{\partial t} G_t^*(\omega) \\
&= G_1^*(\omega) - G_0^*(\omega) \\
&= g_1^*(\omega) - g_0^*(\omega),
\end{aligned}$$

as desired. In the above manipulations, we use the commutation of the differential with the integral and the pullback as well as the fundamental theorem of calculus. □

Corollary 21. *The pullbacks through homotopic maps act identically on the cohomologies, that is, $H^k(g_0^*) = H^k(g_1^*)$. In particular, on a contractible domain, $H^k(id^*) = H^k(c^*)$, where c is the constant map.*

But $H^k(c^*) = 0$. So we have from the corollary that, on contractible domains, $H^k(\mathrm{id}^*) = 0$. But then $H^k(\Omega^k(U)) = 0$, i.e. $\mathrm{im}\, d = \ker d$. And this is precisely what we wanted to prove. □

2.6 Conclusion

It is valuable to consider what, if anything, we have accomplished beyond the Lemma itself. In particular, the ideas here seem somewhat far-removed from those where we started—in the realm of coordinate-invariant objects. What does the Poincaré Lemma tell us? What have we gained by introducing such strange, if elegant, mathematical tools?

To begin with, many more standard proofs of the Lemma are heavily calculational [Sp]; the referenced method proves the theorem only on star-shaped domains, and at the cost of massive amounts of counterintuitive index-juggling. The methods here slightly weaken the hypothesis on the domain and achieve a much cleaner solution.

But more importantly, the tools we have developed have varied applications. One of the better-known such applications occurs in electrodynamics, where Maxwell's equations tell us that, under magneto-static conditions,

$$\nabla \times \vec{\mathbf{E}} = 0,$$

where $\vec{\mathbf{E}}$ denotes the electric field. The Poincaré Lemma implies immediately that there exists a scalar function V such that

$$\vec{\mathbf{E}} = -\nabla V,$$

that is, the electric potential.

2.7 Acknowledgment

The vast majority of the information presented here was demonstrated to the author in some form by Professor Dennis Gaitsgory of Harvard University, and the methods herein owe much to his presentation.

References

[DF] David S. Dummit and Richard M. Foote: *Abstract Algebra*, 3rd Edition. New York: John Wiley and Sons, 2004.

[Sp] Michael Spivak: *Calculus on Manifolds*. New York: Perseus Books Publishing, 1965.

[Wa] Frank W. Warner: *Foundations of Differentiable Manifolds and Lie Groups.* New York: Springer Science+Business Media, 1983.

STUDENT ARTICLE

3 An Introduction to Combinatorial Game Theory

Paul Kominers†
Walt Whitman High School '08
Bethesda, MD 20817
paul.kominers@gmail.com

Abstract

We survey the field of combinatorial game theory. We discuss Zermelo's Theorem, a foundational result on which the theory of combinatorial game strategy is based. We then introduce the simple game of Nim and explain how it, through the theory of Nimbers, is critical to and underlies all of impartial combinatorial game theory.

3.1 Combinatorial Game Theory

3.1.1 History

Combinatorial game theory, founded in the early 20th century, deals with recursive analysis of **combinatorial games**, two-player games having neither chance elements nor concealed information.[1] Combinatorial game theory allows mathematical analysis of games as seemingly simple as Nim and, potentially, those as complex as chess. Additionally, combinatorial games are finite, and the players alternate moves in well-defined plays [Br, De, Fe].

Two foundational works in the field of partizan games[2] are Conway's *On Numbers and Games* [Co] and Berkelamp, Conway, and Guy's *Winning Ways for Your Mathematical Plays* [BCG]; the study of impartial games began with Zermelo's Theorem.

3.1.2 Zermelo's Theorem

Since combinatorial games are finite, they must end either with a win for one player or a draw for both [Br, Fe]. Throughout this paper, the first player to move will be referred to as Player 1 and the second player to move will be referred to as Player 2.

In any combinatorial game, either Player 1 has a winning strategy, Player 2 has a winning strategy, or both players have a strategy that guarantees a draw [Br, Fe]. Because the games have perfect information, if Player 1's opening move is the best possible opening move and Player 2's response is the best possible response, there is no reason for either player to change his strategy in successive games. However, since there is no chance element, such play will always lead to either a draw or a win for one of the players. This notion is formalized in Zermelo's Theorem:

Theorem 1 (Zermelo's Theorem, see [Br])**.** *In a combinatorial game, either one of the players has a formal strategy that guarantees a win, or both players have formal strategies that guarantee at least a draw.*

†Paul Kominers is a high school senior at Walt Whitman High School in Bethesda, Maryland. His mathematical interests lie chiefly in subfields of combinatorics, including extremal combinatorics and combinatorial game theory. He plans to study computer science in college, although a mathematics major is increasingly tempting.

[1]Games with neither chance elements nor concealed information are said to have **perfect information**.

[2]A game which is not impartial is called **partizan**.

3.1.3 Impartiality

A game in which both players have identical sets of available moves at any point in the game is called **impartial**. In other words, if Player 1 and Player 2 are playing an impartial game and Player 2 opts (and is allowed) to skip his move, then Player 1's ideal move is the move which would have been Player 2's ideal move.

3.2 Nim

3.2.1 Definition

The rules of the game of **Nim** are given below.

Game (Nim)**.** *The players are given n piles of matches, where the kth pile has m_k matches. The players take turns choosing and removing any number of matches from a single pile. The player to take the last match wins.*

Nim is perhaps the most important impartial combinatorial game. It is self-evident that it is combinatorial: there is perfect information and the game must end when one player cannot remove any more matchsticks. It is impartial because both players have the same choice of matchsticks to remove; given any Nim game, the set of Player 1's available moves is identical to the set of Player 2's given moves.

In this article, a Nim game of n piles is written as $\{m_1, m_2, \ldots, m_k, \ldots, m_n\}$. For example, the Nim game with $n = 3$ piles of matchsticks such that $m_1 = m_2 = 1$ and $m_3 = 2$ will be written as $\{1, 1, 2\}$. Player 1's moves will be denoted by lines above the target pile; Player 2's moves will be denoted by lines below. As an example, if Player 1 decided to remove one matchstick from pile three, it would be denoted as:

$$\left\{1, 1, \overline{2}\right\} \to \{1, 1, 1\}\,.$$

A complete game might proceed as follows:

$$\left\{1, 1, \overline{2}\right\} \to \{1, 1, \underline{1}\} \to \left\{1, \overline{1}, 0\right\} \to \{\underline{1}, 0, 0\} \to \{0, 0, 0\}\,. \tag{3.1}$$

In sample game (3.1), Player 2 takes the last matchstick and wins.

3.2.2 Symmetric Strategies

Player 2's win in sample game (3.1) resulted from a foolish move by Player 1. A different opening move would have given Player 1 an easy win:

$$\left\{1, 1, \overline{2}\right\} \to \{1, \underline{1}, 0\} \to \left\{\overline{1}, 0, 0\right\} \to \{0, 0, 0\}\,. \tag{3.2}$$

The sample game (3.2) is an example, albeit a simple one, of a game won by a **symmetric strategy**. A symmetric strategy is a strategy in which one player creates a situation in which he can always copy his opponent's move. In so doing, the player is guaranteed a response to each of his opponent's moves. Most importantly, he can make the last move and win.

Imagine that a Nim game has come down to the following two piles on Player 1's move:

$$\{3, 3\}\,.$$

If Player 2 plays competently, Player 1 has lost. No matter how many matches Player 1 removes from one pile, Player 2 can remove the same number from the other. Eventually, Player 1 will remove the last matchstick from one pile. Player 2 will do the same to the other, thereby winning the game. For example, the game could proceed as follows:

$$\left\{3, \overline{3}\right\} \to \{\underline{3}, 2\} \to \left\{2, \overline{2}\right\} \to \{\underline{2}, 0\} \to \{0, 0\}\,.$$

3.2.3 Winning and Losing Positions

Of course, the winning move is not this easy to spot in all Nim games; Nim can be markedly more complex. For example, consider the game

$$\{2, 2, 3, 5, 7, 8, 9\}. \tag{3.3}$$

In this game, Player 1 has a **losing position**, or **safe** position[3]; that is, a game state such that, assuming ideal play from one's opponent, one will always lose. Analogously, a **winning position**, or **unsafe** position, is a position from which, if one plays ideally, one will always win. To properly apply this notion, instead of thinking of a combinatorial game as a series of moves, we must treat it as a series of **inherited positions**. Player 1's making a move must be thought of as Player 1's causing Player 2 to inherit Player 1's position, slightly modified. Player 2, in response, modifies the position slightly and then causes Player 1 to inherit that position. The modified position P' inherited after the original position P is called the **successor** of P. We will show shortly that a player inheriting a losing position can only ever cause his opponent to inherit a winning position and that every winning position has at least one move such that the player's opponent will inherit a losing position.

3.2.4 Nimbers

The notion of inheriting positions is foundational to the theory of **Nimbers**, alternately called **Nim-sums** or **Sprague-Grundy Numbers**, invariants which both give limited information about the game and allow us to find isomorphisms between games. Nimbers are found by the following recursion (see [Br, Fe]):

> The Nimber of any losing position is 0. The Nimber of any current position is the smallest non-negative integer not in the set of Nimbers of positions which can result from the current position. So, for example, if a given position can only go to 0-positions, its Nimber is 1. If it can go to 0- or 1-positions, its Nimber is 2. If it can go to 0- or 2-positions, its Nimber is 1, and if it can go to 1- or 2-positions, its Nimber is 0. Any Nimber greater than 0 indicates a winning position.

There is a unique Nimber for each position in an impartial combinatorial game (see [Br] for a proof). This definition of Nimbers is well-defined; in particular, all losing positions have Nimber 0. We can now show that safe positions can have only unsafe successors and that unsafe positions can always have at least one safe successor.

Suppose that some losing position A exists such that A has a successor A' which is a losing position. As losing positions, both A and A' have Nimber 0. However, if A' has Nimber 0, the smallest non-negative integer in the set of Nimbers of successors of A cannot be 0. Therefore, A cannot be a losing position, contradicting our initial assumption.

Now suppose that some winning position B, with Nimber greater than 0, exists such that B has no losing successor. If B has no losing successor, then it has no successor with Nimber 0. This means that 0 cannot be included in the set of Nimbers of successors of B. Since 0 is the smallest non-negative integer, the smallest non-negative integer not present in a set not containing 0 must be 0. If B's Nimber is 0, then it is a losing rather than a winning position, so a contradiction occurs.

3.2.5 The Solution to Nim

We can now use the theory of Nimbers to show how Player 1 occupies a losing position in sample game (3.3).

In Nim, as an alternative to the described recursion, the Nimber of any position can be computed by the following algorithm (see [Br]):

[3]The choice of the word "safe" may seem incongruous until one considers that, for Player 2, the position that Player 1 is in is perfectly safe.

NIM-SUM ALGORITHM:

1. Convert the number of matches in each pile to binary.

2. Add the binary digits modulo two.

For an example, we will return to our original sample game, (3.1). By our algorithm, we compute the Nim-sum to be

$$\begin{array}{ccc} 1 & \longrightarrow & 01 \\ 1 & \longrightarrow & 01 \\ 2 & \longrightarrow & 10 \\ \hline & & 10 \end{array}$$

Since this sum gives the Nimber of a position (hence the names "Nimber" and "Nim-sums"), it follows that the winning move is the move that reduces the sum to zero (see [Co]). Therefore, we see again that the winning move is to remove the entire pile of two.

This particular game also demonstrates an important property of Nim and of combinatorial games in general: two identical games cancel each other out. Suppose we consider a Nim game of n piles to be the sum of n one-pile Nim games. Since the sum of any two piles of the same size is zero, they do not affect the Nimber of any position.

Therefore, we may once again examine the game $\{2, 2, 3, 5, 7, 8, 9\}$ (or, since the two piles of size 2 cancel each other, the equivalent game $\{3, 5, 7, 8, 9\}$) and see that the Nim-sum is zero. If the game were to be played out, it might proceed something like this (with Nim-sums of positions given under the positions themselves):

$$\begin{array}{ccccccccccc} \{\overline{2},2,3,5,7,8,9\} & & \{0,\underline{2},3,5,7,8,9\} & & \{0,0,\overline{3},5,7,8,9\} & & \{0,0,0,\underline{5},\underline{7},8,9\} & & \{0,0,0,5,4,8,\overline{9}\} & & \{0,0,0,5,4,\underline{8},0\} \\ 0 & \rightarrow & 2 & \rightarrow & 0 & \rightarrow & 3 & \rightarrow & 0 & \rightarrow & 9 \end{array}$$

$$\begin{array}{ccccccccccc} & & \{0,0,0,\overline{5},4,1,0\} & & \{0,0,0,0,\underline{4},1,0\} & & \{0,0,0,0,1,\overline{1},0\} & & \{0,0,0,0,\underline{1},0,0\} & & \{0,0,0,0,0,0,0\} \\ \rightarrow & & 0 & \rightarrow & 5 & \rightarrow & 0 & \rightarrow & 1 & \rightarrow & 0 \end{array}.$$

Note that Player 1 moves arbitrarily, while Player 2 always moves to return the game's Nim-sum to zero.

3.2.6 Nimbers in Isomorphisms

Many combinatorial games are **isomorphic** to each other; in other words, despite different appearances, the games can be shown to be mathematically equivalent.[4] More formally, two isomorphic combinatorial games have identical trees of Nimbers.

In some instances, the isomorphisms are fairly obvious, as in the game of **Rook on a 3-D Board**, from [Br]:

Game (Rook on a 3-D Board)**.** *A rook is placed in the north-east back corner of a three-dimensional $i_1 \times i_2 \times i_3$ gameboard. In turn, players move the rook either south, west, or forward any number of spaces. The player who moves the rook into the south-west corner wins.*

It is clear that this game is equivalent to the Nim game $\{m_1 = i_1, m_2 = i_2, m_3 = i_3\}$. Each dimension takes the place of one pile of matchsticks.

In general, we have that:

Theorem 2. *Any finite impartial game G played such that one move strictly changes one of a finite set of numbers $\{i_1, i_2, ..., i_n\}$ that ends when all elements of the set have reached 0 will have the same outcome as the Nim game $\{m_1 = i_1, m_2 = i_2, ..., m_n = i_n\}$.*

[4] Isomorphism comes from the Greek *iso* meaning "same" and *morph* meaning "form."

Since there is no limit to the amount by which any element of the set is changed, whatever either player increases an element by, the other player can decrease that element by more. The conditions by which G ends are identical to those of Nim. Therefore, each element in the set corresponds to a pile in the isomorphic Nim game.

It is possible to construct some Nim game such that the possible moves produce any sequence of Nimbers. Since two games having identical Nimber trees are isomorphic, there must exist some Nim game that is isomorphic to G. This is formalized in the well-known Sprague-Grundy Theorem:

Theorem 3 (Sprague-Grundy Theorem)**.** *Any impartial game G is isomorphic to some Nim game.*

Formal proofs of both Theorems 2 and 3 can be found in Conway's *On Numbers and Games* [Co].

3.3 Conclusion

We defined combinatorial games and discussed several ideas foundational to their study. We also introduced and illustrated the solution to the game of Nim, which is critical in solving many combinatorial games. Finally, we noted that all impartial combinatorial games are isomorphic to some game of Nim.

There are still many unsolved problems in combinatorial game theory. Games such as Chess and Go are so complex enough that they deny easy analysis (see [De]). Meanwhile, the ease with which new combinatorial games can be created results in an infinite set of new games to work with. These new games often have surprising depth or hidden isomorphisms to better-known games.

References

[BCG] Elwyn Berlekamp, John H. Conway, and Richard Guy: *Winning Ways for Your Mathematical Plays*, 2nd ed., Vols. 1–4. Massachusetts: AK Peters, 2001.

[Br] Mortimer Brown: *Mathematical Games*. Unpublished, 2006.

[Co] John Conway: *On Numbers And Games* 2nd ed. Massachusetts: AK Peters, 2000.

[De] Erik D. Demaine: Playing games with algorithms: algorithmic combinatorial game theory, *Lecture Notes in Computer Science* **2136** (2001) [=Proceedings of the 26th Symposium on Mathematical Foundations in Computer Science], 18–32.

[Fe] Thomas Ferguson: *Impartial Combinatorial Games Notes*. Preprint, 2005 (online at `http://www.math.ucla.edu/tom/Game_Theory/comb.pdf`).

STUDENT ARTICLE

4 The Knot Quandle

Eleanor Birrell[†]
Harvard University '09
Cambridge, MA 02138
ebirrell@fas.harvard.edu

Abstract

Mathematicians have been interested in knot theory, or the study of knots, since the early nineteenth century. However, despite this interest, some basic questions remain unanswered; for example there is no effective way to definitively determine whether or not two knots are the same. In this paper, we will look at a powerful but frequently overlooked knot invariant: the knot quandle. We will show that the knot quandle is a generalization of several more familiar invariants and that it is a complete invariant up to orientations. However, as we will see, determining whether two quandles are isomorphic is computationally intractable, which limits the utility of this otherwise powerful invariant.

4.1 Introduction

Knot theory is a subfield of mathematics that can be described simply as the study of knots, or embeddings of S^1 into $\mathbb{R}^3$. Since its beginning in the nineteenth century, knot theory has appealed to mathematicians for a variety of reasons. It contains many interesting theoretical questions related to both algebra and topology and also has applications to biology, physics, and cryptography.

Despite the range of questions and applications that arise in knot theory, some relatively simple questions remain unanswered. For example, how does one tell whether or not two knots are "the same"? We consider two knots to be the same if there exists an **ambient isotopy** between them; i.e. there is a homotopy of self-diffeomorphisms from $\mathbb{R}^3 \to \mathbb{R}^3$ that transforms one knot into the other. This means that two knots are the same if there exists a continuous deformation of space that takes the first knot to the second (this definition is consistent with the intuitive idea of equivalent knots). Kurt Reidemeister showed that two knots are connected by an ambient isotopy if and only if they are connected by a finite chain of Reidemeister moves [Cr], which are defined on knot diagrams, a special type of projection of the knots into $\mathbb{R}^2$ (see Figure 4.2).

Despite the lack of a complete answer, some progress has been made towards determining when knots are equivalent. One important although frequently overlooked step was the development of the knot quandle. In 1982, David Joyce [Jo] and Sergei Matveev [Mat] independently introduced a **knot invariant** (an object that is invariant under different representations of equivalent knots) that would aid in attempts to determine knot isomorphisms. They called this invariant the **knot quandle** and the **distributive groupoid**, respectively. Although this invariant is not perfect (it fails to distinguish between the right and left handed trefoils, for example), it proves to be complete up to orientation. Moreover, this invariant serves as a generalization of both the knot group and colorability, suggesting that the knot quandle is, perhaps, the most complete and the most fundamental invariant known today. Unfortunately, the difficulty inherent in proving that two quandles are isomorphic has severely limited the utility of this otherwise powerful invariant.

[†]Eleanor Birrell, Harvard '09, is a mathematics concentrator living in Pforzheimer House. She is originally from Los Altos, California where she graduated from Los Altos High School. Her academic interests include algebraic topology and computational complexity.

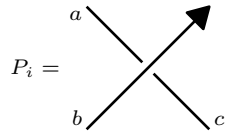

Figure 4.1: A Crossing in a Knot Diagram.

4.2 The Knot Quandle

The knot quandle is defined by associating a quandle structure with a knot K.

Definition 1. A **quandle** is a set S equipped with a binary operation $\triangleright : S \times S \to S$ satisfying the following three axioms:

A1. $x \triangleright x = x$ for all $x \in S$
A2. For all $y, z \in S$, there exists a unique $x \in S$ such that $x \triangleright y = z$
A3. $(x \triangleright y) \triangleright z = (x \triangleright z) \triangleright (y \triangleright z)$ for all $x, y, z \in S$

Note that as the inverse operation $\triangleleft$ defined by $b \triangleleft a = x$ if and only if $x \triangleright a = b$ also fulfills the quandle axioms, a quandle is sometimes defined as a set with *two* binary operations: $\triangleright$ and $\triangleleft$. Distributive properties between the two operations, $(x \triangleright y) \triangleleft z = (x \triangleright z) \triangleleft (y \triangleright z)$ and $(x \triangleleft y) \triangleright z = (x \triangleleft z) \triangleright (y \triangleleft z)$, follow from the quandle axioms. For example, any group G is a quandle with the quandle product $\triangleright$ defined by conjugation, i.e. $a \triangleright b = bab^{-1}$.

The quandle structure can be associated with a knot in two distinct ways. The first method, which we shall refer to as the algebraic definition of the knot quandle, was first introduced by Joyce in [Jo]. In order to give the algebraic definition of the knot quandle, it is important to first understand the concept of a diagram of a knot.

Definition 2. A **diagram** D of a knot K is a projection of K onto a plane such that at most two strands of the knot intersect at any point and such that there are finitely many such intersections. By convention, at each crossing of two strands, one removes a segment of the projected image of the lower strand to convey relative height information. These breaks make the diagram a set of disjoint arcs. The orientation of the knot is also indicated on the diagram. Recall, from the introduction, that two diagrams represent the same knot if and only if they are connected by a finite chain of Reidemeister moves (see Figure 4.2).

Using this definition, it is possible to assign a quandle structure to any diagram of a knot.

Definition 3. Let D be a diagram of an oriented knot K. Let $A_D = \{\text{arcs of } D\}$. For all crossings P_i (as in Figure 4.1), define the relation $a \triangleright_A b = c$. Note that this definition depends on the orientation of the overcrossing arc but not the undercrossing arc. The **algebraic knot quandle** $\Gamma_A(K, D)$ of a given knot diagram D is then defined as $\langle A_D \mid \triangleright_A \rangle$. That is, it is the quandle with generators A_D subject to the relations $\triangleright_A$.

The algebraic knot quandle of a diagram does in fact prove to be a well-defined quandle. As we will show, it is independent of the chosen knot diagram; that is, it is a knot invariant.

Theorem 4. *The algebraic quandle, $\Gamma_A(K, D)$, is independent of D up to isomorphism.*

Proof. Let D_1, D_2 be two diagrams of a knot K. Since D_1 and D_2 are diagrams of the same knot, they are connected by a finite number of Reidemeister moves.

Consider the effect of each of these moves on the knot quandle (given an arbitrary choice of orientation), as in Figure 4.2. (Note that we show 8 such moves as opposed to the usual three, because we must have that the knot quandle is unchanged by Reidemeister moves regardless of orientation.) For each move, we can verify that the quandles obtained from the diagrams before

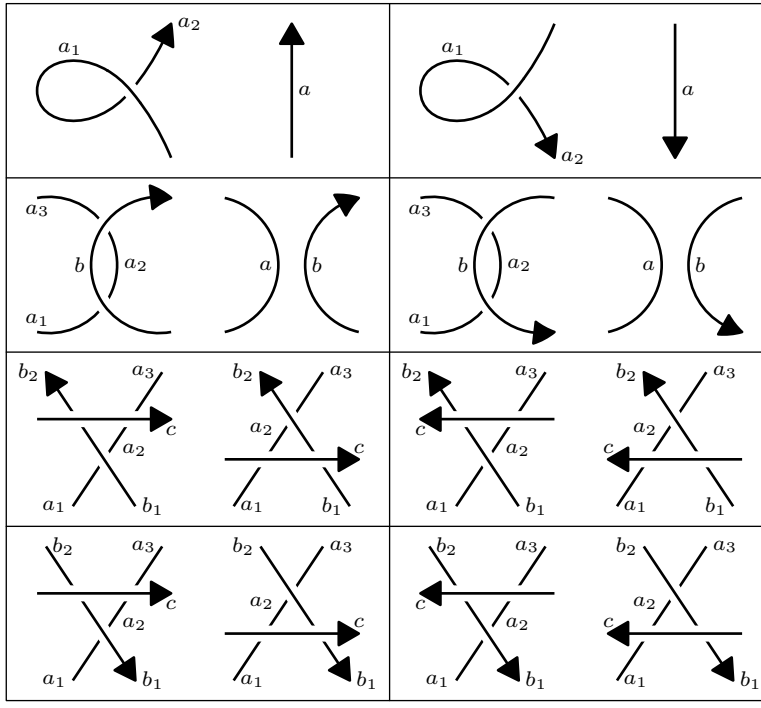

Figure 4.2: Reidemeister Moves.

and after the move are isomorphic. Consider the move depicted in the upper left-hand corner of Figure 4.2. Let Γ the the algebraic quandle corresponding to a knot including a strand a, as in the right-hand side of the diagram illustrating that move. Then Γ is generated by a and some collection of arcs $\{b, c, d, \ldots\}$, and is subject to some collection of relations. Let Γ' be the algebraic quandle corresponding to the same knot after the Reidemeister move has been performed, so that the knot diagram contains strands a_1 and a_2 as depicted on the left-hand side of the illustration. Then Γ' is generated by a_1, a_2 and $\{b, c, d, \ldots\}$, subject to the same relations as Γ (but with a_1 or a_2 substituted for occurrences of a in those relations, as appropriate) plus an additional relation $a_1 \triangleright a_1 = a_2$. But by quandle axiom A1, $a_1 \triangleright a_1 = a_1$. Hence $a_1 = a_2$, so the two quandles are manifestly isomorphic, by mapping $a_1 = a_2$ to a and each of the other generators to itself.

