David Merrell—Probabilistic Verification Exploration

dmerrell@cs.wisc.edu

Rational Householder Reflections

Background

FairSquare computes integrals by iteratively decomposing the region of integration into hyperrectangles that are aligned with the axes. We have been trying to improve this scheme by permitting hyperrectangles that are rotated with respect to the axes.

One of the challenges in this work is posed