We leave it to the reader to check in a similar manner that the other Reidemeister moves also leave the algebraic quandle unchanged, up to canonical isomorphism. (One finds that each type of Reidemeister move corresponds to one of the quandle axioms.) Since D_1 and D_2 are related by a finite sequence of Reidemeister moves, it follows that $\Gamma_A(K, D_1) \cong \Gamma_A(K, D_2)$, as desired. □

There is a second way of associating a quandle structure with a knot. This geometric definition was proposed by Matveev in his paper on distributive groupoids [Mat].

Definition 5. Let K be an oriented knot in $\mathbb{R}^3$. Let $N(K)$ be a small tubular neighborhood of K. Let $E(K) = \overline{\mathbb{R}^3 \setminus N(K)}$. Fix a base point $x \in E(K)$. Let $B_K = \{\text{homotopy}^* \text{ classes of paths in } E(K) \text{ with fixed initial point } x \text{ and an endpoint on } \partial N(K)\}$. Here homotopy* means a homotopy which keeps the base point x fixed, such that the trajectory of the other endpoint is contained in $\partial N(K)$. The reader can check that being homotopic* is in fact an equivalence relation, which we henceforth denote by $\backsimeq$. Given a path b from x to some point y on $\partial N(K)$, let γ be the

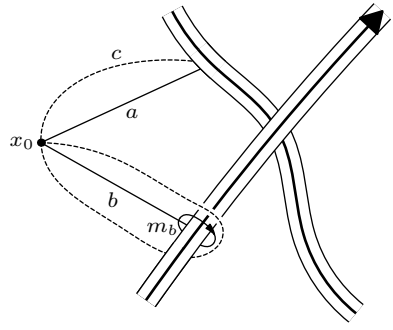

Figure 4.3: The loop c is a representative of $[a] \triangleright [b]$.

shortest line segment from y to the knot, and let the **oriented meridian** m_b be the loop in $\partial N(K)$ based at y, in the plane containing γ and perpendicular to the knot at $\gamma \cap K$.[1] Define an operation $\triangleright : B_K \times B_K \to B_K$ by $[a] \triangleright [b] \mapsto [bm_b b^{-1} a]$, as in Figure 4.3, where b^{-1} denotes the path b traversed backwards, and concatenation of paths denotes the standard composition law for paths from homotopy theory. The **geometric knot quandle** is defined as $\Gamma_B(K, x) = \langle B_D \mid \triangleright_B \rangle$. That is, it is the quandle generated by B_D with relations $\triangleright_B$.

The knot quandle, as defined by Matveev, satisfies the quandle axioms given in Definition 1. It also proves to be a knot invariant.

Theorem 6. *The geometric definition of the knot quandle is independent of the chosen base point, up to isomorphism.*

Proof. Choose two base points $x_1, x_2 \in \mathbb{R}^3 \setminus K$, for a fixed knot K. Let γ be a fixed path from x_2 to x_1; such a path exists because the complement of a knot is path-connected. Define $\phi : \Gamma_B(K, x_2) \to \Gamma_B(K, x_1)$ by $[\beta] \mapsto [\gamma\beta]$. There is a homotopy* $\beta_i \backsimeq \beta_j$ if and only if there is a homotopy* $\gamma\beta_i \backsimeq \gamma\beta_j$, so ϕ is well-defined. By the same token, there is also a well-defined map $\phi^{-1} : \Gamma_B(K, x_1) \to \Gamma_B(K, x_2)$ given by $\phi^{-1}[\beta] = [\gamma^{-1}\beta]$; it is easy to see that ϕ^{-1} inverts ϕ. If m_j denotes an oriented meridian of $N(K)$ at the endpoint of β_j, we can compute as follows, using the definition of $\triangleright$ in the geometric knot quandle:

$$\begin{aligned}
\phi[\beta_i] \triangleright \phi[\beta_j] &= [\gamma\beta_i] \triangleright [\gamma\beta_j] \\
&= [\gamma\beta_j m_j (\gamma\beta_j)^{-1} \gamma\beta_i] \\
&= [\gamma\beta_j m_j \beta_j^{-1} \beta_i] \\
&= \phi[\beta_j m_j \beta_j^{-1} \beta_i] \\
&= \phi([\beta_i] \triangleright [\beta_j]).
\end{aligned}$$

Therefore ϕ is a quandle homomorphism, and by the same reasoning so is ϕ^{-1}. Consequently ϕ is a quandle isomorphism. (Note that this isomorphism is not unique; it depends on the path chosen between x_1 and x_2.) □

There are therefore two distinct ways to associate a quandle structure with a knot. Both definitions give rise to a knot invariant (up to isomorphism). In fact, they give rise to the same invariant: the knot quandle.

[1] While it will not be important for our purposes, it is worth pointing out that the meridians m_b must be chosen with coherent orientations. One way to do this is to require that relative to a fixed orientation of the knot K, the meridians are oriented using the "right hand rule."

Theorem 7. *The algebraic and geometric definitions of the knot quandle are equivalent up to orientation.*

Sketch of proof. Let $\Gamma_A(K)$ be the algebraic knot quandle for a diagram D of K, and let $\Gamma_B(K)$ be the geometric knot quandle. Define a map $\phi : \Gamma_A(K) \to \Gamma_B(K)$ by mapping an arc $a \in \Gamma_A(K)$ to a homotopy class $[s_a]$ in $\Gamma_B(K)$ such that the following conditions hold: (1) The path s_a connects the base point x_0 to a point on the section of $\partial N(K)$ whose distance to the strand of the knot that projects to arc a is minimized; (2) at all points where the projection of s_a intersects D, the path s_a is *above* the strands of the knot it crosses, as in Figure 4.4.

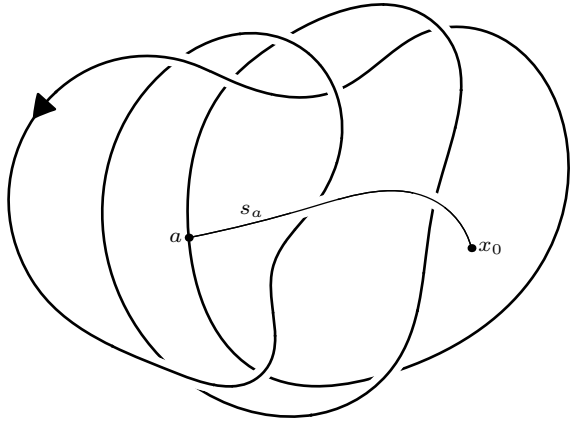

Figure 4.4: Choosing the Path s_a.

Having defined ϕ on the generators of $\Gamma_A(K)$, we extend it to a map on the quandle. It is possible to show that in fact ϕ is a quandle homomorphism.

Define $\psi : \Gamma_B(K) \to \Gamma_A(K)$ as follows. Given $[s] \in \Gamma_B(K)$ choose a representative s such that its projection onto D intersects the diagram in a finite number of points. Let $\{a_1, ..., a_n\}$ be the set of arcs in $\Gamma_A(K)$ that are above s in the diagram D, and let a_0 be the arc containing the endpoint of $[s]$; see Figure 4.5. Set

$$\psi([s]) = ((...((a_0\epsilon_1 a_1)\epsilon_2 a_2)\epsilon_3...)\epsilon_{n-1}a_{n-1})\epsilon_n a_n,$$

where ϵ_i is $\triangleright$ if the crossing between s and a_i is positive in D and ϵ_i is $\triangleleft$ if the crossing is negative. The sign of a crossing is defined as in Figure 4.6.

It is possible to show that ψ also defines a quandle homomorphism, and that ϕ and ψ are mutually inverse. So ϕ is an isomorphism between $\Gamma_A(K)$ and $\Gamma_B(K)$, so $\Gamma_A(K) \cong \Gamma_B(K)$, as desired. □

In order to gain a better understanding of the knot quandle, let us consider the example of the knot 5_2.

Example. Deriving the knot quandle $\Gamma(5_2)$.

Consider the standard diagram of the oriented knot 5_2 (see Figure 4.7a). We assign a generator to each of the arcs of the diagram. Using the algebraic definition of the knot quandle, we assign the relation $x_i \triangleright x_j = x_k$ to each of the crossings. This gives us $\Gamma_A(5_2) = \langle a, b, c, d, e \mid d \triangleright a = e, b \triangleright d = a, a \triangleright b = c, c \triangleright e = d, e \triangleright c = b \rangle$.

Now consider $E(K) = \overline{S^3 \setminus N(K)}$ for a small tubular neighborhood $N(K)$ of K. Fix a base point $x_0 \in E(K)$. The generators of the quandle $\Gamma_B(5_2)$ are the set of equivalence classes of

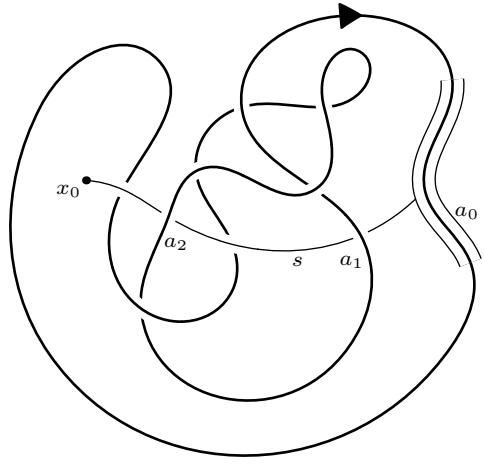

Figure 4.5: Choosing the Representative s.

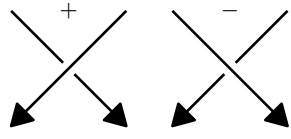

Figure 4.6: The Sign of a Crossing.

paths in $E(K)$, as shown in Figure 4.7b. Equivalently, the generators are the set of paths that end at a distinct arc of the diagram and go over all other arcs by the definition of ϕ in the proof of Theorem 7. The relations of the geometric quandle are given by $[x] \triangleright [y] = [ym_y y^{-1}x]$; therefore, the geometric definition generates the quandle

$$\Gamma_B(5_2) = \langle [a], [b], [c], [d], [e] \mid [d] \triangleright [a] = [am_a a^{-1} d] = [e], [b] \triangleright [d] = [dm_d d^{-1} b] = [a],$$
$$[a] \triangleright [b] = [bm_b b^{-1} a] = [c], [c] \triangleright [e] = [em_e e^{-1} c] = [d], [e] \triangleright [c] = [cm_c c^{-1} e] = [b] \rangle.$$

The equalities of the form $[am_a a^{-1} d] = [e]$ are all geometrically self-evident; given a parameterization of the knot and its tubular neighborhood, it would be straightforward to write down the corresponding homotopies.

The quandles $\Gamma_A(5_2)$ and $\Gamma_B(5_2)$ are trivially isomorphic, so in the case of 5_2 we can see explicitly that the algebraic and geometric quandles are isomorphic, as required by Theorem 7.

4.3 A Complete Invariant

Having established that the knot quandle is actually a well-defined knot invariant, we are left to wonder how useful this new invariant is. In order to answer this, we must look at two independent questions. First, how good is the knot quandle at distinguishing between knots? Second, how easy is it to show whether or not two knots have the same knot quandle? The answer to the first question turns out to be that the knot quandle is extremely good at distinguishing between knots. In fact, it is a complete invariant up to orientation.

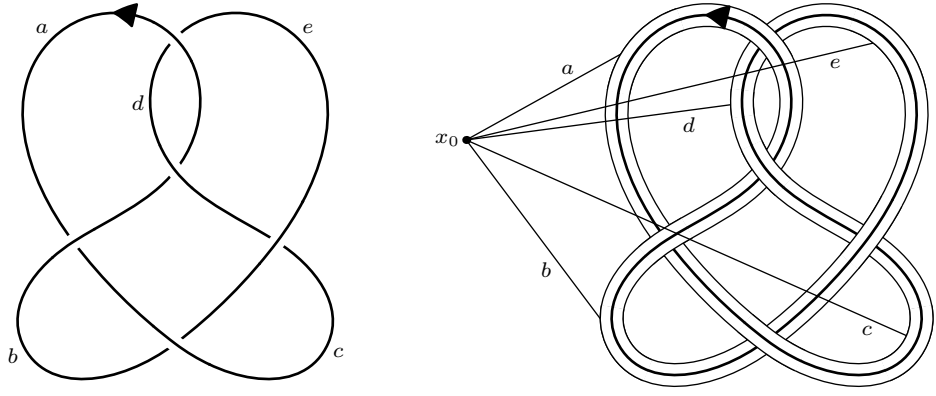

Figure 4.7: (a) Constructing $\Gamma_A(5_2)$; (b) Constructing $\Gamma_B(5_2)$.

In order to prove this, we need to first establish some definitions and background theorems.

Definition 8. A 2-surface F in a 3-manifold M is **compressible** if one of the following conditions holds:

1. F is an embedding of a 2-sphere and bounds the embedding of a 3-ball in M,
2. F is the embedding of a disk in ∂M,
3. F is the embedding of a disc in M and there is an embedded 3-ball in M whose boundary is contained in $F \cup \partial M$,
4. F is not the embedding of a 2-sphere or a disc and there exists an embedded disk $\Delta \subset M$ such that $\Delta \cap F = \partial\Delta$ and such that $\partial\Delta$ is not **nullhomotopic** in F (that is, it is a nontrivial element of the fundamental group $\pi_1(F)$).

Definition 9. If every embedding of a 2-sphere in M is compressible, then M is **irreducible**.

Definition 10. A 3-manifold with boundary M is **boundary irreducible** if its boundary ∂M is not compressible.

Definition 11. A **handlebody** in $\mathbb{R}^n$ is a closed, regular neighborhood of a finite graph in $\mathbb{R}^n$.

Definition 12. A manifold M is called **sufficiently large** if one can embed a handlebody $H \neq S^2$ in M such that the map of fundamental groups induced by the inclusion $H \hookrightarrow M$ is injective.

The proof of the completeness of the knot quandle also relies on four important theorems.

Theorem 13 (Dehn's Lemma)**.** *Let M be a 3-manifold with boundary and let γ be a closed curve on its boundary ∂M. Then if there exists an immersed 2-disc $D \to M$ such that $\partial D = \gamma$, then there exists an embedded disk $D' \subset M$ with the same boundary $\partial D' = \gamma$.*

This theorem was first proven in 1910 by Max Dehn, a German mathematician, but his proof was later discovered to contain holes. It was finally rigorously proven by Christos Papakyriakopoulos in 1956. One of the important consequences of Dehn's lemma (from a knot theoretic point of view) is that it can be used to prove Dehn's Theorem.

Theorem 14 (Dehn's Theorem)**.** *A knot K is the unknot if and only if $\pi_1(\mathbb{R}^3 \setminus K)$ is isomorphic to $\mathbb{Z}$.*

Proof. One direction of Dehn's theorem is not too bad; using the Wirtinger presentation defined in Definition 26 (or see [GP]), the fundamental group of the complement of the unknot has one generator and no relations, and is therefore isomorphic to $\mathbb{Z}$.

The other direction follows from Dehn's lemma. Let K be a knot such that $\pi(\mathbb{R}^3 \setminus K) = \mathbb{Z} = \langle x \rangle$ where x is a loop on $\partial N(K)$. Let μ be a meridian of $N(K)$ and let λ be a **longitude** (a loop in $N(K)$ not in a homotopy class generated by a meridian). The fundamental group $\pi_1(\mathbb{R}^3 \setminus K)$ clearly contains $\langle \mu \rangle = \mathbb{Z}$, so the curve λ must be homotopic to a constant, which implies that there is an immersed disc Δ with $\partial\Delta = \lambda$. By Dehn's lemma, there is a disc embedded in $\mathbb{R}^3 \setminus N(K)$ bounded by λ. By contracting $N(K)$ to K along the radial chords, we obtain a disc embedded in $\mathbb{R}^3 \setminus K$ bounded by K, so K is the unknot. □

Definition 15. Given two knots K_1, K_2 with associated meridians and longitudes m_1, m_2, l_1, l_2 we say a homomorphism $\phi : \pi_1(\mathbb{R}^3 \setminus K_1) \to \pi_1(\mathbb{R}^3 \setminus K_2)$ **preserves peripheral structure** if the image of $\langle [m_1], [l_1] \rangle \subset \pi_1(\mathbb{R}^3 \setminus K_1)$ through ϕ is conjugate to a subgroup of $\langle [m_2], [l_2] \rangle \subset \pi_1(\mathbb{R}^3 \setminus K_2)$. We say that $\langle [m_1], [l_1] \rangle \subset \pi_1(\mathbb{R}^3 \setminus K_1)$ is the **peripheral structure** of the knot K_1, and similarly with K_2; note that we may also view this as simply the normalizer of $[m_1]$, i.e. $N([m_1])$. This definition can be generalized to maps between manifolds—the Waldhausen theorem below holds for this more general definition.

Theorem 16 (Waldhausen Theorem)**.** *Let M, N be irreducible, boundary-irreducible 3-manifolds. Let M be sufficiently large and let $\phi : \pi_1(N) \to \pi_1(M)$ be an isomorphism preserving peripheral structure. Then there exits a homeomorphism $f : N \to M$ that induces ϕ.*

The final theorem was first suggested in 1908, although it remained unproven until 1989.

Theorem 17 (Gordon-Luecke Theorem)**.** *If K_1 and K_2 are unoriented knots in S^3 and there is an orientation preserving homeomorphism between their complements, then K_1 and K_2 are equivalent as unoriented knots.*

Now, for convenience, we will establish some useful lemmas, some of whose proofs we omit, before we proceed to the proof that the knot quandle is a complete invariant.

Lemma 18. *Let K be a non-trivial knot. Let $N(K)$ be a small tubular neighborhood of K, and let $E(K)$ be the closure of $S^3 \setminus N(K)$. Then $E(K)$ is irreducible, boundary irreducible, and sufficiently large.*

Lemma 19. *If K is non-trivial, then there is an injective homomorphism from the fundamental group $\pi_1(\partial N)$ into $\pi_1(\mathbb{R}^3 \setminus K)$.*

Lemma 20. *The peripheral structure of a knot is determined by the knot quandle.*

Proof. The group $\pi_1(\mathbb{R}^3 \setminus K)$ is determined by the knot quandle, as is shown in Section 4.5. We will show that the homotopy class of a meridian of $N(K)$ can be constructed from elements of the geometric knot quandle. Choose any $[x], [y] \in \Gamma(K)$, not necessarily distinct; we can do this because the geometric quandle is nonempty by definition. Now computing in the fundamental groupoid of $E(K)$, define a homotopy class by $[x^{-1}]([y] \triangleright [x])[y^{-1}][x]$. Using the definition of $\triangleright$ in the geometric quandle, we see that this is equal to the homotopy class $[m_x]$ of the meridian m_x associated to $[x]$. Since m_x is a loop, this homotopy class is an element of the fundamental group $\pi_1(\mathbb{R}^3 \setminus K)$. As the knot quandle determines the fundamental group of the knot complement (see Section 4.5) and the homotopy class of a meridian, it determines the normalizer of $[m_x]$ in $\pi_1(\mathbb{R}^3 \setminus K)$, which is precisely the peripheral structure. □

Using this background, it is possible to prove that the knot quandle is a complete invariant up to orientation.

Theorem 21. *The knot quandle $\Gamma(K)$ is a complete invariant up to orientation.*

Proof. Let K_1, K_2 be two knots such that there exists an isomorphism $\psi : \Gamma(K_1) \to \Gamma(K_2)$.

- If K_1 is trivial: The fundamental group of the complement, $\pi_1(\mathbb{R}^3 \setminus K)$ can be derived from the knot quandle $\Gamma(K)$. Therefore, since $\Gamma(K_1) \cong \Gamma(K_2)$, we know that $\pi_1(\mathbb{R}^3 \setminus K_1) \cong \pi_1(\mathbb{R}^3 \setminus K_2)$. The knot K_1 is trivial, so by Dehn's theorem, $\pi_1(\mathbb{R}^3 \setminus K_1) \cong \mathbb{Z}$, which implies that $\pi_1(\mathbb{R}^3 \setminus K_2) \cong \mathbb{Z}$. So by Dehn's theorem K_2 is also trivial.

- If K_1, K_2 are nontrivial: Since $\Gamma(K_1) \cong \Gamma(K_2)$ and the quandle determines the peripheral structure (Lemma 20), K_1 and K_2 have the same peripheral structure. Therefore the isomorphism $\phi : \pi_1(\mathbb{R}^3 \setminus K_1) \to \pi_1(\mathbb{R}^3 \setminus K_2)$ induced by the isomorphism $\psi : \Gamma(K_1) \to \Gamma(K_2)$ preserves peripheral structure. Also, E_1, E_2 are boundary-irreducible, sufficiently large, irreducible 3-manifolds (Lemma 18).

  Therefore, by the Waldhausen theorem, there exists a homeomorphism $f : E(K_1) \to E(K_2)$ that induces ϕ. It follows immediately that there is a homeomorphism between $\mathbb{R}^3 \setminus K_1$ and $\mathbb{R}^3 \setminus K_2$ to which we may apply the Gordon-Luecke theorem, to conclude that K_1 and K_2 are equivalent up to orientation. □

It is again important to note that the knot quandle is not a truly complete invariant; it is only complete up to the orientations of the knot and the ambient space.

Example. Consider both the right- and left-handed trefoils (see Figure 4.8). We know that these are not equivalent knots because their signatures (an invariant defined in [Cr]) are different ($\sigma(3_1) = -2$ and $\sigma(3_1^*) = 2$). However, these two knots have isomorphic knot quandles.

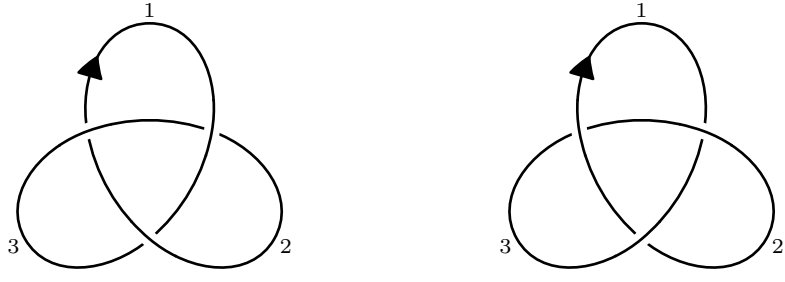

Figure 4.8: (a) The right-handed trefoil knot 3_1; (b) the left-handed trefoil knot 3_1^*.

Using the algebraic definition of the knot quandle, we get that:

$$\Gamma(3_1) = \langle 1, 2, 3 \mid 1 \triangleright 3 = 2, 3 \triangleright 2 = 1, 2 \triangleright 1 = 3 \rangle.$$

Furthermore:

$$\Gamma(3_1^*) = \langle 1, 2, 3 \mid 2 \triangleright 3 = 1, 1 \triangleright 2 = 3, 3 \triangleright 1 = 2 \rangle.$$

Using the quandle axioms, it is easy to see that these give rise to the same quandle. Consider $\Gamma(3_1)$. The relation $1 \triangleright 2$ is not explicitly defined, but $2 \triangleright 2 = 2$ by A1, $3 \triangleright 2 = 1$, and there exists a unique x such that $x \triangleright y = z$ for any y, z (by A2), so we must have $1 \triangleright 2 = 3$. Similarly, we find that $2 \triangleright 3 = 1$ and $3 \triangleright 1 = 2$, and hence $\Gamma(3_1) = \Gamma(3_1^*)$ despite the face that 3_1 and 3_1^* are not equivalent knots.

Nonetheless, the knot quandle is an extremely powerful invariant. However, in order to be computationally useful, we would have to be able to efficiently determine whether or not two knots give rise to isomorphic knot quandles.

4.4 Computational Complexity

Unfortunately, despite the fact that the knot quandle is a complete invariant up to orientation, it does not turn out to be easy to use. The primary problem with the knot quandle is that it is difficult to determine whether or not two quandles are isomorphic. One approach to this problem has been suggested by Sam Nelson and Benita Ho [HN].

Definition 22. The **quandle matrix** associated with a finite quandle Q with n elements, M_Q, is the $n \times n$ matrix whose (i, j)-th entry is given by $x_i \triangleright x_j$,

$$M_Q = \begin{pmatrix} x_1 \triangleright x_1 & \dots & x_1 \triangleright x_n \\ \vdots & \ddots & \vdots \\ x_n \triangleright x_1 & \dots & x_n \triangleright x_n \end{pmatrix}.$$

We can define an equivalence relation on quandle matrices.

Definition 23. Let $\rho \in S_n$ be a permutation of $\{1, ..., n\}$. Set $\rho(M_Q) = A_\rho^{-1}(\rho a_{ij})A_\rho$ where A_ρ is the permutation matrix of ρ and ρa_{ij} is the image of the (i, j)-th entry of M_Q under the permutation ρ, acting on the elements $\{x_1, \dots, x_n\}$ of Q in the natural way. Then we say $\rho(M_Q)$ is **permutation equivalent** or **p-equivalent** to M_Q.

This equivalence allows us to determine whether or not two knot quandles are isomorphic.

Theorem 24. *Two $n \times n$ quandle matrices determine isomorphic quandles if and only if they are p-equivalent by a permutation $\rho \in S_n$.*

Proof. Let $Q = \langle x_1, \dots, x_n \mid \triangleright \rangle$ and $Q' = \langle y_1, \dots, y_n \mid \triangleright \rangle$ be finite quandles and let $M_Q, M_{Q'}$ be their respective matrices. Let $\psi : Q \to Q'$ be an isomorphism of finite quandles. We have that ψ, considered on subscripts, induces a bijection $\rho : \{1, ..., n\} \to \{1, ..., n\}$, so $\rho \in S_n$. The fact that ψ is an isomorphism gives us that $\psi(x_i \triangleright x_j) = \psi(x_i) \triangleright \psi(x_j) = x_{\rho(i)} \triangleright x_{\rho(j)}$, so we obtain a permutation of the quandle matrix for Q' by applying $\rho \in S_n$ to every element in M_Q. Conjugation by the permutation matrix of ρ puts the matrix back in standard form, yielding the equality $\rho(M_Q) = A_\rho^{-1}(\rho(a'_{ij}))A_\rho$ where a'_{ij} is an element in the matrix $M_{Q'}$, as desired.

The argument for the reverse implication is essentially identical. □

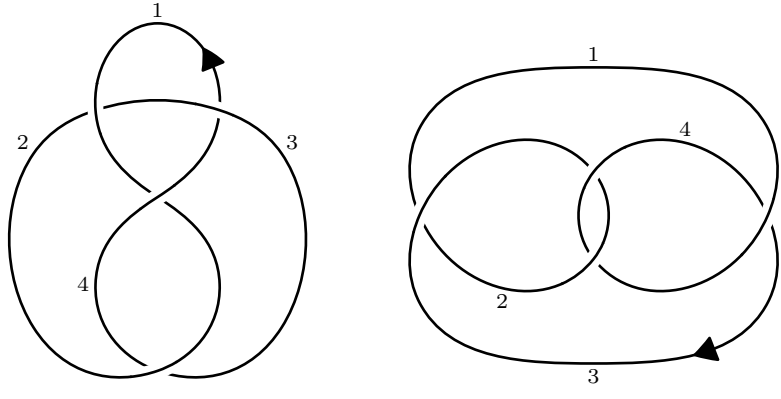

Figure 4.9: Two presentations of the knot 4_1: (a) the knot D_1; (b) the knot D_2.

Example. Consider the two representations of 4_1 in Figure 4.9. Using the algebraic definition of the knot quandle, we derive the following two quandles: $\Gamma(D_1) = \langle 1,2,3,4 \mid 3 \triangleright 1 = 2, 1 \triangleright 4 = 2, 3 \triangleright 2 = 4, 1 \triangleright 3 = 4\rangle$ and $\Gamma(D_2) = \langle 1,2,3,4 \mid 1 \triangleright 3 = 2, 1 \triangleright 2 = 4, 3 \triangleright 4 = 2, 3 \triangleright 1 = 4\rangle$.

We have the two quandle matrices

$$M_1 = \begin{pmatrix} 1&3&4&2\\ 4&2&1&3\\ 2&4&3&1\\ 3&1&2&4 \end{pmatrix},\ M_2 = \begin{pmatrix} 1&4&2&3\\ 3&2&4&1\\ 4&1&3&2\\ 2&3&1&4 \end{pmatrix}.$$

Define $\rho = (1)(2)(34) \in S_4$. Then

$$\begin{aligned} A_\rho^{-1}(\rho(M_1))A_\rho &= \begin{pmatrix} 1&0&0&0\\ 0&1&0&0\\ 0&0&0&1\\ 0&0&1&0 \end{pmatrix}^{-1} \rho\left(\begin{pmatrix} 1&3&4&2\\ 4&2&1&3\\ 2&4&3&1\\ 3&1&2&4 \end{pmatrix}\right)\begin{pmatrix} 1&0&0&0\\ 0&1&0&0\\ 0&0&0&1\\ 0&0&1&0 \end{pmatrix} \\ &= \begin{pmatrix} 1&0&0&0\\ 0&1&0&0\\ 0&0&0&1\\ 0&0&1&0 \end{pmatrix}\begin{pmatrix} 1&4&3&2\\ 3&2&1&4\\ 2&3&4&1\\ 4&1&2&3 \end{pmatrix}\begin{pmatrix} 1&0&0&0\\ 0&1&0&0\\ 0&0&0&1\\ 0&0&1&0 \end{pmatrix} = \begin{pmatrix} 1&4&2&3\\ 3&2&4&1\\ 4&1&3&2\\ 2&3&1&4 \end{pmatrix} = M_2. \end{aligned}$$

So by Theorem 24, $\Gamma(D_1) \cong \Gamma(D_2)$, as M_1 is p-equivalent to M_2. As a result, $K(D_1) = K(D_2)$ up to orientations; they are both diagrams of the knot 4_1. (Note that the ambient isotopy between the two representations is easy to visualize explicitly: lift strand 3 in Figure 4.9(b) and pull it up to the top of the picture.)

This method is reasonably effective; it allows us to determine whether two knots are equivalent if they generate small, finite, easily understood quandles. The problem is that determining whether two matrices are permutation-equivalent is thought to be a computationally intractable problem.

There is no currently known algorithm for determining whether or not two matrices are permutation equivalent in polynomial time. Therefore, as the size of the knot quandle increases, determining whether two knot quandles are isomorphic becomes difficult. Since all small knots have already been completely classified, this renders the knot quandle relatively useless as a computational tool.

The difficulty in actually making use of the quandle, despite the fact that it is a complete invariant up to orientation, has caused mathematicians to resort to less complete, but more useful, invariants that can be derived from the knot quandle.

4.5 The Knot Group

The concept of the fundamental group gives rise to the first such knot invariant: the knot group.

Definition 25. The **knot group** of a knot K is defined as the fundamental group of its complement, $\pi_1(\mathbb{R}^3 \setminus K)$. For convenience, however, we shall denote the knot group by $\pi_1(K)$.

Just as there are two definitions for the knot quandle, there are two equivalent definitions for the knot group.

Definition 26. Let K be a knot. Let D be an oriented diagram of K. Let A_D be the set of arcs of D. For each crossing P_i as in Figure 4.1, define the relation R_i by $c = bab^{-1}$. The **knot group** of K with respect to D is the group $\langle A_D \mid R_i\rangle$. This representation of the knot group is called the **Wirtinger presentation**.

To prove that the Wirtinger presentation of a knot is independent of the choice of diagram, one can simply check that it is invariant under Reidemeister moves. The proof that the Wirtinger presentation gives rise to the same group as the fundamental group of the complement of a knot is more complex and is omitted.

Using the Wirtinger presentation, we can show that the knot quandle determines the knot group up to isomorphism. In this manner, the knot quandle gives rise to an invariant that is easier to calculate than the quandle itself, but which is unfortunately less exact; while the knot quandle is a complete invariant up to orientation, the knot group is not.

Theorem 27. *The knot group $\pi_1(K)$ is determined by the knot quandle $\Gamma(K)$.*

Proof. Let K be a knot. Choose a diagram D of K. We will define an association ψ from knot quandles to knot groups. Let $\Gamma(K) = \Gamma_A(K, D) = \langle A_D \mid R_i \rangle$. Define a group $\psi(\Gamma(K)) = G = \langle A_D \mid S_i \rangle$ where S_i is defined as follows: if R_i is the relation $a \triangleright b = c$, then S_i is the relation $bab^{-1} = c$. By definition, G is the Wirtinger presentation of the knot group, so the knot group can be derived from the knot quandle. □

This method of determining the knot group given a knot quandle can be more easily seen by working through the example of the knot 5_2.

Example. Consider the knot 5_2, as in Figure 4.7a. As seen in Section 4.2,

$$\Gamma(5_2) = \langle a, b, c, d, e \mid d \triangleright a = e, b \triangleright d = a, a \triangleright b = c, c \triangleright e = d, e \triangleright c = b \rangle.$$

The group $\psi(\Gamma(5_2))$ determined from $\Gamma(5_2)$ as in the proof of Theorem 27 is $\langle a, b, c, d, e \mid ada^{-1} = e, dbd^{-1} = a, bab^{-1} = c, ece^{-1} = d, cec^{-1} = b \rangle$. This is exactly the knot group $\pi_1(5_2)$, as can easily be checked using Wirtinger presentation.

However, because the information about the peripheral structure is not retained by the map ψ from the knot quandle to the knot group, the latter is not a complete invariant.

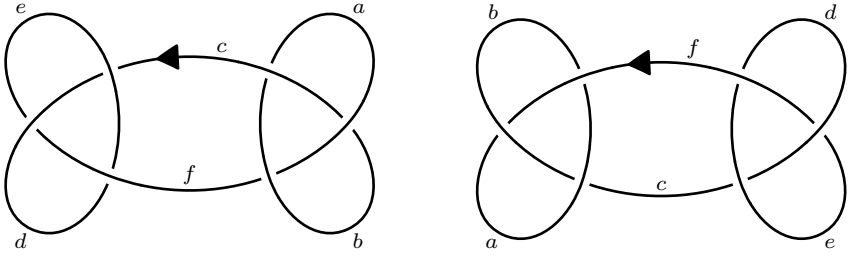

Figure 4.10: Two inequivalent knots with the same knot group: (a) $3_1\#3_1$, and (b) $3_1\#3_1^*$.

Example. Consider the diagrams of the knots $3_1\#3_1$ and $3_1\#3_1^*$, in Figure 4.10.

Using Wirtinger presentations, we calculate that:

$$\pi_1(3_1\#3_1) = \langle a, b, c, d, e, f \mid aba^{-1} = c, faf^{-1} = b, bfb^{-1} = a, cdc^{-1} = e, ece^{-1} = e, ded^{-1} = f \rangle,$$
$$\pi_1(3_1\#3_1^*) = \langle a, b, c, d, e, f \mid aba^{-1} = c, faf = b, bfb^{-1} = a, ede^{-1} = c, dfd^{-1} = e, fef^{-1} = d \rangle.$$

These knot groups simplify to:

$$\pi_1(3_1\#3_1) = \langle b, c, d \mid bcb = cbc, cdc = dcd \rangle,$$
$$\pi_1(3_1\#3_1^*) = \langle a, e, f \mid afa = faf, efe = fef \rangle.$$

These knot groups are clearly isomorphic, despite the fact that the knots $3_1\#3_1$ and $3_1\#3_1^*$ are inequivalent. Therefore, as claimed, the knot group (unlike the knot quandle from which it is derived) is not a complete invariant up to orientation.

4.6 Colorability

The second useful invariant that can be derived from the knot quandle is colorability.

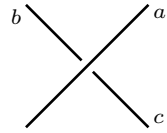

Figure 4.11: Each crossing in a 3-coloring must satisfy $2a - b - c \equiv 0 \mod 3$.

Definition 28. A diagram, D, of a knot K is **3-colorable** if its arcs can be labeled with elements of the color set $\{0, 1, 2\}$ such that at each crossing P_i we have the relation $2a - b - c \equiv 0 \mod 3$, as in Figure 4.11, and at least two distinct colors are used across the entire knot.

Note that because we are dealing with only three colors, we can simplify out understanding of 3-colorability by saying that a knot diagram D is 3-colorable if at each crossing P_i either the arcs a, b, c are all the same color or they are all different colors and such that at least two distinct colors are used across the entire knot. As you can check, these two definitions are equivalent.

Theorem 29. *3-colorability is a knot invariant.*

Proof. Consider each of the Reidemeister moves in Figure 4.2. With appropriate relabeling (in the case of the second move), you can check that none of these moves change the 3-colorability of the diagram. So 3-colorability is a knot invariant, as claimed. □

It turns out that it is possible to derive the 3-colorability of a knot K from the knot group $\pi_1(K)$.

Theorem 30. *A knot K is 3-colorable if and only if there exists a surjective homomorphism $\phi : \pi_1(K) \to D_3$ from the knot group to the dihedral group on three elements.*

We omit the proof, as it is technical and has little to do with the knot quandle.

This method of deriving the colorability of a knot from the knot group can be taken one step further to give us a way to derive the colorability directly from the knot quandle.

3-colorability is, however, a very weak invariant. For example, it fails to distinguish between the trefoil and the cinquefoil 5_1 (neither of them is 3-colorable). However, 3-colorability can be generalized to p-colorability for any prime p. This defines a labeling of the arcs of the diagram by the elements of $\{0, ..., p-1\}$. The crossing relations are defined again by $2x-y-z = 0 \mod p$. As in the case of 3-coloring, a knot is p-colorable if and only if there exists a surjective homomorphism between the knot quandle and the dihedral group D_p. Therefore, by combining colorability by p colors for all prime p, one defines a much stronger invariant that is, in some sense, determined by the knot quandle (or the knot group).

4.7 Conclusion

A quandle structure can be naturally associated with any knot. This gives rise to the knot quandle, an invariant that is complete up to orientation. Although it is unable to distinguish between the right- and left-handed trefoils, the knot quandle can distinguish between all knots that are not related by a change in orientations, including mutants (for example $3_1\#3_1$ and $3_1\#3_1^*$).

Although the knot quandle is extremely powerful, it is not computationally useful because it is difficult to determine when two quandles are isomorphic. Therefore, mathematicians spend more time considering the weaker but more useful invariants that can be derived from the knot quandle. Two of these invariants, the knot group and p-colorability, have already been discussed. Other invariants can also be derived from the knot quandle, including the Alexander and Conway polynomials.

Despite its limited usefulness, the knot quandle proves to be a interesting invariant. Not only does it provide an almost complete invariant, but it also serves as a generalization for a collection of more familiar (and more useful) knot invariants.

Acknowledgments

I am very grateful to Professor Elizabeth Denne, under whose guidance this paper was written. I would also like to acknowledge the aid and advice of Zachary Abel, Scott Kominers, Sam Lichtenstein, Daniel Litt, Alison Miller, and Charles Nathanson. I am especially indebted to Daniel Litt for his invaluable assistance throughout the editorial process.

References

[Cr] Peter Cromwell: *Knots and Links*. Cambridge, UK: Cambridge University Press, 2004.

[GP] N.D. Gilbert and T. Porter: *Knots and Surfaces*. New York: Oxford University Press, 1994.

[HN] Benita Ho and Sam Nelson: Matrices and finite quandles, *Homology, Homotopy and Applications* **7**#1 (2005), 197–208.

[Jo] David Joyce: A classifying invariant of knots, the knot quandle, *Journal of Pure and Applied Algebra* **23**#1 (1982), 37–65.

[Lic] W.B. Raymond Lickorish: *An Introduction to Knot Theory*. New York: Spring-Verlag, 1997.

[Liv] Charles Livingston: *Knot Theory*. Washington D.C.: Mathematical Association of America, 1993.

[Man] Vassily Manturov: *Knot Theory*. Boca Raton, Florida: CRC Press, 2004.

[Mat] Sergei Matveev: Distributive Groupoids In Knot Theory, *Math. USSR Sb*. **47**#1 (1984), 73-83.

[Mu] James R. Munkres: *Topology*. Upper Saddle River, NJ: Prentice Hall, 2000.

STUDENT ARTICLE

5 Problems of Circle Tangency

Gregory Minton†
Harvey Mudd College '08
Claremont, CA 91711
gminton@hmc.edu

Abstract

This article presents a (very) brief overview of geometric problems involving tangent circles. In addition to defining the technique of inversion, we give two example problems with full solutions and suggest another challenge problem related to Pappus circles.

5.1 Introduction

Geometric problems dealing with tangent circles have a long history and arise in surprising places. For example, the broad challenge "given three objects (where an object may be a line, circle, or point), draw a circle which is tangent to each" is known as **Apollonius' Problem** after the 3rd century BCE Greek geometer Apollonius, who wrote two works considering the problem. The **Descartes circle theorem** is a special case (the hardest special case) of Apollonius' Problem [Cox]. During Japan's isolationist period between the mid 17th and 19th centuries, inscribed geometry problems known as **sangaku** which often dealt with circle tangency were hung from religious buildings [RF]. In more modern use, a particular set of circles with rational centers known as **Ford circles** may be used to prove the **Hurwitz theorem**:

Theorem 1 (Hurwitz Theorem)**.** *If $k \geq 1/\sqrt{5}$, then for each irrational number w there are infinitely many fractions p/q satisfying*

$$\left|\frac{p}{q} - w\right| < \frac{k}{q^2}. \tag{5.1}$$

If $k < 1/\sqrt{5}$, then there exist irrationals w for which (5.1) has only finitely many solutions $\frac{p}{q} \in \mathbb{Q}$.

For more information on Ford circles, the reader is referred to L.R. Ford's original article [Fo].

Instead of providing a full historical overview of the subject or presenting new connections, this article demonstrates two particular problems of circle tangency and solves them using two very different strategies. Along the way, we will encounter **inversions of the plane**, an operation in planar geometry which is interesting but often overlooked. The paper is divided into three sections. First, we present and solve one problem without inversion. We then define inversions and list their important properties. Finally, we present another problem and solve it using inversion. The author hopes these two examples will help convince the reader of the beauty of the subject.

†Gregory Minton is a senior at Harvey Mudd College, Claremont, California who is looking forward to pursuing a PhD in graduate school. Some of his mathematical interests include complex analysis and representation theory. When not working on math, he enjoys playing racquetball and shooting pool.

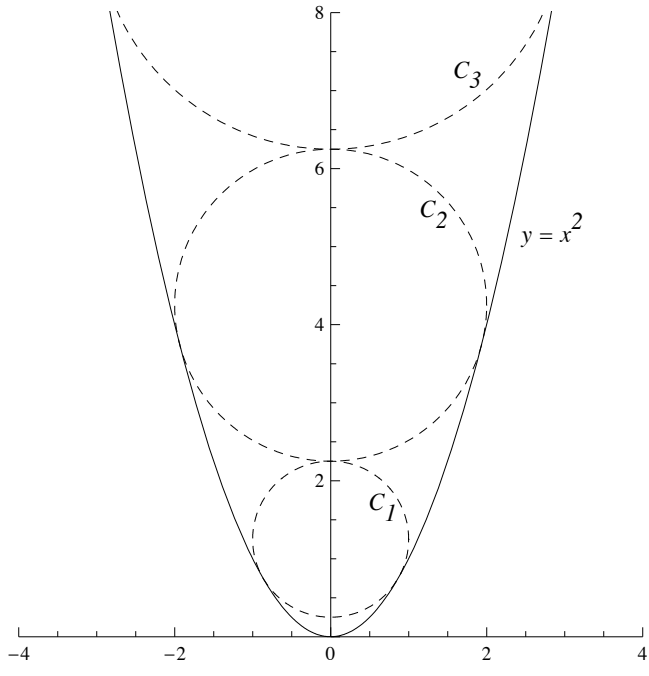

Figure 5.1: Mutually Tangent Circles Inscribed in a Parabola.

5.2 Circles in the Parabola

This section is dedicated to a solution of the following problem. To the author's knowledge, this problem was first printed as Exercise 63/64 in Section 7.1 of [SM].

Problem 1. *Let P be the parabola $y = x^2$, and let C_1 be the circle of radius 1 which is tangent to P at two points. Iteratively define a sequence of circles $\{C_2, C_3, \dots\}$, where each C_n is located above and tangent at a point to C_{n-1} and tangent at two points to P. Find the radius of C_n.*

Solution. We present an analytic solution which relies only on high school algebra techniques. Let C_n be the circle centered at (x_n, y_n) and having radius r_n. Notice that by symmetry, we must have $x_n = 0$ for C_n to be tangent to P at two points. Now the curves C_n, P are described by the equations

$$x^2 + (y - y_n)^2 = r_n^2, \tag{5.2}$$

$$y = x^2. \tag{5.3}$$

Substituting (5.3) into (5.2) gives the intersection condition

$$y + (y - y_n)^2 = r_n^2$$

$$y^2 + (1 - 2y_n)y + (y_n^2 - r_n^2) = 0. \tag{5.4}$$

Tangency of the curves means that the intersection equation (5.4) must have a double root. This implies the discriminant is zero:

$$0 = (1 - 2y_n)^2 - 4(1)(y_n^2 - r_n^2) = 1 - 4y_n + 4r_n^2,$$

so $y_n = \frac{1}{4} + r_n^2$.

Since C_n and C_{n-1} are tangent, the sum of their radii equals the distance between their centers. Writing this equation and substituting the above relation between y_n and r_n, we get

$$\begin{aligned} r_n + r_{n-1} &= y_n - y_{n-1} \\ r_n + r_{n-1} &= \left(\frac{1}{4} + r_n^2\right) - \left(\frac{1}{4} + r_{n-1}^2\right) \\ r_n + r_{n-1} &= r_n^2 - r_{n-1}^2 \\ 1 &= r_n - r_{n-1}. \end{aligned}$$

Thus, since $r_n = r_{n-1} + 1$ and $r_1 = 1$, the radius r_n of the nth circle is n. □

5.3 Inversions in the Plane

Inversion in the plane is a geometric "reflection" technique that uses a circle in lieu of a linear axis of symmetry. The formal definition of such a transformation is

Definition 2. Given a circle C with center O and radius r, the **inversion about** C is the map taking each point P to the point on ray $\overrightarrow{OP}$ of distance $r^2/|OP|$ from O.

Notice that inversion about the circle C fixes all the points of C and exchanges the groups of points that lie inside and outside the circle. In particular, the center O is taken to infinity. Because it is not well-defined at the center of the circle, inversion should really be viewed as a map on the **one-point compactification** of the plane formed by adding a point at infinity. This level of detail is not important for our purposes, but we shall speak occasionally about the point at infinity.

Notice that we obtain the identity transformation if we perform two inversions about the same circle; thus inversions are bijections of the (compactification of the) plane. Furthermore, though we will only employ inversions of the plane $\mathbb{R}^2$, it is worth noting that inversions may be defined naturally in higher-dimensional space $\mathbb{R}^n$, where the object to be inverted about is an $(n-1)$-dimensional scaling and translate of the sphere S^{n-1}. If we consider the inversions about two different circles centered at the same point, then the inversions are dilations of each other. Thus, the radius is not as interesting as the center; we often refer to inversion about a point P as shorthand for inversion about the circle of radius 1 centered at P.

The following result gives the primary use of inversions for circle problems:

Theorem 3. *Any inversion in the plane maps circles to circles, where lines are viewed as circles through infinity.*

Thus, since inversions are bijections which preserve the class of circles, they map a set of tangent circles to another set of tangent circles. Clever choice of the center for inversion could help make this new set particularly easy to work with.

The interested reader is referred to [Coo] for a more in-depth discussion of inversion, as well as a proof of Theorem 3.

5.4 Fibonacci Circles

This section is dedicated to an inversion solution to the following problem. This problem may also be found as Exercise 61/62 in Section 7.1 of [SM]; however, it is not original to this text.

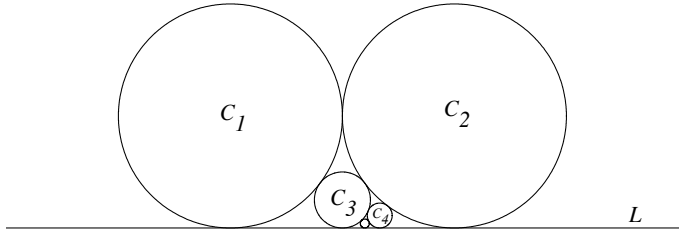

Figure 5.2: Fibonacci Circles.

Problem 2. *Let L be a line, and let C_1, C_2 be two circles of radius 1 tangent to each other and both tangent to L. Define a sequence of circles $C_3, C_4, \ldots$ where C_n is tangent to C_{n-1}, C_{n-2}, and L. Find the radius of C_n.*

Solution. Label by r_n the radius of C_n; thus $r_1 = r_2 = 1$. Consider now C_n, $n \geq 3$, and let P be the point at which C_{n-1} intersects C_{n-2}. Let I be the inversion of the plane about P. Since inversions are bijections taking circles to circles, C_{n-1} and C_{n-2} must each map to circles which only intersect at $I(P)$, the point at infinity. Thus, $I(C_{n-1})$ and $I(C_{n-2})$ are parallel lines.

The minimum distance from P to $I(C_{n-1})$ is the inverse of the maximum distance from P to C_{n-1}, which is $2r_{n-1}$. Similarly, the maximum distance from P to C_{n-2} is $2r_{n-2}$, so $I(C_{n-1})$ and $I(C_{n-2})$ are parallel lines with separation $(2r_{n-1})^{-1} + (2r_{n-2})^{-1}$.

Since L is a line (circle through infinity) which is tangent to C_{n-1} and C_{n-2}, $I(L)$ is a circle through P tangent to the lines $I(C_{n-1})$ and $I(C_{n-2})$. Further, since C_n is the circle (not passing through P) tangent to C_{n-1}, C_{n-2}, and L, $I(C_n)$ must be a circle, not passing through P, tangent to the lines $I(C_{n-1})$ and $I(C_{n-2})$ as well as to the circle $I(L)$. This fully determines the situation, which is shown in the figure below.

Note that any circles tangent to both $I(C_{n-1})$ and $I(C_{n-2})$ must fit in between the lines and have diameter given by the line separation, $(2r_{n-1})^{-1} + (2r_{n-2})^{-1}$. Let $\ell = \frac{1}{2}(2r_{n-1})^{-1} + \frac{1}{2}(2r_{n-2})^{-1}$ be the associated radius.

Denote by B the line through the centers of $I(L)$ and $I(C_n)$; B is the line halfway between $I(C_{n-1})$ and $I(C_{n-2})$. Note that P will be at least as close to $I(C_{n-1})$ as to $I(C_{n-2})$, since we expect $r_{n-1} \leq r_{n-2}$ by geometric intuition; this can be proven by induction along with the following calculations. Now P has distance $d_1 = (2r_{n-1})^{-1}$ from $I(C_{n-1})$, so the distance from P to line B is $\ell - d_1$. Let the closest point on B to P be D, that is, let D be the intersection of B with the perpendicular to B through P. Using the Pythagorean theorem, we can see that the distance from the center O of $I(L)$ to D is $\sqrt{\ell^2 - (\ell - d_1)^2} = \sqrt{2\ell d_1 - d_1^2}$. Now the distance between O and the center O' of $I(C_n)$ is 2ℓ, and this distance is along B, so the distance from D to O' is $\sqrt{2\ell d_1 - d_1^2} + 2\ell$. Finally, using the Pythagorean theorem again, we can see that the distance from P to O' is

$$|PO'| = \sqrt{(\ell - d_1)^2 + \left(\sqrt{2\ell d_1 - d_1^2} + 2\ell\right)^2} = \sqrt{5\ell^2 + 4\ell\sqrt{2\ell d_1 - d_1^2}}.$$

Since O' is the center of $I(C_n)$, a circle of radius ℓ, the least distance from P to $I(C_n)$ is $|PO'| - \ell$ and the maximum distance is $|PO'| + \ell$. If we perform another inversion about P, we map $I(C_n)$ back to C_n. The inversions of the previous minimum and maximum distance from P to $I(C_n)$ will give, respectively, the maximum and minimum distances from P to C_n; the difference between these is the diameter of C_n.

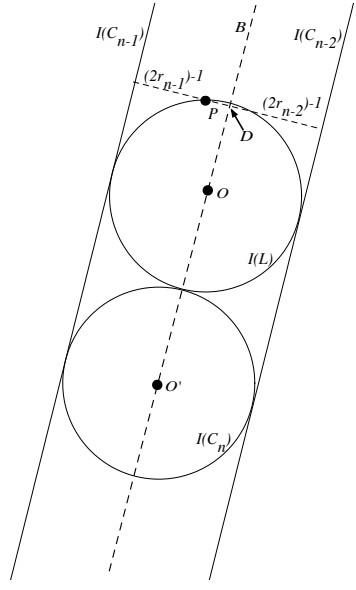

Figure 5.3: Inverted Fibonacci Circles.

Thus, we see that

$$\begin{aligned} 2r_n &= \frac{1}{|PO'| - \ell} - \frac{1}{|PO'| + \ell} \\ &= \frac{(|PO'| + \ell) - (|PO'| - \ell)}{|PO'|^2 - \ell^2} \\ &= \frac{2\ell}{4\ell^2 + 4\ell\sqrt{2\ell d_1 - d_1^2}}. \end{aligned} \tag{5.5}$$

After cancelling a 2ℓ in (5.5) and making the substitutions $d_1 = (2r_{n-1})^{-1}$ and $\ell = \frac{1}{2}(2r_{n-1})^{-1} + \frac{1}{2}(2r_{n-2})^{-1}$, we obtain

$$\begin{aligned} 2r_n &= \frac{1/2}{\frac{1}{2}(2r_{n-1})^{-1} + \frac{1}{2}(2r_{n-2})^{-1} + \sqrt{((2r_{n-1})^{-1} + (2r_{n-2})^{-1})(2r_{n-1})^{-1} - (2r_{n-1})^{-2}}} \\ 4r_n &= \frac{r_{n-1}r_{n-2}}{r_{n-2}/4 + r_{n-1}/4 + \sqrt{r_{n-1}r_{n-2}/4}} \\ r_n &= \frac{r_{n-1}r_{n-2}}{r_{n-1} + 2\sqrt{r_{n-1}r_{n-2}} + r_{n-2}} \\ r_n &= \frac{r_{n-1}r_{n-2}}{(\sqrt{r_{n-1}} + \sqrt{r_{n-2}})^2}. \end{aligned} \tag{5.6}$$

Taking an inverse and a square root of both sides of (5.6), we see that

$$\frac{1}{\sqrt{r_n}} = \frac{1}{\sqrt{r_{n-1}}} + \frac{1}{\sqrt{r_{n-2}}}, \tag{5.7}$$

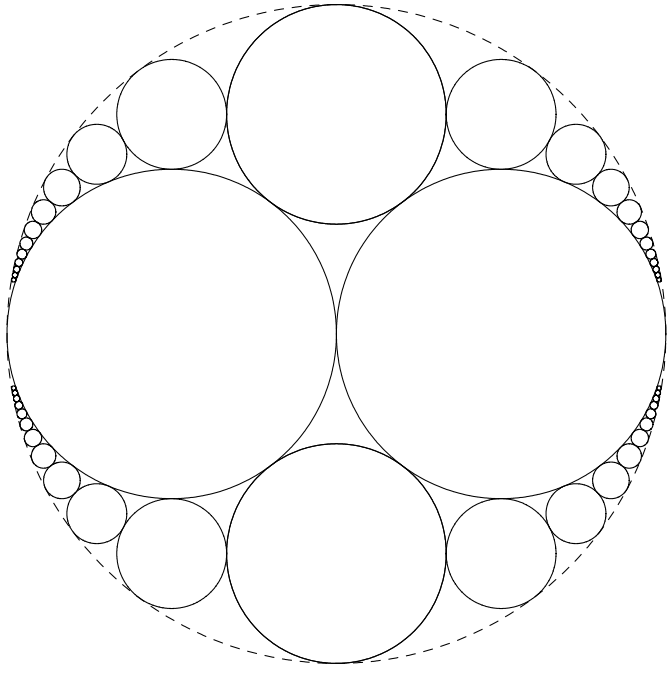

Figure 5.4: Multiple Pappus Chains.

which is the relation we were seeking. We can now complete the problem: let F_n be the nth Fibonacci number (indexed so that $F_1 = F_2 = 1$) and notice that $1/\sqrt{r_n}$ satisfies the Fibonacci relation. Thus, since $r_i = 1/F_i^2$ for $i = 1, 2$, this pattern holds in general; the radius of the nth circle is $1/F_n^2$. □

For the sake of honesty, we should observe that this relation may be derived more quickly by using analytic geometry techniques; inversion is not required. In fact, the relation (5.7) can be seen in a Japanese *sangaku* from 1824 [RF].

5.5 Conclusion

This paper has presented two examples which provide a taste for the variety of approaches in solving circle tangency problems. In particular, the author has found inversions of the plane to be extremely powerful. We close with a challenge: The reader is invited to find the total area of all the solid circles in Figure 5.4, below.

In Figure 5.4, the bounding, dotted circle has radius 2 and the two largest solid circles have radius 1. This problem is related to the so-called "Ancient Theorem" of Pappus, which has been studied in depth by Jakob Steiner [Coo]. Related constructions have also been found in Japanese *sangakus* [RF]. For more information and examples, we refer the reader to Martin Gardner's survey article [Ga].

5.6 Acknowledgment

The author wishes to dedicate this article to James Albrecht, in thanks both for posing the Fibonacci circles problem and in general for his camaraderie over the years.

References

[Coo] Julian Lowell Coolidge: *A Treatise on the Circle and the Sphere*. New York: Chelsea, 1971.

[Cox] Harold Scott MacDonald Coxeter: The problem of Apollonius, *Amer. Math. Monthly* **75** (1968), 5–15.

[Fo] L. R. Ford: Fractions, *Amer. Math. Monthly* **45** (1938), 586–601.

[Ga] Martin Gardner: Tangent circles, pages 149–166 in *Fractal Music, Hypercards and more*. New York: W.H. Freeman and Co., 1992.

[RF] T. Rothman and H. Fukagawa: Japanese temple geometry, *Sci. Amer.* **278** (1998), 85–91.

[SM] Robert T. Smith and Roland B. Minton: *Calculus: Concepts and Connections*. McGraw-Hill, 2004.

FACULTY FEATURE ARTICLE

6

Solving Large Classes of Nonlinear Systems of PDEs by the Method of Order Completion

Elemér Elad Rosinger[†]
University of Pretoria
Pretoria, 0002 South Africa
eerosinger@hotmail.com

"... provided also if need be that the notion of a solution shall be suitably extended."

—Hilbert's 20th Problem

6.1 Preliminaries

One of the sharpest divides in the history of technological and mathematical development occurred with Newton's development of calculus. Prior to this development, we were not able to understand, let alone model rigorously, the motion of even a single massive particle unless it was moving along a straight line or a circle and was doing so with constant velocity. Even Galileo's discoveries about gravitation and Kepler's laws of planetary motion were merely *empirical*. In other words, these so-called laws were not based on any scientific principles or corresponding rigorous theories and were instead just formulae fitted to observed data.

Newton's three laws of motion and the methods of calculus he invented enabled the practice of science as we know it today. In particular, Newton's methods allowed the fundamental laws of nature to be formulated as **differential equations**. In fact, much of modern science would be impossible to formulate, let alone apply technologically, without the use of ordinary and partial differential equations, denoted respectively by ODEs and PDEs.

A simple example serves to illustrate the immense leap brought about by Newton's calculus. **Newton's Second Law** states, in modern terms, that the motion of a massive particle along a straight line satisfies the property "mass times acceleration is equal to force," $ma = F$. In terms of the position $x(t)$ of a particle, acceleration is the second derivative: $a(t) = \ddot{x}(t)$. Thus the Second Law takes the form of the second-order ordinary differential equation $m\ddot{x}(t) = F(t)$ in the

[†]Prof. Elemér Elad Rosinger earned his doctorate in mathematics in 1972 at the University of Bucharest, Romania, in the subject of functional analysis, under Prof. G. Marinescu. Starting in 1960, he went through a variety of academic and industrial jobs, until 1973, when he moved away from Eastern Europe. During 1974-1979 he was with the department of applied mathematics, and subsequently, computer science at Haifa Technion, Israel. In the years 1980-1983, he enjoyed interdisciplinary research at CSIR (Council for Scientific and Industrial Research), Pretoria, South Africa. Since 1983, he has been with the department of mathematics and applied mathematics at the University of Pretoria. Over the years, he has visited several dozen universities and research institutes on five continents. His research interests range over a number of fields, among them, nonlinear mathematics, optimization, relativity, quantum mechanics and the foundations of physics.

He has published nine research monographs and several dozen papers. Among the more notable publications he has produced is the first complete presentation of Hilbert's Fifth Problem. He has discovered powerful solution methods for general nonlinear systems of partial differential equations. Among his wider interests are reading and writing books, both fiction and non-fiction, on topics including philosophy and metaphysics. When young, he was involved in classical music, swimming and martial arts. Now, he is focusing on daily walking.

position $x(t)$. A significant implication of this law is that one requires precisely *two* independent initial conditions to determine a unique solution in a particular physical situation. One may, for example, give the initial position $x(0)$ and the initial velocity $\dot{x}(0)$.

While this may seem obvious, Newton's innovation was in fact quite revolutionary. For approximately two millennia prior to Newton, the prevailing view of such motion, stated by Aristotle, held velocity, rather than acceleration, to be proportional to the force applied. Thus, according to Aristotle, the *first*-order differential equation $\dot{x}(t) = cF(t)$, where $c > 0$ is some constant, would describe motion; he engaged in no experimentation to test this notion. Using Newton's calculus, the trouble with such a first-order equation is so obvious that one need not conduct any experimentation to perceive it. Indeed, as the equation is first-order in the position $x(t)$, it only allows *one single* initial condition for the unique determination of its solution. But this *contradicts* the empirically known fact that one can give two objects with the same initial position two different initial velocities, resulting in two different trajectories.

Considerable mathematical effort has been expended in solving such equations, especially PDEs. The modern era of PDE theory started in the early 20th century, when methods of functional analysis were introduced. This trend became very strong starting in the 1930s, when a large variety of Sobolev spaces proved to be particularly convenient for finding solutions to PDEs. Finally, in the late 1940s, with the introduction of Schwartz distributions, the functional analytic methods became ubiquitous.

Sobolev spaces are complete normed spaces, that is, **Banach spaces**; they are more sophisticated than the **Lebesgue spaces** L^p, with $1 \leq p \leq \infty$, since their norms involve not only the generalized functions which are their elements but also various derivatives of these generalized functions. A motivation for such sophisticated norms is that, under the usual norms of the L^p spaces, the derivative is not a bounded operator and thus is not continuous. Under certain conditions, the generalized functions in Sobolev spaces turn out to be the usual smooth functions.

As for the more general spaces $\mathcal{D}'$ or $\mathcal{S}'$ of **Schwartz distributions**, these are no longer normed spaces, so their topologies are far more complicated. Indeed, those topologies are locally convex, thus considerably more general than those of normed spaces. These more general spaces of distributions prove to have advantages beyond those of Sobolev spaces when solving certain classes of PDEs.

Let us now be more explicit about what it means to solve a differential equation, from both a mathematical and physical perspective. We take two simple yet nontrivial examples for illustrative purposes.

First, let us consider the motion of a particle under gravitation. We let the variable x denote position and denote by x_0 the position of the particle at time $t = 0$. Then, as gravitational force is constant, Newton's Second Law gives the second-order ordinary differential equation $\ddot{x}(t) = 1$ after an appropriate normalization of units. A general solution of this equation exists for all $t \in \mathbb{R}$, given by $x(t) = t^2/2 + v_0 t + x_0$, where v_0 is the initial velocity of the particle. Now two facts are important to note here: First, such a general solution exists for all $t \in \mathbb{R}$. Second, the general solution describes all possible free falls, rather than merely the free fall of one particular particle. To find this description of the particle's motion, we specify its initial position x_0 together with its initial velocity v_0. In general, we are interested in **existence** and **uniqueness** of solutions, the latter under specific additional conditions which are required by, for example, physical reality.

As a second example, let us consider one of the most basic PDEs of fluid dynamics, namely the **nonlinear shock-wave equation** $U_t(t,x) + U(t,x)U_x(t,x) = 0$, where $t \geq 0, x \in \mathbb{R}$. Under certain conditions, this equation can be seen as describing the motion of a fluid within an infinitely long tube parallel to the x-axis. Here, $U(t,x)$ represents the velocity at time t of a particle of fluid which is at the point of coordinate x. It is well known that the general solution of this equation exists; the issue is how to determine a *unique* solution which corresponds to a specific physical situation.

The mathematical answer, and one which makes physical sense, is that one should give an initial condition which describes the velocity of the fluid at time $t = 0$ and along the whole x-axis. Namely, one has to give $U(0,x) = u(x)$, for $x \in \mathbb{R}$; in this case, the initial condition is a function defined on all of $\mathbb{R}$.

There is also a *third* problem, especially concerning the solution of PDEs, namely, determining the **regularity** of solutions. From this point of view, the nonlinear shock-wave equation can already give a good example, even if it is one of the simplest nontrivial nonlinear equations of major physical interest. Namely, we note that this equation is of *first*-order, since only first-order partial derivatives of the unknown function U appear. Therefore, we would expect that any well-behaved solution is given by a function U of t and x such that both the partial derivatives U_t and U_x which appear in that equation exist. Indeed, if we are given a function U of t and x which is not differentiable with respect to both variables, then we simply cannot verify in a straightforward manner whether U is a solution of that equation. So, by a **classical solution** to a PDE, we mean a function for which all partial derivatives which appear in the PDE exist; this condition is therefore the natural definition of regularity.

Here, however, things get rather complicated. Indeed, even in the case of the above nonlinear shock-wave equation, many physically relevant solutions are not at all classical. Such solutions are called **shock waves**, and their existence and physical relevance is precisely the reason for the name of the equation. In fact, such shock wave solutions U not only lack the partial derivatives U_t and U_x, but even fail to be continuous. Nonetheless, such solutions are physically realistic; for example, they can model the effects of a sonic boom.

It follows that even though we would like solutions to be regular in the classical sense, important practical considerations oblige us to deal with solution which are less regular than the classical ones. Such non-regular solutions are then called **generalized solutions**. Of course, developments show that such generalized solutions do nevertheless satisfy the respective PDEs in certain suitable senses. Consequently, the problem of regularity of solutions of PDEs means, in practice, to find solutions which are in some sense "generalized as little as possible" and thus are as near to classical solutions as possible. The case of the shock waves shows that this is not always a trivial issue.

The history of solving ordinary and partial differential equations is impressively rich and complex. Its complexity should in no way be surprising, since most of the fundamental laws of nature which such equations model are not as simple as that of a free-falling particle. The history is also remarkable, given its sometime paradoxical ways of progressing.

For instance, in spite of the fact that solving PDEs is significantly harder than solving ODEs, the first general existence, uniqueness and regularity result for solutions was that of Cauchy-Kovalevskaya for arbitrary nonlinear systems of analytic PDEs, obtained in the early 1870s. Furthermore, there are two instructive facts with respect to this result. First, the "hardest" mathematics used in the proof of that theorem is the summation of a convergent geometric series. Thus in particular, its proof used no topology, let alone functional analysis of any kind. Second, the subsequent century-long development of topology and functional analysis was not able to improve even slightly upon the original result of the Cauchy-Kovalevskaya theorem when one considers this theorem in its own terms of nonlinear generality or upon the strength of its existence, uniqueness and regularity results. In this regard, the first time an extension in of the Cauchy-Kovalevskaya theorem (its own terms) was obtained was in [Ro7].[1] Once again (and surprisingly), functional analytic methods were not used.

As it happens, it took approximately two decades following the Cauchy-Kovalevskaya theorem before a correspondingly general existence and uniqueness result for ODEs was obtained by Picard and Lindelöf, who used a sophisticated fixed point argument, typical of methods in functional analysis. The relative strength of the Picard-Lindelöf result is that it is valid not only for analytic ODEs, but also for those ODEs which are far less smooth, for instance, ODEs that are continuous with some mild Lipschitz-type conditions.

As far as ODEs are concerned, methods of solving such equations are well-established, and the main remaining concerns are of a numerical nature related to improvements in the approximation of such solutions. With respect to the solution of PDEs, since the introduction of Sobolev spaces, and in general, of the Schwartz distributions, functional analytic methods have attained a near-monopoly, with hardly any other significant methods developed until the late 1970s. Furthermore, it became common to claim that it is simply not possible mathematically to develop a general

[1] See also [Ro8], [Ro12], where a global version of that theorem was presented.

existence, uniqueness and regularity theory for solving PDEs.

Instead, it is claimed, one must focus on specific types of such equations, each with its own highly specific solution method. Thus, the claim is that present day mathematics is in fact incapable of developing any relevant **type-independent** PDE theory with respect to the existence, uniqueness and regularity of solutions. Recent expressions of that strongly entrenched view can be seen in advanced textbooks of noted specialists in PDEs. For example Arnold's text [Ar], starts with the statement (italics added):

> In contrast to ordinary differential equations, there is *no unified theory* of partial differential equations. Some equations have their own theories, while others have no theory at all. The reason for this complexity is a more complicated geometry...

Similarly, Evans' text [Ev], starts his Examples on page 3 with the somewhat more cautious statement (italics added):

> There is no general theory known concerning the solvability of all partial differential equations. Such a theory is *extremely unlikely* to exist, given the rich variety of physical, geometric, and probabilistic phenomena which can be modelled by PDE. Instead, research focuses on various particular partial differential equations ...

The historical facts, however, show the relevance of general, type-independent results concerning PDEs. Indeed, in the context of arbitrary **analytic nonlinear** systems of PDEs, such a general, type-independent result was obtained on the existence, uniqueness, and analytic regularity of solutions back in the 1870s with the classical Cauchy-Kovalevskaya theorem. In the context of **linear constant coefficient** PDEs, a general type-independent existence result was already obtained in the early 1950s by Malgrange, and independently by Ehrenpreis, concerning the so-called elementary solutions of such equations, related to the well-known Green functions.

The severe limitations of the functional analytic methods in solving even linear PDEs came most unexpectedly and shockingly to the fore fifty years ago, with the celebrated 1957 Hans Lewy impossibility result [Le], concerning the nonexistence of solutions of PDEs. Indeed, Lewy showed that the rather simple linear first-order PDE in three independent variables and with first degree polynomial coefficients

$$(D_x + iD_y - 2(x+y)D_z)U(x,y,z) = f(x,y,z), \qquad (x,y,z) \in \mathbb{R}^3 \tag{6.1}$$

does not have distribution solutions in any neighborhood of any point in $\mathbb{R}^3$, for a large class of smooth right-hand terms f. In 1967, Shapiro gave a similar example of a smooth linear PDE which does not have solutions in Sato's **hyperfunctions**.

Recently, however, type-independent existence, uniqueness, and regularity results on solutions of large classes of nonlinear systems of PDEs, with possibly associated initial and/or boundary value problems have been introduced [OR]. The method of solution, a first in the literature, is based on the **order completion** (see Appendix) of suitable spaces of usual functions on the Euclidean domains of definition of the respective PDEs. As a general and hence type-independent regularity result, the solutions obtained can be assimilated with **Hausdorff continuous** functions (see Appendix) on the domains of the PDEs (see [An], [Ro1, Ro9, Ro11, Ro17], [Wa1, Wa2, Wa3, Wa4]).

Thus, one can do away with the use of various generalized functions, such as Schwartz distributions or elements of Sobolev spaces.[2] However, the latest methods of improved regularity being developed for the general order completion method indicate the possibility of gaining more insight into this problem [Wa1, Wa2, Wa3, Wa4]. An important fact to note is that the order completion solution method does not involve functional analysis; thus, it does not make use of various Sobolev or other spaces of distributions or generalized functions which usually provide solutions to PDEs.

[2] Among others, the usual 3-dimensional Navier-Stokes equations are included as a particular case of the nonlinear systems of PDEs which can be solved by the order completion method; the resulting solutions are Hausdorff continuous (a somewhat weak regularity condition).

Instead, the solutions obtained are no longer generalized functions and can be assimilated with Hausdorff continuous functions.

The power of the order completion method is shown in three facts. First, this method is the first in the literature to overcome the celebrated 1957 Hans Lewy impossibility . In fact, it overcomes Lewy's result in the case of very general nonlinear PDEs, far beyond the simple linear PDE in (6.1). Second, the order completion solution method allows a particularly convenient treatment of initial and/or boundary value problems associated with PDEs, which, as is well-known, is an advantage over functional analytic methods [OR, Chapter 8]. Third, and perhaps most importantly, the concept of order is more basic than that of algebraic structure. Indeed, the dichotomy between linear and nonlinear PDEs, which singles out the nonlinear ones as incomparably harder to solve, manifests itself on the algebraic level: more precisely, in terms of vector spaces. Therefore, that unfortunate dichotomy between linear and nonlinear PDEs is simply unavoidable with functional analytic methods. On the other hand, the order completion method does not distinguish between linear and nonlinear PDEs, solving both types of differential equation with equal ease (see [OR], [Ro1, Ro9, Ro11, Ro17], [Wa1, Wa2, Wa3, Wa4]).

The order completion method in [OR] brought a considerable improvement with respect to the extent that general type-independent existence, uniqueness and regularity results concerning solutions of large classes of nonlinear systems of PDEs can be obtained. Thus, it appears that the limitations claimed on PDE theory are in fact only limitations on the functional analytic methods used.

6.2 Main Ideas of the Order Completion Solution Method

The solution method is divided into two parts. The proof of the **existence** and **uniqueness** of solutions follows the method of order completion introduced and first developed in [OR]. The proof of the **regularity** of solutions is a consequence of recent results regarding the structure of the **Dedekind order completion** of spaces of continuous functions $\mathcal{C}(X)$, where X is a topological space with some weak conditions on it [An]. The respective regularity results have further been developed and improved in [Ro1, Ro9, Ro11, Ro17], [Wa1, Wa2, Wa3, Wa4].

For simplicity of presentation, we shall consider single nonlinear PDEs.[3] Let us therefore consider nonlinear PDEs of the general form

$$F(x, U(x), \ldots, D_x^p U(x), \ldots) = f(x), \qquad x \in \Omega \subseteq \mathbb{R}^n \tag{6.2}$$

with $p \in \mathbb{N}^n$, $|p| \leq m$. Here, the domain Ω is an open, not necessarily bounded subset of $\mathbb{R}^n$, while the orders $m \in \mathbb{N}$ of the PDEs are fixed but otherwise arbitrary, and solutions are functions $U : \Omega \to \mathbb{R}$.

The unprecedented generality of these nonlinear PDEs comes, above all, from the class of functions F which define the left-hand terms, and which are only assumed to be **jointly continuous** in all of their arguments. The right hand terms f are also required to be continuous.[4]

Regardless of the above generality of the nonlinear systems of PDEs considered, one can find for them solutions U defined on the whole of the respective domains Ω. These solutions U have the **type-independent**, or **universal regularity**, property that they can be assimilated with Hausdorff continuous functions.

It follows in this way that, when solving systems of nonlinear PDEs of the generality of those in (6.2), one can dispense with the various customary spaces of distributions, hyperfunctions, generalized functions, Sobolev spaces, and so on. Instead, one can stay within the realms of "usual

[3] The extension to systems of such nonlinear PDEs and associated initial and/or boundary value problems can, rather surprisingly, be done easily, this being one of the major advantages of the order completion method (see [OR]).

[4] However, it turns out that in the most general case, both F and f can have certain discontinuities as well (see [OR]).

functions," that is, **interval-valued functions** (see Appendix).[5]

Let us now associate with each nonlinear PDE in (6.2) the corresponding nonlinear partial differential operator defined by its left hand side, namely

$$T(x, D)U(x) = F(x, U(x), \dots, D_x^p U(x), \dots), \qquad x \in \Omega. \tag{6.3}$$

The fact that $T(x, D)$ is an **operator** simply means that the nonlinear PDE in (6.2) can be written in the simple form

$$T(x, D)U(x) = f(x), \qquad x \in \Omega. \tag{6.4}$$

Two facts about the nonlinear PDEs in (6.2) and the corresponding nonlinear partial differential operators $T(x, D)$ in (6.3) are important and immediate:

- The operators $T(x, D)$ can *naturally* be seen as functions acting in the **classical context**, namely, between classical spaces of functions

$$T(x, D) : \mathcal{C}^m(\Omega) \to \mathcal{C}^0(\Omega). \tag{6.5}$$

Unfortunately on the other hand:

- The mappings in this natural classical context (6.5) are typically not surjective even in the case of linear $T(x, D)$, and they are even less so in the general nonlinear case of (6.2), (6.4).

In other words, linear or nonlinear PDEs in (6.2) typically cannot be expected to have classical solutions $U \in \mathcal{C}^m(\Omega)$, for arbitrary continuous right-hand terms $f \in \mathcal{C}^0(\Omega)$, as illustrated by a variety of well-known examples, some of them rather simple ones (see [OR, Ch. 6]).

Furthermore, it can often happen that non-classical solutions do have a major applicative interest and thus have to be sought out beyond the confines of the classical framework in (6.5). One of the simplest such examples comes from the aforementioned shock wave solutions of the nonlinear shock wave equation. In fact, non-classical solutions can be critically important, even in the case of linear PDEs.[6] Thus we are led to the necessity of considering generalized solutions U to PDEs like those in (6.2), that is, solutions $U \notin \mathcal{C}^m(\Omega)$, which therefore are no longer classical. This means that the natural classical mappings (6.5) must in certain suitable ways be extended to **commutative diagrams**:

$$\begin{array}{ccc} \mathcal{C}^m(\Omega) & \xrightarrow{T(x,D)} & \mathcal{C}^0(\Omega) \\ \downarrow \subseteq & & \subseteq \downarrow \\ X & \xrightarrow{\widetilde{T}} & Y \end{array} . \tag{6.6}$$

The generalized solutions are now found as

$$U \in X \setminus \mathcal{C}^m(\Omega), \tag{6.7}$$

instead of the classical solutions $U \in \mathcal{C}^m(\Omega)$, which may easily fail to exist. A further important point is that one expects to reestablish certain kinds of surjectivity properties typically missing in (6.5); for example,

$$\mathcal{C}^0(\Omega) \subseteq \widetilde{T}(X). \tag{6.8}$$

[5] Furthermore, when proving the existence and the mentioned type of regularity of such solutions, one can dispense with methods of functional analysis. However, functional analytic methods can possibly be used in order to obtain further regularity or other desirable properties of such solutions. Therefore, the order completion method does not aim to abolish functional analytic methods in solving PDEs, but rather to improve significantly on the well-known—yet so often disregarded—severe limitations of such methods.

[6] Such, for example, as those whose solutions are given by **Green functions**.

Here, it is important to note the following two facts. First, the extended spaces X and Y need not be minimal. Indeed, one is interested in solving not only one particular PDE, or one single system of PDEs. On the other hand, as the history of PDE theory has clearly shown, we cannot expect to find some sort of universally valid unique extensions X or Y. Moreover, such extensions may often depend on the PDEs solved, although different PDEs may still be solvable in the same extensions. Second, and following from the above, we should not always ask the surjectivity condition in its strongest possible form, $\widetilde{T}(X) = Y$. Instead, depending on the particulars of the situation, it may be sufficient to ask only that $\widetilde{T}(X)$ is a large enough subset of Y, such as that specified in (6.8) above.

Before going further, let us recall that extensions of mappings through commutative diagrams similar to (6.6) have been associated with solving equations—even if not explicitly—ever since ancient times (see [OR, chap. 12]). For example, it is well-known that for all $x \in \mathbb{R}$, $x^2 \neq -1$. That is, $x^2 + 1 = 0$ has no solution in $\mathbb{R}$. However, as we all know, that equation does have a solution in $\mathbb{C}$. This fact can be formulated in the following extension of a mapping through a commutative diagram. Namely, let us define the mapping $T : \mathbb{R} \ni x \mapsto T(x) = x^2 \in \mathbb{R}$. Then we have the commutative diagram

$$\begin{array}{ccc} \mathbb{R} & \xrightarrow{T} & \mathbb{R} \\ \big\downarrow{\scriptstyle \subseteq} & & {\scriptstyle \subseteq}\big\downarrow \\ \mathbb{C} & \xrightarrow{\widetilde{T}} & \mathbb{C} \end{array} . \tag{6.9}$$

Here, of course, T is *not* surjective, since $-1 \in \mathbb{R} \setminus T(\mathbb{R})$. On the other hand, $\widetilde{T} : \mathbb{C} \ni x \mapsto \widetilde{T}(x) = x^2 \in \mathbb{R}$ has the property that $-1 \in \widetilde{T}(\mathbb{C})$.

6.3 Constructing the Order Completion

Since we solve PDEs through order completion, let us see how close we can come to satisfying the equality in (6.2), in the sense of order. For that purpose, it is useful to consider, for each $x \in \Omega$, the set of real numbers

$$\mathbb{R}_x = \{F(x, \xi_0, \ldots, \xi_p, \ldots) \mid \xi_p \in \mathbb{R}, \text{ for } p \in \mathbb{N}^n, |p| \leq m\}. \tag{6.10}$$

Clearly, for fixed $x \in \Omega$, $\mathbb{R}_x$ is the range in $\mathbb{R}$ of $F(x, \ldots)$, and since F is jointly continuous in all its arguments, it follows that $\mathbb{R}_x$ is a nonempty interval which is bounded, half-bounded, or is the whole of $\mathbb{R}$. This latter case, which can happen often with nonlinear PDEs in (6.2), will be easier to deal with, as we will see in (6.12) below.

Clearly, in the case of non-degenerate linear PDEs in (6.2), the latter case is ubiquitous. Now given $x \in \Omega$, it is obvious that a necessary condition for the existence of a classical smooth solution $U \in \mathcal{C}^m$ of (6.2) in a neighborhood of x is

$$f(x) \in \mathbb{R}_x. \tag{6.11}$$

Consequently, for the time being, we shall make the assumption that the right hand term functions f in the nonlinear PDEs in (6.2) satisfy the somewhat stronger version of condition (6.11) given by

$$f(x) \in \text{interior}(\mathbb{R}_x), \qquad \text{for } x \in \Omega. \tag{6.12}$$

Clearly, whenever we have

$$\mathbb{R}_x = \mathbb{R}, \qquad \text{for } x \in \Omega, \tag{6.13}$$

then (6.12) is satisfied. And as mentioned, this is the case with all nontrivial linear PDEs, as well as with most of the nonlinear PDEs of practical interest.

We now formulate the basic and rather simple **local approximation** result on how nearly we can satisfy the equality in (6.2) and (6.4). A remarkable fact is that the proof of this local approximation result, as well as of its global version in Proposition 2 in the sequel, is surprisingly elementary.

Proposition 1 ([OR], Lemma 2.2)**.** *Given $f \in \mathcal{C}^0(\Omega)$, then for all $x_0 \in \Omega \subset \mathbb{R}^n$, $\epsilon > 0$, there exists $\delta > 0$ and a polynomial P in n variables with real coefficients such that in a δ-ball around x_0, we have*

$$f(x) - \epsilon \leq T(x, D)P(x) \leq f(x). \tag{6.14}$$

In view of the several successive quantifiers in the above approximation result, let us briefly elucidate it in a somewhat less formal manner. Our first aim is, of course, to prove the existence of solutions U of the nonlinear PDE in (6.4). The order completion method obtains such existence results in two steps. First, it shows that the nonlinear PDE in (6.4) can be satisfied *approximately* as nearly as we want. Second, it shows that—in their totality as a set—such approximate solutions do in fact define an exact solution, provided that we build a convenient order completion of both the domain and range of the nonlinear partial differential operator $T(x, D)$ in (6.3) and do so as in the commutative diagram (6.6). This is similar to the order completion of the rationals used to obtain the reals and in fact uses a method analogous to Dedekind cuts. As it happens, however, with the nonlinear PDE in (6.4), we are not looking for one single number, but for a whole function $U : \Omega \to \mathbb{R}$. Furthermore, it is much easier first to approximate the solution of that nonlinear PDE in (6.4) only *locally*, that is, in a suitable neighborhood of any given point $x_0 \in \Omega$. This approximation is precisely what the above proposition accomplishes. Namely, for every given $x_0 \in \Omega$ and $\epsilon > 0$, it delivers such a simple (in fact, polynomial) function P, together with a neighborhood of x_0 described by a corresponding $\delta > 0$, with the **two-sided approximation property**

$$f(x) - \epsilon \leq T(x, D)P(x) \leq f(x), \qquad x \in \Omega, \|x - x_0\| \leq \delta.$$

Let us briefly give a proof of Proposition 1:

Proof of Proposition 1. Let any $x_0 \in \Omega$ be given. Then for suitable $\epsilon > 0$, (6.12) yields

$$\xi_p \in \mathbb{R}, \qquad \text{for } p \in \mathbb{N}^n, |p| \leq m$$

such that

$$F(x_0, \xi_0, \ldots, \xi_p, \ldots) = f(x_0) - \frac{\epsilon}{2}$$

Therefore, there exists a polynomial P in the variable $x \in \mathbb{R}^n$, such that

$$D^p P(x_0) = \xi_p, \qquad \text{for } p \in \mathbb{N}^n, |p| \leq m,$$

which means that

$$T(x_0, D)P(x_0) - f(x_0) = -\frac{\epsilon}{2}.$$

However, both F and f are assumed to be continuous, so the function $\Omega \ni x \mapsto T(x, D)P(x) - f(x) \in \mathbb{R}$ is continuous as well. Therefore (6.14) follows immediately. □

And now, the **global approximation** version of the inequality property in (6.14) is given by

Proposition 2 ([OR], Prop. 2.2)**.** *Suppose $f \in \mathcal{C}^0(\Omega)$. Then for all $\epsilon > 0$, there exists $\Gamma_\epsilon \subset \Omega$ closed and nowhere dense and $U_\epsilon \in \mathcal{C}^m(\Omega \setminus \Gamma_\epsilon)$ such that*

$$f - \epsilon \leq T(x, D)U_\epsilon \leq f \tag{6.15}$$

on $\Omega \setminus \Gamma_\epsilon$.

Remark. It is easy to see that the inequalities in (6.14) and (6.15) can be replaced with

$$f(x) \leq T(x, D)P(x) \leq f(x) + \epsilon, \tag{6.16}$$

$$f \leq T(x, D)U_\epsilon \leq f + \epsilon, \tag{6.17}$$

as the proofs of (6.16) and (6.17) follow after the corresponding obvious minor changes in the proofs of the above two propositions.

We now proceed to the order completion, based on MacNeille's construction, using **Dedekind cuts** (see [OR, Ma, Lu]); such cuts require the above sharp inequalities. Let us briefly recall here Dedekind's original construction of $\mathbb{R}$ from $\mathbb{Q}$. While this construction is simpler than that of MacNeille, as $\mathbb{Q}$ is totally ordered, it is largely analogous.

Dedekind calls a **cut** in $\mathbb{Q}$ any partition of $\mathbb{Q}$ into two subsets A and B which satisfy $x < y$ for all $x \in A, y \in B$. For instance, the cut which defines $\sqrt{2} \in \mathbb{R} \setminus \mathbb{Q}$ is given by $A = \{x \in \mathbb{Q} \mid x^2 < 2\}$ and $B = \{y \in \mathbb{Q} \mid y^2 > 2\}$. Thus, if we want effectively to construct A, for example, then one way to obtain it is by the union $A = \bigcup_{\epsilon>0}\{x \in \mathbb{Q} \mid 2 - \epsilon \leq x^2 \leq 2\}$ which also gives an approximation process (from below) for $\sqrt{2}$. The approximation results given here model this process, albeit with polynomials in place of rationals.

Note that in Proposition 2, as well as in its version corresponding to the above inequality (6.17), we can have in addition the property

$$\text{mes}(\Gamma_\epsilon) = 0 \tag{6.18}$$

where mes denotes the usual Lebesgue measure.[7]

As seen from the proof of Proposition 2 (see [OR, pp. 18-20]), the functions U_ϵ can in fact be chosen as piecewise polynomials in $x \in \mathbb{R}^n$.

The considerable power of the order completion method in solving very general classes of nonlinear systems of PDEs comes from the fact that in the above order approximation results (6.15) and (6.17), one does not need more than the continuity of the functions F and f which define the nonlinear PDEs (6.2). Due to the inevitable presence of the closed, nowhere dense subsets of singularities Γ_ϵ, one can in fact allow even certain *discontinuities* in these functions F and f (see [OR]).

And now, the construction of commutative diagrams (6.6) follows easily, [OR], [Ro1, Ro9, Ro11, Ro17]. Indeed, the order approximations in (6.15) or (6.17) lead to the construction of the spaces X and Y as the **Dedekind order completion** (see Appendix) of spaces of piece-wise smooth functions corresponding in a natural manner to $\mathcal{C}^m(\Omega)$ or $\mathcal{C}^0(\Omega)$. Then (this is nontrivial), the mappings $\widetilde{T}$ turn out to be **order isomorphic embeddings**.

6.4 General Existence Result

Once we reformulated the problem of solving PDEs in terms of the commutative diagrams (6.6), all the subsequent results concerning existence, uniqueness and regularity of solutions are obtained in terms of such diagrams.

One of the typical **main existence results** concerning the solutions of the nonlinear PDEs in (6.2) is presented in the following theorem (see [OR, pp. 38-64] for a proof):

Theorem 3. *In the commutative diagram* (6.6)*, we have*

$$\widetilde{T}(X) = Y. \tag{6.19}$$

That is, $\widetilde{T}$ is surjective.

[7] It should be noted that the presence of the closed, nowhere-dense singularity sets Γ_ϵ in the global inequalities (6.15) and (6.17) proves not to be a hindrance. In fact, the presence of such closed, nowhere dense singularity sets is rather deeply-rooted, as it is connected with the flabbiness of related sheaves of functions, or the global version of the classical Cauchy-Kovalevskaya theorem on analytic nonlinear PDEs (see [OR, chap. 7] and the literature cited there).

This means that, given any nonlinear PDEs in (6.2), for every right hand term $f \in Y$, there exists a solution $U \in X$, satisfying the relation $\widetilde{T}(U) = f$.

However, as mentioned following (6.8) and (6.9), it is not always convenient to expect, let alone require, that one has equality in (6.19). Instead, what happens often, and turns out to be satisfactory in applications, is a weaker form of (6.5), namely, one in which $\widetilde{T}(X)$ can be proved to be large enough.

It is important to note that the spaces Y for which nonlinear PDEs are now solved by Theorem 3 include many highly discontinuous functions on Ω (see [OR, pp. 74-93]).

What is particularly interesting is that, in view of (6.19), a large variety of linear and nonlinear PDEs can be solved, in spite of the fact that the respective PDEs are known not to have solutions in distributions or in Sobolev spaces. Among such PDEs is the celebrated 1957 Hans Lewy impossibility example (6.1). In this regard, it was for the first time in [OR, chap. 6, 8] that this Hans Lewy example of a PDE not solvable in distributions or Sobolev spaces was nevertheless solved (through the method of order completion).

The correspondence between the solutions obtained in (6.19) and the usual classical solutions, (whenever the nonlinear PDEs in (6.2) may have classical solutions) follows easily from the way the commutative diagrams (6.6) are constructed. In other words, whenever the nonlinear PDEs in (6.2) happen to have classical solutions $U \in \mathcal{C}^m(\Omega)$, then they are also solutions in the sense of (6.19).

Recently, significant further improvements of the regularity of solutions were obtained through a refinement of the order completion method, [Wa1, Wa2, Wa3, Wa4]. The respective results indicate that the order completion method has considerable potential in attaining stronger regularity results than are currently known even without the use of functional analytic methods.

As far as the generality of the existence result of solutions, this was already attained to such an extent in [OR] that, at present, there appears to be no need for further extensions.

Finally, let us mention that the order completion method turns out to be significantly more powerful in solving large classes of nonlinear PDEs than the earlier introduced **nonlinear algebraic** method.[8]

Indeed, while the earlier nonlinear algebraic method can solve large classes of smooth linear or nonlinear PDEs, it falls short, even if not by much, in overcoming the Hans Lewy impossibility result (6.1). There are two main shortcomings of that algebraic method. First, the nonlinear PDEs which it can solve are significantly less general than those solved by the order completion method. Second, the solutions delivered by the algebraic method tend to have rather weak regularity properties, since they are given by generalized functions which are in spaces far larger than the Schwartz distributions or the Sobolev spaces. This is a sharp contrast with solutions delivered by the order completion method: solutions which are Hausdorff continuous functions.

Having said this, of course, it is important to note that the algebraic method in solving nonlinear PDEs is powerful enough to offer the first complete solution of Hilbert's Fifth Problem (see [Ro15]).

6.5 Appendix

6.5.1 Order Completion

A given poset $(X, \leq)$ is called **order complete** if and only if $\sup A, \inf A \in X$, for every $A \subseteq X$. If $\sup A \in X$ (respectively, $\inf A \in X$), only for every upper, (respectively, lower) bounded $A \subseteq X$, then $(X, \leq)$ is called **Dedekind order complete**. Clearly, $\mathbb{R}$ with its usual order is Dedekind order complete but not also order complete. On the other hand, the extended real line $\overline{\mathbb{R}} = [-\infty, \infty]$, as well as the closed intervals $[a, b] \subset \mathbb{R}$ are both Dedekind order complete and order complete.

[8] See [Ro1]–[Ro17], Zbl717*35001, MR92d:46098, MR89g:35001, Bull.AMS, Jan.1989, 96-101, and also subject 46F30 at `http://www.ams.org/msc/46Fxx.html`

Given two posets $(X, \leq)$ and $(Y, \leq)$, a mapping $\psi : X \to Y$ is called an **order isomorphic embedding** if and only if, for $x, x' \in X$, we have $x \leq x' \iff \psi(x) \leq \psi(x')$. If in addition ψ is also surjective, then it is called an **order isomorphism**.

The fundamental result with respect to order completion was obtained in 1937 by MacNeille: it states that for every poset $(X, \leq)$ which does not have a smallest or a largest element, there is an order complete poset $(\widetilde{X}, \leq)$ in which X is order-dense, that is, for every $\widetilde{x} \in \widetilde{X}$, there exists a subset $A \subseteq X$, such that $\widetilde{x} = \sup A$, [Ma] (see also [Lu] or [OR, Appendix]). Furthermore, $(\widetilde{X}, \leq)$ is *unique* up to **order isomorphism**.

A remarkable fact about MacNeille's result is that the order completion $\widetilde{X}$ is obtained in a manner which is a direct generalization of the construction of $\mathbb{R}$ from $\mathbb{Q}$ by Dedekind cuts. Thus MacNeille's method is called the **Dedekind order completion of** $(X, \leq)$, although it delivers an order complete poset $(\widetilde{X}, \leq)$, rather than a Dedekind order complete poset.

6.5.2 Hausdorff Continuous Functions

Let us denote by $\mathbb{A}$ the set of all functions $f : \mathbb{R} \ni x \mapsto [a, b]$, where $-\infty \leq a \leq b \leq \infty$. Thus such functions have **closed interval** values, and the respective intervals can be infinite at one, or at both ends. Usual or extended real valued functions $f : \mathbb{R} \to \overline{\mathbb{R}}$, where $\overline{\mathbb{R}} = [-\infty, \infty]$, can naturally be seen as particular cases of such interval valued functions, if we consider them as having values at $x \in \mathbb{R}$, as given by the intervals $[f(x), f(x)]$ reduced to single points.

Now to every function $f \in \mathbb{A}$ we associate two functions $If, Sf : \mathbb{R} \to \overline{\mathbb{R}}$, defined for $x \in \mathbb{R}$, as follows:

$$If(x) = \sup_{V \in \mathcal{V}_x} \inf\{z \in f(y) \mid y \in V\},$$
$$Sf(x) = \inf_{V \in \mathcal{V}_x} \sup\{z \in f(y) \mid y \in V\},$$

where $\mathcal{V}_x$ is the set of neighborhoods of x.

Lastly, we also associate to f the function $Ff \in \mathbb{A}$, defined for $x \in \mathbb{R}$, by

$$Ff(x) = [If(x), Sf(x)].$$

Then an interval valued function $f \in \mathbb{A}$ is called **Hausdorff continuous**, if and only if it satisfies the following minimality condition: for every function $g \in \mathbb{A}$, we have for all $x \in \mathbb{R}$ such that $g(x) \subseteq f(x)$

$$Fg(x) = f(x).$$

We denote by $\mathbb{H}$ the set of all Hausdorff continuous functions.

Surprisingly, Hausdorff continuous functions have many of the important properties of usual continuous functions. For instance, if $f, g \in \mathbb{H}$ and if A is a *dense* subset of $\mathbb{R}$, then $f = g$ on A implies $f = g$ on $\mathbb{R}$.

As for the discontinuities of Hausdorff continuous functions, the following property is fundamental. Let $f \in \mathbb{H}$. Then we define $\underline{f}(x), \overline{f}(x)$ such that

$$f(x) = [\underline{f}(x), \overline{f}(x)],$$

where $\underline{f}, \overline{f} : \mathbb{R} \longrightarrow \overline{\mathbb{R}}$, and $\underline{f}(x) \leq \overline{f}(x)$, for $x \in \mathbb{R}$. Let us now consider the set

$$\Gamma_f = \{x \in \mathbb{R} \mid \underline{f}(x) < \overline{f}(x)\},$$

that is, the points $x \in \mathbb{R}$ where the value of the function f is a genuine interval rather than a real or extended real number. Then it can be shown that Γ_f is meager.

Such regularity properties of Hausdorff continuous functions are particularly important in the context of solving PDEs though the order completion method. Obviously, the above definition of Hausdorff continuous functions can be extended to functions defined on any topological space with suitable properties, and thus in particular, to any open set in Euclidean space.

References

[An] Roumen Anguelov: Dedekind order completion of $C(X)$ by Hausdorff continuous functions, *Quaestiones Mathematicae* **27** (2004), 153–170.

[Ar] Vladimir Igorevich Arnold: *Lectures on PDEs*. Springer Universitext, 2004.

[Ev] Lawrence Craig Evans: *Partial Differential Equations*. American Mathematical Society, 1998 (AMS Graduate Studies in Mathematics, **19**).

[Le] Hans Lewy: An example of smooth linear partial differential equation without solutions, *Ann. Math.* **66**#2 (1957), 155–158.

[Lu] W. A. J. Luxemburg , A. C. Zaanen: *Riesz Spaces I*. Amsterdam: 1971.

[Ma] H. M. MacNeille: Partially ordered sets, *Trans. AMS* **42** (1937), 416–460.

[OR] M. B. Oberguggenberger, Elemér E. Rosinger: *Solution of Continuous Nonlinear PDEs through Order Completion*. Amsterdam: Elsevier, 1994 (North-Holland Mathematics Studies **181**).

[Ro1] Elemér E. Rosinger: Can there be a general nonlinear PDE theory for the existence of solutions?, arXiv:math.AP/0407026.

[Ro2] Elemér E. Rosinger: Characterization for the solvability of nonlinear PDEs, *Trans. AMS* **330** (1992), 203-225.

[Ro3] Elemér E. Rosinger: Differential Algebras with Dense Singularities on Manifolds, *Acta Applicandae Mathematicae* **95**#3 (2007), 233–256.

[Ro4] Elemér E. Rosinger: *Distributions and Nonlinear Partial Differential Equations.* New York: Springer, 1978 (Springer Lectures Notes in Mathematics **684**).

[Ro5] Elemér E. Rosinger: Division of Distributions, *Pacif. J. Math.* **66**#1, (1976), 257–263.

[Ro6] Elemér E. Rosinger: Embedding of the D' distributions into pseudotopological algebras, *Stud. Cerc. Math.* **18**#5 (1966), 687–729.

[Ro7] Elemér E. Rosinger: *Generalized Solutions of Nonlinear Partial Differential Equations*. Amsterdam: Elsevier, 1987 (North Holland Mathematics Studies **146**).

[Ro8] Elemér E. Rosinger: Global version of the Cauchy-Kovalevskaia theorem for nonlinear PDEs, *Acta Appl. Math.* **21** (1990), 331-343.

[Ro9] Elemér E. Rosinger: Hausdorff continuous solutions of arbitrary continuous nonlinear PDEs through the order completion method, arXiv:math.AP/0405546.

[Ro10] Elemér E. Rosinger: How to solve smooth nonlinear PDEs in algebras of generalized functions with dense singularities, *Applicable Analysis* **78** (2001), 355–378.

[Ro11] Elemér E. Rosinger: New method for solving large classes of nonlinear systems of PDEs, arxiv:math/0610279.

[Ro12] Elemér E. Rosinger: *Nonlinear Partial Differential Equations, An Algebraic View of Generalized Solutions*. Amsterdam: Elsevier, 1990 (North Holland Mathematics Studies **164**).

[Ro13] Elemér E. Rosinger: *Nonlinear Partial Differential Equations, Sequential and Weak Solutions*. Amsterdam: Elsevier, 1980 (North Holland Mathematics Studies **44**).

[Ro14] Elemér E. Rosinger: Nonsymmetric Dirac distributions in scattering theory, pages 391–399 in *Springer Lecture Notes in Mathematics* **564**, New York: Springer, 1976.

[Ro15] Elemér E. Rosinger: *Parametric Lie Group Actions on Global Generalized Solutions of Nonlinear Partial Differential Equations and an Answer to Hilbert's Fifth Problem.* Dordrecht, London, Boston: Kluwer Acad. Publ., 1998.

[Ro16] Elemér E. Rosinger: Pseudotopological spaces, the embedding of the D′ distributions into algebras, *Stud. Cerc. Math.* **20**#4 (1968), 553–582.

[Ro17] Elemér E. Rosinger: Solving general equations by order completion, arxiv:math/060845.

[Wa1] Jan Harm van der Walt: The uniform order convergence structure on $ML(X)$, arXiv:math/0708.2785.

[Wa2] Jan Harm van der Walt: On the completion of uniform convergence spaces and an application to nonlinear PDEs, arXiv:math/0709.0574.

[Wa3] Jan Harm van der Walt: The order completion method for systems of nonlinear PDEs: Pseudo-topological perspectives, arXiv:math/0706.3990.

[Wa4] Jan Harm van der Walt: Generalized solutions to nonlinear first order Cauchy problems, arXiv:math/0709.1994.

FEATURE

7

MATHEMATICAL MINUTIAE

Irrational Numbers and the Euclidean Algorithm

Brett Harrison[†]
Harvard University '10
Cambridge, MA 01238
bharris@fas.harvard.edu

Remember in middle school when we first learned the difference between **rational** and **irrational** numbers? Informally, we were told that irrational numbers could not be represented as fractions of integers. But now we will see how a number theoretic algorithm based on the simple concept of division can yield fraction representations of irrational numbers.

First, we define the standard division algorithm in $\mathbb{Z}$, the set of integers.

Definition 1 (Division Algorithm)**.** The **division algorithm** in $\mathbb{Z}$ states that for all $a, b \in \mathbb{Z}$, there exist $q, r \in \mathbb{Z}$ such that

$$a = bq + r, \qquad 0 \leq r < b.$$

This algorithm precisely matches our intuition about division in the integers. By recursively applying the division algorithm, we obtain the famous algorithm of Euclid:

Definition 2 (Euclidean Algorithm)**.** The **Euclidean algorithm** in $\mathbb{Z}$ is a repeated division process, beginning with the division algorithm on two integers a and b and proceeding as follows:

$$\begin{aligned}
a &= b \cdot q_1 + r_1, & 0 &\leq r_1 < b,\\
b &= r_1 \cdot q_2 + r_2, & 0 &\leq r_2 < r_1,\\
r_1 &= r_2 \cdot q_3 + r_3, & 0 &\leq r_3 < r_2,\\
&\vdots & &\vdots\\
r_{k-2} &= r_{k-1} \cdot q_k + r_k, & 0 &\leq r_k < r_{k-1}.
\end{aligned}$$

The Euclidean algorithm stops when the remainder in the division algorithm is 0. In the above representation, k is the number of steps in the algorithm, r_{k-1} is the last non-zero remainder, and $r_k = 0$. (A proof that the Euclidean algorithm eventually stops for every pair of integers a and b is left as an exercise to the reader.) We introduce one last definition:

Definition 3 (Fraction Sequence)**.** The **fraction sequence** $[a_1, a_2, a_3, \ldots, a_n]$ is equal to the following continued fraction expansion:

$$a_1 + \cfrac{1}{a_2 + \cfrac{1}{a_3 + \cfrac{1}{\cdots + \cfrac{1}{a_n}}}}$$

[†]Brett Harrison, Harvard '10, is a computer science concentrator. In the past, Brett has conducted research in graph theory, number theory, and Galois theory; he is currently most interested in artificial intelligence. He is also an avid musician who enjoys performing as much as he can. Brett is a founding member of The HCMR, and is currently the Design Director and Webmaster for the organization.

It is a wonderful result in elementary number theory that if a and b are integers, then $\frac{a}{b} = [q_1, q_2, q_3, \ldots, q_k]$, where $q_1, q_2, q_3, \ldots, q_k$ are the sequence of quotients from the Euclidean algorithm on a and b. Furthermore, the division and Euclidean algorithms can be extended for irrational numbers, yielding a similar result for the representing irrational numbers as fraction sequences. I encourage the reader to read more about this extension of the Euclidean algorithm in Niven, Zuckerman, and Montgomery's *An Introduction to the Theory of Numbers* [NZM].

As a consequence, we can write many irrational numbers as infinite fraction sequences. For example, for the golden ratio $\phi = \frac{1+\sqrt{5}}{2}$, the Euclidean algorithm will tell us that $\phi = [\overline{1}] = [1, 1, 1, \ldots]$. This representation also suggests a useful way of approximating irrational numbers, i.e. by computing a finite portion of the infinite fraction sequence

$$\phi \approx [1, 1, 1, 1] = 1 + \cfrac{1}{1 + \cfrac{1}{1 + \cfrac{1}{1}}} \approx 1.667.$$

We can compute as many terms of the fraction sequence we would like to find better and better approximations. So the next time your middle school algebra teachers tell you that you cannot represent irrational numbers with fractions, you had better tell them otherwise!

References

[NZM] Ivan Niven, Herbert S. Zuckerman, and Hugh L. Montgomery: *An Introduction to the Theory of Numbers*, 5th ed. Wiley, 1991.

FEATURE

8

STATISTICS CORNER

Presidential Election Polls: Should We Pay Attention?

Robert W. Sinnott[†]
Harvard University '09
Cambridge, MA 02138
rsinnott@fas.harvard.edu

8.1 Introduction

Every four years, Americans are bombarded with polls tracking the presidential race to the White House. What many Americans and political pundits do not realize, however, is how to synthesize the massive amounts of statistical information thrust upon them. Despite the huge amounts of energy, money, and punditry devoted to predicting election outcomes, successful and accurate prediction is still considered elusive. In the examination provided in this feature, it will become apparent how inaccurate and misleading most media opinion poll reports are and how even professional, respected poll results are often misinterpreted by news broadcasters.

Outside the media, political election prediction has become a far more accurate practice. Econometric time series models based on historical data and causal understanding allow for more consistent, precise prediction of election outcomes than even midday exit polling results can offer. This feature includes a brief explanation of these methods as well as their surprising conclusion.

The use of statistics to tell the difference between systematic aberrations and random noise, however, is even more important than the science's predictive ability. This feature concludes with a discussion of the statistically significant findings regarding voting machine error during the 2004 election.

8.2 Presidential Election Polling

In 2004, midday exit poll results strongly suggested that Senator Kerry would become the next president of the United States. For the previous two weeks, broadcasters Fox News and CNN [Po] presented polls showing Kerry holding a 1-3% lead over his republican opponent. By the end of election day, however, Bush had carried 50.75% of the vote, securing another four years in office.

The question many ask is: how can 4% of the American voting population change their mind about something this important in under a week? The simple answer is that they do not. That 4% difference is the result of two kinds of measurement error: bias and sampling error. It is not necessarily due to a change in public opinion.

Bias arises from flaws in the method of data collection. These flaws can be non-neutral survey questions (**survey bias**), non-random samples of the population (**sample bias**), or even non-random refusal to take the survey (**non-response bias**). Presidential polling has had a long and colorful history of biased polling. Perhaps the most famous was the *Literary Digest* [Li] poll of October 31, 1936. Despite polling 2.3 million people (nearly 2% of the US population at the time), the *Literary Digest* predicted that Alf Landon would carry some 370 electoral college votes and 57% of the

[†]Robert W. Sinnott, Harvard '09, is a statistics concentrator.

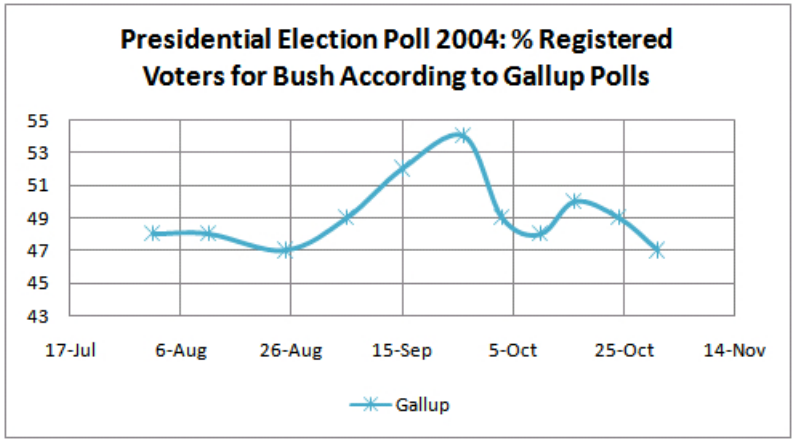

Figure 8.1: *Source: PollingReport.com*

popular vote. Three days later, Franklin D. Roosevelt carried a record 523 of 531 electoral college votes and a full 60.8% of the popular vote.

How could the *Literary Digest* be so far off? Sampling bias. The *Literary Digest* created its survey list by combining telephone and automobile ownership listings. During the Great Depression, owners of automobiles and telephones were hardly a random sample of Americans. At the same time a young upstart by the name of George Gallup randomly polled a mere $5,000$ Americans and correctly predicted FDR's landslide victory. This comparison shows how powerfully survey biases affect results, as even a sample size 500 times larger than the Gallup poll could not counteract the negative effects of sample bias.

Fortunately, such historical mistakes have drawn attention to the potential biases of survey reporting. Currently, only politically motivated and naïve survey groups do not actively correct their methods through sample stratification, imputation, and a large variety of other statistical techniques developed to minimize these biases.

Although the statistical polling companies' methods have vastly improved, the reporting of their results still largely misinterprets the facts. There are two common varieties of misinterpretation: **cherry-picking** data and ignoring **sampling error**.

Leading up to the 2004 election, political pundits would commonly show graphs such as in Figure 8.1 below, citing them as evidence of "dramatic changes in voter opinion."

By looking at the results of this Gallup poll [Po] data, it seems that 5% of voting Americans had changed their minds several times leading up to the election. What most pundits did not include is the other political polling results from the same period. The results from six other respected polling firms are shown in Figure 8.2. By carefully selecting ("cherry-picking") survey results, patterns emerge where none existed before.

Looking at Figure 8.2, it would be hard to discern any pattern at all.

The two logical questions are then, why is there so much variation between polling companies, and why are each company's poll results so inconsistent? The variation between polling companies likely comes from differences in their polling processes, i.e. bias. The poll result inconsistencies potentially come from changes in public opinion, but as will be shown, more likely come from random sampling error.

Sampling error is a measurement of the uncertainty that stems from inferring the state of a population from a study of a random sampling of that same population. News anchors generally report data along the lines of: "47% of Americans support X, plus or minus 2%." In this instance,

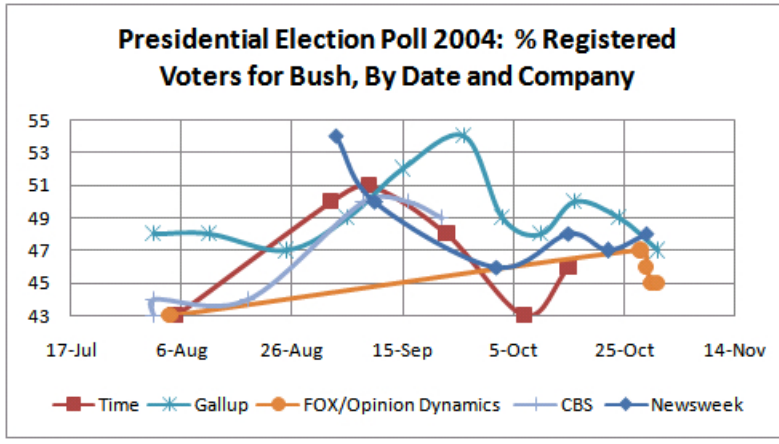

Figure 8.2: *Source: PollingReport.com*

the measure of uncertainty is the 2% margin of error. The problem is that most news reporters do not include this margin of error in their report.

There are dozens of ways of quantifying sampling error in different situations. In political polling the **maximum margin of error** (MMOE) is the standard sampling error measurement. The MMOE is calculated using an approximation to a normal distribution to find the length of the 95% confidence interval. In common terms, the MMOE is the number of percentage points from the estimate for which a statistician is confident the population's actual percentage (called a **population parameter**) will be 95% of the time. For example, a MMOE of 4% for Bush opinion poll result of 48% means that for 95% of sample polls following the same methodology of poll reported, the actual population parameter will be within four percentage points of 48%.

Note a common misconception: this does not mean that the population parameter is within 4% of the estimate 95% of the time. It means that the answer given by the sampling methodology is within 4% of the actual population parameter 95% of the time. The population parameter is fixed, from the perspective of frequentist survey sampling. The estimate and confidence interval generated by the polling data is not.

Intuitively, it makes sense that the larger a sample is, the more certain one should be about the population parameter's true value and the more accuracy (the less sampling error) it should have. This intuition is correct, as evinced by the formula for the MMOE. Using a sample of size n, the calculation for the MMOE is

$$\text{Maximum Margin of Error (95\%)} = 1.96 \cdot \sqrt{\frac{0.5^2}{n}} = \frac{0.98}{\sqrt{n}}.$$

As a result, sample sizes of $n = 2,400$ provide a MMOE of 2%, $n = 600$ provides a MMOE of 4%, and so on. The above calculation is the the maximum of the standard Margin of Error Calculation:

$$\text{Margin of Error (95\%)} = t_{n-1,1-\alpha/2} \times \sqrt{\frac{p(1-p)}{n}}$$

with probability $p = 50\%$, significance $\alpha = 5\%$, and the asymptotic approximation of the t-distribution value $t_{n-1,1-\alpha/2} = 1.96$.[1]

[1]The t-distribution is a special distribution used in statistics in place of the normal distribution when the

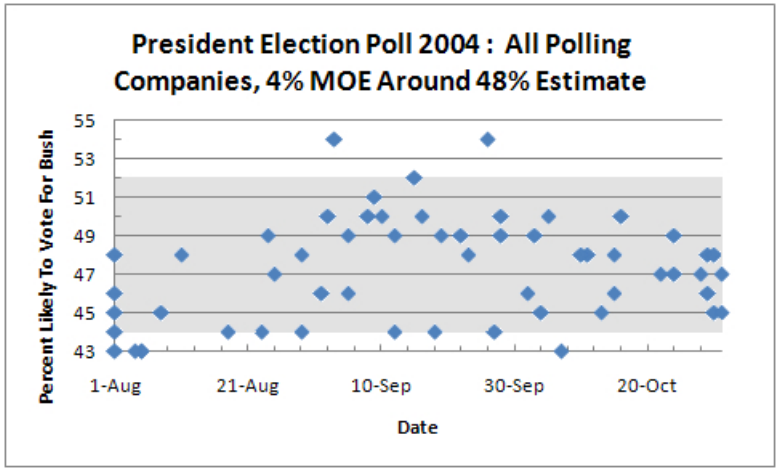

Figure 8.3: *Source: PollingReport.com*

In the figures shown above, the MMOE averages for the different polls to be slightly less than 4%. In Figure 8.3, we show the cumulative results for several polls for Bush as well as a constant line at 48% with a shaded MMOE around it. As you can see, much of the variation can be completely explained by sampling error, with only a few outlying points.

By looking at polling data from this perspective, one must ask whether polling serves any purpose whatsoever (besides giving the political pundits something to talk about). There is no simple answer; however, given the time and labor put into such polling, it seems reasonable to ask if there is a better way to predict the presidential election results.

8.3 Predicting Presidential Elections Without Polls

The advent of statistical computation has allowed for a variety of non-traditional subjects to be analyzed and studied quantitatively. Presidential polling happens to be a subject for which the analysis has been especially successful.

There are a variety of different prediction models used. Not surprisingly, they have varying levels of success. Professor David Walker [Wa] of Georgetown University provides an example of one such model in his 2006 article "Predicting Presidential Election Results."

Walker generalizes Ray Fair's [Fa] work on predicting elections based on non-polling data. The motivation for this approach is to try to use economic and political data to predict public opinion and thereby avoid biasing it through survey methods. Surprisingly, the predictive model is exceedingly simple:

$$\text{Vote} = b_0 + b_1 * \text{Grow} + b_2 * \text{Inflat} + b_3 * \text{Warx} + b_4 * \text{Gdnews} + \epsilon_{t-1}$$

where Grow, Inflat, and Gdnews are, respectively, values calculated based on the economic growth, monetary inflation, and political news at the time of the election. Also, Warx is also used as an

actual mean and standard deviation of the population is not known. When $n \to \infty$ the two distributions are equivalent. The useful part about both the t-distribution and the normal distribution is that the standard deviation and the mean are unaffected by each other, allowing statisticians to use them to calculate probabilities in a wide variety of situations. It is equivalent to a normal distribution times $\sqrt{n}$ and then divided by a sum of n squared unit normal distributions

indicator of whether the country is actively engaged in a war. After calibrating this model against sixteen previous elections, Walker's model [Wa] predicted Bush would carry the 2004 election with 52.3% of the popular vote, a mere 1.6% off from the actual vote totals.

More impressively, other more complex econometric models such as those from Hibbs, Abromowitz, and Wlezien and Erikson predicted Bush victories by 53%, 53.7%, and 52.3% respectively (see [Wa]). However, these models only used data available from before the previous August, a full three months before the elections! These models are beyond the scope of this feature, but they give a flavor of the power of advanced econometric techniques.

Of course, as cautioned above, the selection of these econometric models for this feature is akin to the cherry-picking of poll data from before. Take the predictive merits of these modeling techniques with a grain of salt.

Without question, however, there is the potential for successful election prediction in using econometric models. If you are interested in trying something like one of these models out on your own, you should examine Fair's model [Fa] online, at `http://fairmodel.econ.yale.edu/vote2008/index2.htm`.

8.4 Anomalies in 2004 Election Results

Perhaps the most scintillating use of statistics in the 2004 presidential election was in the comparison of exit poll results to the official vote results. In 2004, an exit poll was conducted by Edison Media Research and Mitofsky International, two highly regarded public opinion polling firms. Immediately after the election, their exit poll results showed a 3% Kerry lead, whereas the official results showed a Bush victory of 2.5% (see [USC]). As one can imagine, such a large discrepancy drew national attention. Several possibilities were proposed for the large sampling error, ranging from random sampling error to insidious conspiracy theories.

After a quick analysis of the data, one can quickly conclude that random chance was not the culprit for the discrepancy between the exit polls and the official results. In their analysis, seven of 50 state results were found to have t-values of less than -2.7, ie, each state had less than 1% probability of having such a large error. Cumulatively, the possibility of having all of these states having errors this large occur is astronomically small, less than 1×10^{-7} (see [USC]).

Once again, a more likely cause for this discrepancy comes from sampling bias. That is to say, more Democrats may have taken part in exit polls than Republicans. As the example of the Literary Digest in 1936 demonstrates, such a large discrepancy could easily be explained by non-response bias. However, ongoing research [USC] currently suggests that if anything, more Republicans take part in exit polls than Democrats.

The research on this topic is ongoing, but one thing is for sure: statistics have shown their uses in finding the counterintuitive task: finding non-random information within the random events that occur every day.

References

[Fa] Ray C. Fair: The effect of economic events on votes for President: 2000 update. `http://fairmodel.econ.yale.edu/RAYFAIR/PDF/2002DHTM.HTM`.

[GK] Andrew Gelman and Gary King: Why are American presidential election campaign polls so variable when votes are so predictable?. *B. J. Pol. Sci.* **23** (1993), 409–451.

[Li] Landon, 1,293,669; Roosevelt, 972,897. *Literary Digest* **31** (Oct.1936), 5–6.

[Po] PollingReport.com: *National 2-way Trial Heat Summary*. `http://pollingreport2.com/wh2004a.htm#2way` [Accessed on October 11, 2007].

[USC] US Counts: US count votes: Study of the 2004 presidential election exit poll discrepancies. http://uscountvotes.org/ucvanalysis/uscountvotes_remitofsky-edison.pdf.

[Wa] David A. Walker: Predicting presidential election results. *Applied Econometrics* **38** (2006), 483–490.

[WE1] Christopher Wlezien and Robert S. Erikson: The horse race: What polls reveal as the election campaign unfolds, *International Journal of Public Opinion Research* **19**#1 (2006), 74–88.

[WE2] Christopher Wlezien and Robert S. Erikson: Presidential polls as a time series. *Public Opinion Quarterly* **63** (1999), 163–177.

[WE3] Christopher Wlezien and Robert S. Erikson: Campaign effects in theory and practice. *American Politics Research* **29**#5 (2001), 419–436.

[Wo] Jeffrey M. Wooldridge: *Econometric Analysis of Cross Section and Panel Data*. Cambridge: MIT Press, 2002.

FEATURE

9

APPLIED MATHEMATICS CORNER

Fireflies & Oscillators

Pablo Azar†
Harvard University '09
Cambridge, MA 02138
azar@fas.harvard.edu

9.1 Introduction

In 1990, R. Mirollo and S. Strogatz [MS] presented a coupled oscillator model explaining how synchronicity arises from self-organization. In this article, I will describe this model and its main result, as well as some applications to computer science. As the methods of pure mathematics become more and more complicated, it is very exciting to see how nature can still be explained and imitated using simple models.

9.2 Synchronization and Self-Organization

When ants forage for food and birds flock together, they form patterns that betray an ability to communicate. **Self-organization** is the idea that these patterns arise from local interactions (see [CDF]). Instead of following a master plan or being guided by a leader, individuals react to their local environments by following simple rules. The aggregation of these small actions is what leads to the final complicated pattern. Because fish look at their nearest neighbors to decide on their directions and velocities, they end up swimming in tight formations that are useful for evading predators. Ants decide where to forage for food by perceiving the local levels of pheromones deposited by other members of their colonies.

In this exposition, I will present a model of synchronous firefly flashing: males flash in unison to attract females. The purpose of this feature is to serve as an introduction to the study of self-organization and an invitation for the reader to investigate this topic further.

9.3 From Facts to Model

I will start by presenting some basic facts about fireflies and by showing a possible translation of these facts into a working mathematical model. An interesting first fact is the following:

Fact 1. *A firefly can produce periodic pulses of light.*

We can use this fact to start building the model. Let t denote time. Since the flashes are periodic, there must be a periodic function $f(t)$ with image $[0, 1]$ that measures how close the firefly is to flashing. When $f(t) = 1$ the firefly flashes and f falls back to 0. If the period of the firefly is ω then we have $f(t + \omega) = f(t)$. Though this model looks good, it has a fatal flaw: it does not explain synchronization. Two fireflies with the same ω will only synchronize if they start their cycles at the same moment, that is, only if they were synchronized from the beginning! To fix this, we must incorporate another fact into our model:

†Pablo Azar, Harvard '09, is an applied mathematics concentrator from Buenos Aires, Argentina. He is a founding member of The HCMR and currently serves on The HCMR's staff.

Fact 2. *A firefly can control the frequency of its pulses.*

To incorporate this fact, let $\phi(t)$ denote the time kept by the firefly's internal clock, expressed as a function of the time measured by objective clocks in the environment. The firefly flashes when $f(\phi(t)) = 1$, and f is periodic as a function of ϕ. This model also describes the behavior of a circuit that consists of a capacitor and a light bulb: $\phi(t)$ is the charge accumulated in the capacitor at time t, while $f(\phi(t))$ is the voltage. When the voltage reaches a certain threshold, the capacitor is discharged, illuminating the light bulb. A change in $\phi(t)$, which signifies a change in the speed at which the firefly thinks time is passing, is analogous to a change in the amount of energy it is going to use to charge the capacitor. There are two more things to be said about this analogy. The first is that as charge accumulates, voltage increases. Thus, f should be increasing as a function of ϕ. The second is that the first units of charge should have a larger impact than the succeeding units of charge, so that the rate of increase of f diminishes as f increases. It will not hurt to assume that f is smooth, so we can state these assumptions as $\frac{df}{d\phi} > 0$, $\frac{d^2f}{d\phi^2} < 0$. These assumptions, of course, are only valid for within one cycle. They are not valid in the isolated moments of time when the firefly flashes and f drops immediately from 1 to 0.

The first model suggests that this charge always accumulates at a constant rate, but we want the firefly to be able to change this rate. The reason for this is

Fact 3. *A firefly responds to other fireflies.*

When one firefly sees other fireflies flash, it will want to accelerate the rate at which it accumulates charge, in order to bring itself closer to flashing.[1] Since the function f measures how close the firefly is to flashing, increasing f by a constant amount ϵ is a good response to a neighbor's flash. Of course, the firefly only controls the argument $\phi(t)$. This increase must be achieved by updating $\phi(t)$ to some $\phi'(t)$ for which $f(\phi'(t)) = f(\phi(t)) + \epsilon$. Note that, since f is concave, the change $\phi'(t) - \phi(t)$ must become larger as $f(t)$ approaches 1. If this boost brings the voltage to something larger than one, the firefly flashes, the potential is reset to zero, and the individual becomes synchronized with its neighbors.[2]

Why do we have this strange rule? Suppose that the firefly followed instead the policy of increasing the argument $\phi(t)$ always by a constant magnitude δ to $\phi'(t) = \phi(t) + \delta$. Then, because f is concave, the corresponding increases in f would get smaller as f is closer to 1. This is very inefficient if we want to achieve synchronicity quickly.

9.4 From Model to Facts

So far, our model is based on a few simple facts and some modeling assumptions. These assumptions are not carved in stone. For example, this model assumes that the effect of a firefly on its

[1] Why would a firefly want to do this? In this feature we explain how male fireflies synchronize to attract females. The question of why this would be evolutionary advantageous leads to an interesting intersection between game theory and biology. These models on evolutionary theory not only attempt to explain evolution, but also lend themselves to many applications such as optimization. If you are interested, some seminal references are Richard Dawkins' *The Selfish Gene* for the biology/game theory aspect and John H Holland's *Adaptation in Natural and Artificial Systems* for some non-biological applications.

[2] More formally, if we give each firefly an index i from 1 to n then:

$$f(\phi_i(t)) = 1 \implies \lim_{s \to t^+} f(\phi_j(s)) = \min(f(\phi_j(t)) + \epsilon, 1), \text{ for } i \neq j.$$

Note that what is hidden behind this equation is a change in $\phi_j(t)$. All fireflies share the same potential function f, but each can only accelerate its own internal, subjective time $\phi_j(t)$. The firefly must hence find a new value $\phi'_j(t)$ such that $f(\phi'_j(t)) = f(\phi_j(t)) + \epsilon$. Now, we had reasons to assume that f was strictly increasing. Hence, f has an inverse g and we can compute $\phi'_j(t) = g(\min(f(\phi_j(t)) + \epsilon))$. The reader who believes that we are assuming the problem away by introducing more notation should take relief by knowing that in applications, f is a familiar concave function like $\log(x)$, which has a familiar inverse e^x.

neighbors is a discrete boost, but other models work just as well assuming continuous boosts. Also, we have to watch out that this model does not lead to a dead end like the first one shown does. In fact, the model achieves what we want and more. First, it explains self-synchronization. Assume that we have n fireflies with initial charges $\phi_1(0), ..., \phi_n(0)$, such that each firefly boosts its neighbors' potentials by ϵ_i when it flashes. The values $(\phi_1(0), ..., \phi_n(0), \epsilon_1, ..., \epsilon_n)$ are the parameters of the model. Mirollo and Strogatz [MS] prove the following:

Theorem 4. *Let a model with n fireflies have parameters $\phi_1(0), ..., \phi_n(0), \epsilon_1, ..., \epsilon_n$. Assume that when a firefly flashes, it boosts all its neighbors. Then, for all initial parameters except those on a set of measure zero, there exists a time t_{conv} such that $f(\phi_1(t_{conv})) = f(\phi_2(t_{conv})) = ... = f(\phi_n(t_{conv}))$. At this moment, the fireflies are synchronized.*

This is equivalent to the following biological fact:

Fact 5. *If all fireflies in a group influence each other, they will eventually become synchronized.*

From a few simple facts about fireflies, one can build a mathematical model that explains how they achieve synchronization. But we can do even better. One of the most unrealistic assumptions in the Mirollo-Strogatz model is that one firefly influences all the others. However, self-organization relies on individuals reacting to their *local* environment. Fourteen years after Mirollo and Strogatz obtained Theorem 4, Lucarelli and Wang [LW] modified their model to obtain synchronization under the assumption that a firefly can only be influenced by its nearest neighbor.

There is one more thing that vouches for the Mirollo-Strogatz model of firefly synchronization: its simplicity and flexibility make it very applicable. The model can explain nature and can also imitate it. Based on the Mirollo-Strogatz model, Werner-Allen, Tewari, Patel, Welsh, and Nagpal [WTP] have developed the **Reachback Firefly Algorithm** to induce synchronicity in sensor networks.

The applications of this reasoning go far beyond fireflies. For example, suppose that a robot had an artificial eye with millions of tiny sensors. It would be inefficient for each sensor to send its information to the central processing unit at idiosyncratic intervals. We would rather have all sensors send their information simultaneously. Similarly, suppose that we had thousands of small, cheap robots exploring a foreign planet and controlled by a central base. We would want all robots to send their information simultaneously, especially if the central base had to make very quick decisions.

Finally, a word must be said about simplicity. What first attracted me to biologically inspired models was that they explained nature with new but simple concepts. Applied mathematics does not necessarily move forward by becoming more abstract and complicated. Sometimes, new developments arise with just simple math and common sense.

References

[CDF] Scott Camazine, Jean-Louis Deneubourg, Neigel R. Franks, James Sneyd, Guy Theraulaz, and Eric Bonabeau: *Self-Organization in Biological Systems*. Princeton University Press, 2003.

[LW] Dennis Lucarelli and I-Jeng Wang: Decentralized synchronization protocols with nearest neighbor communication. *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems* (2004), 62–68.

[MS] Renato E. Mirollo and Steven H. Strogatz: Synchronization of Pulse-Coupled Biological Oscillators, *SIAM J. on App. Math.* **50** #6 (1990), 1645–1662.

[WTP] Geoffrey Werner-Allen, Geetika Tewari, Ankit Patel, Matt Welsh, and Radhika Nagpal: Firefly-inspired sensor network synchronicity with realistic radio effects. *Proceedings of the Third International Conference on Embedded Networked Sensor Systems* (2005), 142–153.

FEATURE

10

MY FAVORITE PROBLEM

Bert and Ernie

Zachary Abel[†]
Harvard University '10
Cambridge, MA 02138
zabel@fas.harvard.edu

10.1 Allow Me to Introduce Myself

The primary driving force behind my mathematical career up to now—through recreation, competition, and research—has been problem solving. This problem solving process allows a two-way channel of communication between myself and my mathematical experience. First, solving a problem allows me to draw from my current cache of **intuitions**, **bits of knowledge** (or in the words of Scott Kominers, **knowledgecules**), and **ideas** that might be applicable to the current challenge, thus facilitating directed reflection into my current mathematical understanding. In turn, the effort exerted in solving (or at least working on) the problem only adds to this cache and strengthens this same understanding. This cyclic (singly generated!) process of reflection and growth turns the act of solving problems—and indeed of studying mathematics—into a beautiful and highly personal experience.

While perusing my problem repertoire searching for a "favorite," I had trouble pinpointing a leader because different partial orderings—trickiness, elegance, difficulty (unsolved?), cuteness...—lead to vastly different maximal elements. So I finally decided to choose a problem that best illustrates the process above: specifically, one that illustrates the powerfully personal nature of mathematical study.

And now, we present the problem:

Bert and Ernie. *Bert is thinking of an ordered quadruple of integers (a, b, c, d). Ernie, hoping to determine these integers, hands Bert a 4-variable polynomial $P(w, x, y, z)$ with integer coefficients, and Bert returns the value of $P(a, b, c, d)$. From this value alone, Ernie can always determine Bert's original ordered quadruple. Construct, with proof, one polynomial that Ernie could have used.*

To simplify discussion, allow me to strip the PBS language. The following problem is equivalent:

No More Bert and Ernie. *Find, with proof, a polynomial $P \in \mathbb{Z}[w, x, y, z]$ so that $P : \mathbb{Z}^4 \hookrightarrow \mathbb{Z}$ is injective.*

We are thus given two tasks: extract multiple pieces of information from a single integer (namely $P(a, b, c, d)$), and do so using integer polynomials. If we were dealing with polynomials but not necessarily integers, we could use a polynomial like

$$P(w, x, y, z) = w\sqrt{2} + x\sqrt{3} + y\sqrt{5} + z\sqrt{7}$$

[†]Zachary Abel, Harvard '10, is a computer science and mathematics concentrator. He is an avid problem solver and researcher, with interests in such varied fields as computational geometry, number theory, partition theory, category theory, and applied origami. He is a founding member of The HCMR and currently serves as Problems Editor, Issue Production Director, and Graphic Artist.

and rely on the fact that the set $\{\sqrt{n} \mid n \in \mathbb{N} \text{ is squarefree}\}$ is linearly independent over the rationals. Or, if we were restricted to integers but not to polynomials, we could easily set

$$P(w, x, y, z) = 2^w 3^x 5^y 7^z$$

and use unique factorization to recover the exponents. The difficulty arises from the combination.

This double condition forces a convergence of algebraic and number-theoretic ideas. Beyond that, however, anything goes: the problem does not further restrict the range of useful directions of exploration. Since a wide variety of ideas can be usefully applied to the problem, a solver reveals a great deal about his or her problem solving process simply by writing down the final proof.

10.2 Draw from Knowledgecules

Anyone who has asked about the cardinality of $\mathbb{N}^2$ (or $\mathbb{Q}$) and been shocked to find that it equals that of $\mathbb{N}$ has undoubtedly stumbled across the enumeration of the $\mathbb{N}^2$ as depicted in Figure 10.1. Perhaps they have also noticed that this enumeration can be written down as a rational polynomial:

$$\ell(x, y) = \frac{(x + y - 2)(x + y - 1)}{2} + y.$$

This means that $2\ell : \mathbb{N}^2 \to \mathbb{N}$ is an injective integer polynomial! Having this fact in our repertoire, all we have to do now is find a way to replace $\mathbb{N}^2$ with $\mathbb{Z}^4$.

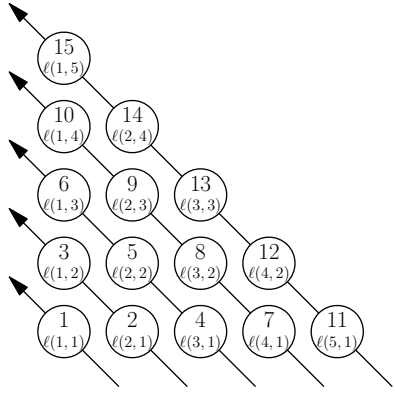

Figure 10.1: A proof that $|\mathbb{N}^2| = |\mathbb{N}|$ by enumerating $\mathbb{N}^2$ along diagonals.

Solution 1 by Nick Wage, '10. Use the notation $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. We construct the following polynomials in order:

$$\begin{aligned}
A &: \mathbb{N}_0^2 \to \mathbb{N}, & (x, y) &\mapsto 2\ell(x + 1, y + 1) \\
B &: \mathbb{N}_0^4 \to \mathbb{N}, & (w, x, y, z) &\mapsto A(A(w, x), A(y, z)) \\
C &: \mathbb{Z}^2 \to \mathbb{N}, & (x, y) &\mapsto B(x^2, (x + 1)^2, y^2, (y + 1)^2) \\
P_1 &: \mathbb{Z}^4 \to \mathbb{N}, & (w, x, y, z) &\mapsto C(C(w, x), C(y, z)).
\end{aligned}$$

It is clear from these definitions that A, B, C, and P_1 are integer-coefficient polynomials. We now show that each, in turn, is also injective, so that P_1 is the desired polynomial.

Polynomial $A(x, y) = 2\ell(x + 1, y + 1)$ is injective, as mentioned in the previous paragraph. This means there exists a pair of well-defined left inverses A_x and A_y (not necessarily polynomial) so that $A_x(A(x, y)) = x$ and $A_y(A(x, y))$. So for any $(w, x, y, z) \in \mathbb{N}_0^4$ we have

$$A_x(A_x(B(w,x,y,z))) = w, \qquad A_x(A_y(B(w,x,y,z))) = x,$$
$$A_y(A_x(B(w,x,y,z))) = y, \qquad A_y(A_y(B(w,x,y,z))) = z.$$

This means that we can decipher (w, x, y, z) from $B(w, x, y, z)$, whence B is injective. To see that C is injective, suppose we know the values of x^2 and $(x + 1)^2$ (which we will, by B's injectivity) for some integer x. If $(x+1)^2 < x^2$, x must be negative, i.e. $x = -\sqrt{x^2}$, and if $(x+1)^2 > x^2$ then x is non-negative, i.e. $x = \sqrt{x^2}$. So we may uniquely determine x from $C(x, y) = B(x^2, (x + 1)^2, y^2, (y + 1)^2)$, and likewise for y. So C is injective. Finally, P_1's injectivity follows from that of C by the same argument used for B. So this P_1 does indeed solve the problem. □

In order to turn 2ℓ into a full solution (i.e. to replace $\mathbb{N}$s with $\mathbb{Z}$s in the domain), this solver used the clever injective map $\mathbb{Z} \hookrightarrow \mathbb{N}^2,\ x \mapsto (x^2, (x + 1)^2)$. Other methods of injecting $\mathbb{Z}$ into $\mathbb{N}$ may be used for alternate solutions. For example, having encountered (during some dabbling in partition theory) **Euler's Pentagonal Number Theorem**, namely the result

$$\prod_{k=1}^{\infty}(1 - x^k) = \sum_{i=-\infty}^{\infty}(-1)^k x^{k(3k+1)/2} = 1 - x - x^2 + x^5 + x^7 - x^{12} - x^{15} + x^{22} + x^{26} - \cdots,$$

(note that no power of x is hit twice by the sum), I recognized that $m(x) = \frac{x(3x+1)}{2}$ is such an injective polynomial $\mathbb{Z} \hookrightarrow \mathbb{N}_0$. This gives rise to the next solution.

Solution 2 by the author. Define $n(x) = 2m(x) + 1 = x(3x + 1) + 1$, which has integer coefficients and strictly positive values. Note that $n(x) = n(y)$ for any integers x and y implies that $x = y$ or $x + y = -\frac{1}{3}$, i.e. $x = y$. Thus, $n : \mathbb{Z} \hookrightarrow \mathbb{N}$ is injective. So we may define the injective polynomial $D : \mathbb{Z}^2 \hookrightarrow \mathbb{N}$ by $(x, y) \mapsto 2\ell(n(x), n(y))$, and the final solution $P_2(w, x, y, z) = D(D(w, x), D(y, z))$ is found as above. □

10.3 Draw from Intuition

Consider the simple **floor** (or **greatest integer**) **function** $\lfloor\cdot\rfloor$. Given any number r, this function tells us two important pieces of information: macroscopically, the unit interval in which r lies, namely the half-open interval $[\lfloor r\rfloor, \lfloor r\rfloor + 1)$; and microscopically, how far away we landed from that integer, $r - \lfloor r\rfloor$. Intuitively, we have a number of larger targets equipped with small, disjoint regions for error. This separation into primary and error terms comes most directly from analysis, by approximating a function locally with its linear derivative and quadratic error. But the method has certainly been put to good use in other ways: for example, Hamming's "error correcting codes" tightly pack disjoint balls in $\mathbb{F}_2^n$ in order to (1) identify a code word even if there were a few errors in the transcription, and (2) locate those errors.

This disjoint wiggle-room intuition can be very beneficial for the problem at hand by dividing the given output into a primary, large value with a relatively tiny error term added on, both of which are thus uniquely determined. The fuzzy idea of exploiting big and small terms (perhaps by throwing in huge exponents and coefficients!) that comes from the above intuition is certainly enough to solve the problem, especially when one leans on the intuitive notions of big and small as elucidated by the infamous "Big 'O' Notation." A rather nice way of combining these ideas is explained below.

Solution 3 by the author. Suppose we have two positive integers a and b with $a < b$. I claim that the value $b^2 + a$ uniquely identifies both a and b. Indeed, since $b^2 + 1 \le b^2 + a \le b^2 + b - 1$, and the intervals $[c^2 + 1, c^2 + c - 1]$ and $[(c + 1)^2 + 1, (c + 1)^2 + (c + 1) - 1]$ are disjoint for

all positive integers c (since $c^2 + c - 1 < (c+1)^2 + 1$), the value $b^2 + a$ falls into at most one such interval, which uniquely determines b. The value of a follows. We thus obtain an injection $E : \mathbb{Z}^2 \hookrightarrow \mathbb{N},\ (x, y) \mapsto n(x) + (n(x) + n(y))^2$ (where n is the injection $\mathbb{Z} \hookrightarrow \mathbb{N}$ from above), and as usual the final solution $P_3(w, x, y, z) = E(E(w, x), E(y, z))$ is immediate. □

In fact, the lemma used above may be strengthened:

Lemma 1. *For fixed n and k, the equation $k = a_n^n + a_{n-1}^{n-1} + \cdots + a_1$ has at most one solution in positive integers $a_1, \ldots, a_n$ subject to the condition $0 < a_1 < a_2 < \cdots < a_n$.*

Proof. We have the inequality

$$\begin{aligned} a_n^n + 1 < a_n^n + \cdots + a_1 &< a_n^n + a_n^{n-1} + \cdots + a_n \\ &< a_n^n + \tbinom{n}{n-1} a_n^{n-1} + \cdots + \tbinom{n}{1} a_n + 2 = (a_n + 1)^n + 1, \end{aligned}$$

so the value of $a_n^n + \cdots + a_1$ falls into at most one interval of the form $[c^n + 1, (c+1)^n]$, which uniquely defines the value of a_n. The result then follows by induction. □

We therefore obtain a slightly more elegant(?) solution.

Solution 4 by Scott Kominers '09 and the author. The polynomial $P_4 : \mathbb{Z}^4 \to \mathbb{N}$ defined by

$$\begin{aligned} P_4(w, x, y, z) = n(w) &+ \big(n(w) + n(x)\big)^2 + \big(n(w) + n(x) + n(y)\big)^3 \\ &+ \big(n(w) + n(x) + n(y) + n(z)\big)^4 \end{aligned} \tag{10.1}$$

is injective by Lemma 1. □

Note that not only is the error-term intuition useful for solving the problem, but the proof itself—even simply line (10.1) alone—clearly elucidates the idea that the solver had in mind.

10.4 Draw from Idea

Speaking of ideas, here are two more, very ill-formed ideas toward different solutions, both derived from number theory (as a healthy break from the analysis and algebra influences above). We are told to extract multiple pieces of information from a single integer, so we can probably do something with the **representations** of this integer: indeed, the (decimal or binary) **digits** of a number give lots of distinct pieces of information, as do the (prime) **divisors**.

It is often the case that one can store valuable polynomial information in the digits of a number. For example, it is known that if $a_n \ldots a_1 a_0$ is the base-10 expansion of a prime q, then the polynomial $p(x) = a_n x^n + \cdots + a_1 x + a_0$ is irreducible [BFO], and the proof relies heavily on the fact that $p(10) = q$. For this problem, though, I found it difficult to directly apply this digit-storing idea, as most of the potential solutions I found along these lines ended up being exponential, not polynomial. However, a slight generalization of the idea of **digit** leads perhaps to base Fibonacci,[1] base factorial,[2] or base -4 representations,[3] none of which I could get to work here. Another slight generalization of the digit notion—never fully giving up on this idea, but instead running as far as necessary with it—perhaps leads to quadratic form representation theory and representation as sums of (two or more) squares, which, unfortunately, is not usually unique (except for primes $p \equiv 1$ or $2 \bmod 4$ in the two squares case). This in turn recalls Waring's problem and representations as sums of higher powers. Finally, perhaps, one considers sums of *different* powers, and is thus lead to a solution similar to Solution 4. As the increasing powers a_i^i recall the increasing

[1] e.g. IMO 1993 #5
[2] e.g. AIME-II 2000 #14
[3] e.g. USAMO 1996 #4

exponents of the base in usual base-number representation, we are really not too far away from the initial base-digit idea.

The other idea was that of (prime) factors. Suppose we had an injective polynomial $p(x)$ that only output primes. Then knowing the value of $n = p(x) \cdot p(y)$ would uniquely tell you the (unordered) set $\{x, y\}$ by simply factoring n. Unfortunately, such a polynomial does not exist,[4] but the idea may still be useful, and prime factors and polynomials can mix well together in other ways. For example, given any $x \in \mathbb{Z}$, all prime divisors of $x^2 + 1$ must be of the form $p \equiv 1$ or $2 \bmod 4$; cyclotomic polynomials Φ_m directly generalize this fact by only allowing prime divisors of $\Phi_m(x)$ to be congruent to $1 \bmod m$ with finitely many exceptions.[5] However, as I have not yet found a solution down this road, let us move back to the original *divisors* idea, this time throwing out the *prime* part.

Divisors always come in pairs, but how can we distinguish a particular pair? Can we distinguish between pairs of divisors if there is only one nontrivial pair? (Yes, but then $n = pq$ is a product of two primes, and we are back to primes.) What if both elements in the divisor pair are equal? Then $n = a^2$, and we do not get multiple values. But what if the divisors are *almost* equal? I.e., what if we pick the *closest* pair of divisors? In this case, the pair is certainly uniquely defined.

Along these lines, we would like to be able to say that if $m \gg r$, and if a and b with $a < b$ have $ab = m(m + r)$, then $a \leq m < m + r \leq b$, i.e. the closest pair of divisors of $m(m + r)$ is in fact $(m, m + r)$. Indeed, a lemma of this form is not difficult to prove:

Lemma 2 (1998 St. Petersburg City Mathematical Olympiad)**.** *Let n be a positive integer. Show that any number greater than $n^4/16$ can be written in at most one way as the product of two of its divisors having difference not exceeding n.*

Proof (method by Titu Andreescu and Dorin Andrica). Suppose $a < c \leq d < b$ with $ab = cd = t$ and $b - a \leq n$. Note that

$$(a + b)^2 - n^2 \leq (a + b)^2 - (b - a)^2 = 4ab = 4cd = (c + d)^2 - (d - c)^2 \leq (c + d)^2,$$

so that $(a + b)^2 - (c + d)^2 \leq n^2$. But as $a + b > c + d$ (since the function $f : x \mapsto x + t/x$ decreases for $t < \sqrt{x}$, which means $f(a) > f(c)$), we find

$$n^2 \geq (c + d + 1)^2 - (c + d)^2 = 2c + 2d + 1.$$

Finally, the AM-GM inequality gives

$$t = cd \leq \left(\frac{c + d}{2}\right)^2 \leq \frac{(n^2 - 1)^2}{16} < \frac{n^4}{16},$$

proving the claim. □

Thus armed, we arrive at our last solution to the Bert and Ernie problem:

Solution 5 by the author. Consider the polynomial $F : \mathbb{Z}^2 \to \mathbb{N}$ defined by

$$F(x, y) = (n(x) + n(y))^2 \cdot \left((n(x) + n(y))^2 + n(x)\right).$$

As $F(x, y) > n(x)^4 > n(x)^4/16$, and as the difference of the two factors is exactly $n(x)$, Lemma 2 proves that, for fixed x and y, the factorization $F(x, y) = ab$ with $a \leq b$ and $b - a \leq n(x)$ is unique. Taking (a, b) to be the pair of divisors of $F(x, y)$ that are closest together, we must obtain

$$a = (n(x) + n(y))^2, \qquad b = (n(x) + n(y))^2 + n(x).$$

[4] Indeed, if an integer polynomial f has only prime outputs, then since the prime $q = f(1)$ divides all numbers $f(1 + kq)$ for $k \in \mathbb{Z}$, all of these prime values must be $\pm q$. But then f takes one of those values infinitely many times, so f is constant.

[5] Specifically, if $p \mid \Phi_m(x)$, then either $p \equiv 1 \bmod m$ or $p \mid m$, a result apparently proved by Legendre [Ga].

From here, x and y can be reconstructed, whence F is injective. Then

$$P_5(w, x, y, z) = F(F(w, x), F(y, z))$$

solves the problem. □

The preceding solution is an attempt to illustrate some of my thinking while engaging the two ideas mentioned above. It follows a depth-first-like traversal through the directed graph of free idea association, moving to a related node if the current ideas become exhausted or seem unfruitful. Whether or not my brain actually thinks in depth-first terms (or perhaps it is performing greedy best-first, or even A^* search), exploring these associations can certainly be a useful exercise. Almost certainly, the resulting graph traversal will differ greatly from person to person.

10.5 A Parting Challenge: The Bert and Ernie Contest

Now that you have seen my thoughts and approaches for this problem, I would love to see yours! I hereby present the following challenge:

The Bert and Ernie Contest. *Show us how* you *would solve my favorite problem.*

You are invited to submit a different solution based on ideas or premises not discussed here, along with a short description of the methods and approaches used. Successful submissions will be acknowledged both on The HCMR's website and in future issues; the most novel and illuminating will be published. Submissions for this Bert and Ernie Contest should be directed to me (Zachary Abel), either at `hcmr-problems@hcs.harvard.edu` or at the address on the inside front cover.

10.6 Acknowledgments

I am very grateful to, among others and in no particular order, Menyoung Lee, Daniel Litt, Brett Harrison, Scott Kominers, Eleanor Birrell, Brian Basham, Eddie Keefe, Zachary Galant, Ernie Fontes, Alex Zhai, Professor Noam D. Elkies, Dr. Grant Mindle, and Dr. Barbara Currier, for stimulating discussion relating to the Bert and Ernie problem, and of course to Scott Kominers and Nick Wage for sharing their solutions. I would also like to thank my father, Dr. Bruce J. Abel who inspired me to invent this problem. I am once again indebted to Daniel Litt, Scott Kominers, and Eleanor Birrell for their helpful comments on earlier drafts of this article. Finally, I thank you in advance for your future submissions to The Bert and Ernie Contest.

References

[BFO] J. Brillhart, M. Filaseta, and A. Odlyzko: On an irreducibility theorem of Cohn, A., *Canadian Journal of Mathematics* **33** #5 (1981), 1055–1059.

[Ga] Yves Gallot: Cyclotomic polynomials and prime numbers (2000, Revised 2001), `http://pagesperso-orange.fr/yves.gallot/`

FEATURE

11 Problems

The HCMR welcomes submissions of original problems in any fields of mathematics, as well as solutions to previously proposed problems. Proposers should direct problems to `hcmr-problems@hcs.harvard.edu` or to the address on the inside front cover. A complete solution or a detailed sketch of the solution should be included, if known. Solutions to previous problems should be directed to `hcmr-solutions@hcs.harvard.edu` or to the address on the inside front cover. Solutions should include the problem reference number, as well as the solver's name, contact information, and affiliated institution. Additional information, such as generalizations or relevant bibliographical references, is also welcome. Correct solutions will be acknowledged in future issues, and the most outstanding solutions received will be published. To be considered for publication, solutions to the problems below should be postmarked no later than *March 1, 2008*. An asterisk beside a problem or part of a problem indicates that no solution is currently available.

F07 – 1. Consider $\triangle ABC$ an arbitrary triangle and P a point in its plane. Let D, E, and F be three points on the lines through P perpendicular to the lines $\overline{BC}$, $\overline{CA}$, and $\overline{AB}$, respectively. Prove that if $\triangle DEF$ is equilateral and if P lies on the Euler line of $\triangle ABC$, then the center of $\triangle DEF$ also lies on the Euler line of $\triangle ABC$.

Proposed by Cosmin Pohoata (Bucharest, Romania) and Darij Grinberg (Germany).

F07 – 2. Professor Perplex has rounded up his $n > 0$ hat-game seminar students and made the following ominous announcement:

> "I have assigned each of you a hat according to a uniform probability distribution, which I will put on your head after allowing you time to discuss a strategy. Hats come in $h > 0$ different colors, but some colors might be reused and others might not be used at all. Each student will be given a list of the h colors. Nobody will be able to see his or her own hat, but everyone will have the opportunity to observe all the other hats. Then, you will all be instructed to simultaneously write down one of the colors. If any student correctly identifies the color of his or her own hat, then there will be no final exam this semester. Otherwise, I will assign a week-long haberdashery final."

What is the probability that the students have to take a final, assuming best play?

Proposed by John Hawksley (Massachusetts Institute of Technology '08) and Scott Kominers '09.

F07 – 3. Find all integer monic polynomials $f(x)$ such that

(i) $f(x) = f(1 - x)$ and

(ii) all complex zeros of f lie in the disk $|z| < \sqrt[5]{2}$.

Proposed by Vesselin Dimitrov '09.

F07 – 4. Let $a, b \geq 0$ be two nonnegative numbers. Find the limit

$$\lim_{n\to\infty} \sum_{k=1}^{n} \frac{1}{n+k+b+\sqrt{n^2+kn+a}}.$$

Proposed by Ovidiu Furdui (University of Toledo).

F07 – 5. For $i = 1, \ldots, n$, let $f_i : (\mathbb{Z}/m\mathbb{Z} \cup \{\star\})^n \to (\mathbb{Z}/m\mathbb{Z} \cup \{\star\})^n$ be given by

$$f_i(\vec{x}) = \begin{cases} (\star, x_2+1, x_3, \ldots, x_n) & i = 1 \text{ and } x_1 = 1, \\ (x_1, \ldots, x_{i-2}, x_{i-1}+1, \star, x_{i+1}+1, x_{i+2}, \ldots, x_n) & 1 < i < n \text{ and } x_i = 1, \\ (x_1, \ldots, x_{n-2}, x_{n-1}+1, \star) & i = n \text{ and } x_n = 1, \\ (x_1, \ldots, x_n) & \text{otherwise,} \end{cases}$$

where $\star + 1 = \star$ and $\vec{x} = (x_1, \ldots, x_n)$. Find necessary and sufficient conditions on $(x_1, \ldots, x_n) \in (\mathbb{Z}/m\mathbb{Z})^n$ such that there exists a sequence $\{i_k\}_{k=1}^n$ for which

$$f_{i_n}(\cdots(f_{i_1}(\vec{x}))) = (\star, \ldots, \star).$$

Proposed by Paul Kominers (Walt Whitman HS '08), Scott Kominers '09, and Zachary Abel '10.

The following two problems from the Spring 2007 issue received a total of one submission: a correction for S07 – 4 by Alon Amit (Google), for which we are most grateful. Since these problems defied solution, we are re-releasing them for one more issue. Their solutions will appear in Spring 2008.

S07 – 3. The incircle Ω_{ABC} of a triangle ABC is tangent to BC, CA, AB at P, Q, R respectively. Rays PQ and BA intersect at M, rays PR and CA intersect at N, and the incircle Ω_{MNP} of triangle MNP is tangent to MN and NP at X and Y respectively. Given that X, Y and B are collinear, prove:

(a) Circles Ω_{ABC} and Ω_{MNP} are congruent, and

(b) these circles intersect each other in 60° arcs.

Proposed by Zachary Abel '10.

S07 – 4 (Corrected)**.** For a prime p, let $\mathbb{Z}_{(p)} \subset \mathbb{Q}$ denote the localization of the integral domain $\mathbb{Z}$ at the prime ideal (p); that is, the subring of $\mathbb{Q}$ consisting of the rational numbers with denominators prime to p. The canonical homomorphism $\mathbb{Z} \to \mathbb{F}_p$ induces a canonical homomorphism $\phi_p : \mathbb{Z}_{(p)} \to \mathbb{F}_p$, the reduction modulo p homomorphism with kernel the maximal ideal $p\mathbb{Z}_{(p)}$ of the local ring $\mathbb{Z}_{(p)}$. For example, $\phi_5(1/2) = 3 \in \mathbb{F}_5$.

Let V be the set of primes p for which $\{\frac{3^n-1}{2^n-1} \mid n \in \mathbb{N}\} \subset \mathbb{Z}_{(p)}$.

(a) Characterize the set V.

(b) Let P be the set of primes, and define the set $W \subset P$ of **Wieferich primes** to be the set of primes p such that $p^2 \mid 2^{p-1} - 1$. It has been conjectured that, as x tends to infinity, the size of $\{p \in W \mid p \leq x\}$ is $O(\log\log x)$.

Show that V and $P \setminus V$ are both infinite sets, assuming the above conjecture for the former.

(c) Show that, for every $p \in V$, the map $\mathbb{N} \to \mathbb{F}_p$ given by $n \mapsto \phi_p\big((3^n-1)/(2^n-1)\big)$ is periodic.

For example, $5 \in V$, and the corresponding map $\mathbb{N} \to \mathbb{F}_5$ is $2, 1, 3, 2, 2, 1, 3, 2, 2, 1, 3, 2, \ldots$.

Proposed by Vesselin Dimitrov '09.

FEATURE

12 Solutions

How to Chop a Hyperbox

S07 – 1. How many hyperplane cuts are necessary to divide a $3 \times 5 \times 7 \times 9 \times 11$ rectangular solid into $3 \cdot 5 \cdot 7 \cdot 9 \cdot 11$ distinct $1 \times 1 \times 1 \times 1 \times 1$ hypercubes, if previously separated pieces can be rearranged between cuts?

Proposed by Joel Lewis '07.

Solution by Alon Amit (Google). Each cut can, at best, double the number of solid pieces, so an obvious lower bound is $\lceil \log_2(V) \rceil$, where V is the volume of the rectangular solid (henceforth "the box"). However, edges of odd length cannot be efficiently halved, so we are led to the following:

Proposition 3. *Let V be a box of dimensions $(a_1, a_2, \ldots a_n)$. The box can be cut into unit cubes using $L(V)$ cuts, and no fewer, where $L(V)$ is*

$$L(V) = L(a_1, \ldots, a_n) = \sum_{i=1}^{n} \lceil \log_2(a_i) \rceil.$$

Proof. For a fixed dimension $n \geq 1$, we prove this by induction on the value of L. We have $L = 0$ if and only if the box is a unit cube to begin with, so the claim holds in this case. We now assume the claim holds for all boxes with L-value less than L, and prove it for a box V with $L(V) = L \geq 1$. We need to show that V can be fully chopped with the advertised number of cuts and that any chopping procedure requires at least that many cuts.

Since $L \geq 1$, V has at least one edge whose length a_i is greater than 1. We cut the box across this edge, as close to the middle as possible. Namely, letting $b_i = \lceil a_i/2 \rceil$, we cut V into two boxes $V_1 = (a_1, \ldots, b_i, \ldots a_n)$ and $V_2 = (a_1, \ldots, a_i - b_i, \ldots, a_n)$.

Note that $L(V_1) = L(V) - 1$ and $L(V_2) \leq L(V_1)$. By the inductive hypothesis, V_1 can be chopped with $L(V_1)$ cuts. Moreover, V_2 can be chopped with that same number of cuts (or less), and these can be performed simultaneously with those of V_1: simply rearrange the pieces of V_2 that need to be cut at each stage along the same hyperplane used for cutting V_1. We are thus able to fully chop both V_1 and V_2 with $L(V) - 1$ cuts. Together with the initial cut, then, we have cut V in $L(V)$ cuts.

Furthermore, *any* procedure for fully chopping V must start with a single cut creating two pieces, each identical to V in all dimensions save one, and the larger of which is at least half as large as V. It follows that the first cut creates a piece W with $L(W) \geq L(V) - 1$. By the inductive hypothesis, W cannot be fully chopped with less than $L(V) - 1$ cuts, so the original box V requires at least $L(V)$ cuts. □

It is interesting to note that the proof shows a bit more than claimed: to optimally chop a box, all one needs to do is choose a non-trivial edge in each piece currently on hand and simultaneously cut them all near or at their middle. No further cleverness is required in choosing the sides or the cut locations.

For the box in the original problem, the number of cuts required is

$$L(3, 5, 7, 9, 11) = 2 + 3 + 3 + 4 + 4 = 16.$$ □

Also solved by Sergey Ioffe (Google), and the proposer.

πs in Odd Places

S07 – 2. Suppose $f : [0, 1] \to \mathbb{R}$ is an integrable function such that $f(x)y + f(y)x \le x^2 + y^2$. Show that $\int_0^1 f(x)\,dx \le \frac{\pi}{4}$. (One example of such a function is $f(x) = x$.)

Proposed by Scott Kominers '09.

Solution by Garret Dan Vo (Montana State University, Bozeman '10). Integrating the given inequality $x \cdot f(y) + y \cdot f(x) \le x^2 + y^2$ over the unit square $(x, y) \in [0, 1] \times [0, 1]$ gives the following stronger bound (by Fubini's Theorem):

$$\int_0^1 f(x)\,dx = \frac{1}{2}\int_0^1 f(y)\,dy + \frac{1}{2}\int_0^1 f(x)\,dx = \int_0^1\int_0^1 \big(x \cdot f(y) + y \cdot f(x)\big)\,dx\,dy$$
$$\le \int_0^1\int_0^1 (x^2 + y^2)\,dx\,dy = \frac{2}{3} < \frac{\pi}{4}. \qquad \square$$

Solution by Noam D. Elkies (Harvard University). Setting $x = y$ in the original constraint gives

$$x \cdot f(x) \le x^2,$$

whence $f(x) \le x$ for $x > 0$. Thus, we have for any such f that

$$\int_0^1 f(x)\,dx \le \int_0^1 x\,dx = \frac{1}{2} < \frac{\pi}{4}. \qquad \square$$

Solution by the proposer. Making the substitution $(x, y) = (\cos t, \sin t)$ for $t \in [0, \frac{\pi}{2}]$ reduces the constraint to $\sin t \cdot f(\cos t) + \cos t \cdot f(\sin t) \le 1$. Integrating this gives the desired bound:

$$\int_0^1 f(x)\,dx = \frac{1}{2}\int_0^{\frac{\pi}{2}} \big(\sin t \cdot f(\cos t) + \cos t \cdot f(\sin t)\big)\,dt \le \frac{1}{2}\int_0^{\frac{\pi}{2}} dt = \frac{\pi}{4}. \qquad \square$$

Also solved by Sherry Gong '11, John Hawksley (Massachusetts Institute of Technology '08), Sergey Ioffe (Google), Daniel Litt '10, Greg Price '06–'07, The Northwestern University Math Problem Solving Group, Manuel Silva (New University of Lisbon), and Arnav Tripathy '11.

Yet Another Mean Inequality

S07 – 3. (a) Prove that for distinct positive real numbers a and b, the following inequality holds:

$$\frac{a+b}{2} \ge \frac{a^{\frac{a}{a-b}} b^{\frac{b}{b-a}}}{e} \ge \frac{a-b}{\ln a - \ln b}.$$

(b*) Show that both inequalities are strict.

Proposed by Shrenik Shah '09.

Solution to parts (a) and (b) by Greg Price '06–'07. Consider the function $f : (-1, 1) \to \mathbb{R}$, $f(x) = \frac{1}{2x} \ln \frac{1+x}{1-x}$, with $f(0) = 1$. Observe that f is analytic and that it is given by the power series

$$f(x) = \frac{1}{2x}(\ln(1 + x) - \ln(1 - x)) = 1 + \frac{x^2}{3} + \frac{x^4}{5} + \cdots = \sum_{n=0}^{\infty} \frac{x^{2n}}{2n+1}$$

on the entire interval $(-1, 1)$.

We will need three facts about f. First, $f(x) \geq 1$, with equality only at $x = 0$; this follows immediately from the power series. Second, $f(x) \leq (1 - x^2)^{-1/2}$, with equality only at $x = 0$; this follows again from the power series, as the right-hand side is given by

$$1 + \frac{1}{2}x^2 + \frac{1}{2} \cdot \frac{3}{4}x^4 + \cdots = \sum_{n=0}^{\infty} \frac{(2n)!}{(2^n n!)^2} x^{2n}$$

on the whole interval $(-1, 1)$, and for all $n > 0$ we have $\frac{1}{2n+1} < \frac{(2n)!}{(2^n n!)^2}$ by a simple induction. Third, $f'(x) = \frac{1}{x(1-x^2)} - \frac{f(x)}{x}$; this follows from an elementary differentiation of $f(x) = \frac{1}{2x} \ln \frac{1+x}{1-x}$.

Now, given distinct positive reals a, b, we wish to prove

$$\frac{a+b}{2} > \frac{a^{\frac{a}{a-b}} b^{\frac{b}{b-a}}}{e} > \frac{a-b}{\ln a - \ln b}.$$

Let $x = \frac{a-b}{a+b}$; dividing through by $\frac{a+b}{2}$, the desired inequalities become

$$1 > (1 - x^2)^{1/2} e^{f(x)-1} > \frac{1}{f(x)},$$

which we now wish to prove for nonzero x in $(-1, 1)$. Since the three expressions are always positive, we may pass to their logarithms; since they are invariant under $x \mapsto -x$ we may consider only $x > 0$; since they are equal at $x = 0$, it will suffice to show that their respective logarithmic derivatives obey the same inequalities. So we wish to show for $x \in (0, 1)$ that $0 > -\frac{x}{1-x^2} + f'(x) > -\frac{f'(x)}{f(x)}$, or equivalently, multiplying by x, subtracting from 1, and employing our computation of the derivative f', that

$$1 < f(x) < \frac{1}{f(x)(1 - x^2)}.$$

But the left inequality follows from our first fact above, and the right from our second. We have proved the desired inequalities. □

Editor's note. The proposer noted that the inequality in the problem was obtained by taking the limit of the AM-GM-HM inequality for an arithmetic sequence of n terms from a to b as n goes to infinity. Paolo Perfetti (Università degli Studi Di Roma "Tor Vergata") pointed out that part (a) is a direct consequence of problem E3142 in *The American Mathematical Monthly* **95**#3, proposed by Zhang Zaiming. The solution provided there, by Ricardo Perez Marco, additionally proves part (b).

Also solved by Vishal Lama (Southern Utah University) and Paolo Perfetti (Università degli studi di Roma "Tor Vergata"). The proposer solved part (a), only. Partial solutions to part (a) were provided by Avery Carr (University of Memphis) and The Northwestern University Math Problem Solving Group.

FEATURE

13

ENDPAPER

Being a Mathematician

Prof. Véronique Godin[†]
Harvard University
Cambridge, MA 02138
godin@math.harvard.edu

My name is Véronique and I am a mathematician. I have probably been one for years now, but it took me a long time to acknowledge it.

As an undergraduate, I always thought of mathematicians as these weird people who roamed around the department. There was this very polite professor who would always apologize to the trash can whenever he ran into it. And also the professor who stayed seated throughout his lectures. Of course this meant that he could only write on the half-circle of the blackboard that was accessible from that chair.

As I moved to graduate school, mathematicians were my eccentric friends. There was the graduate student who thought that having a girlfriend and having a car were equivalent. Apparently the type of car you drove completely determined the type of women that you ended up with. (He was very depressed when he had to settle for an old Cadillac.)

There was also the graduate student who would play eBay to win. He was very proud of the fact that he had never lost an auction. He would proudly talk about his perfect record to anyone who would listen.

Mathematicians also included the foreign student who once took me out for a drive in his car. When I pointed out to him that he was running low on gas, he simply told me, "I don't know how to put gas in my car. Whenever I run out, someone helps me." And then he explained, "I know the theory behind putting gas in my car, but the details elude me."

But now I am a mathematician. I came to this conclusion recently on my way back to the US. The American custom agent asked me if I had my passport and I said, "Of course I do!" He waited. I waited. We waited some more. Finally he asked me, "Can I please see your passport?"

And someday, you will also be real mathematicians. And the world will not understand.

[†]Prof. Véronique Godin is a faculty member of the Harvard Mathematics Department.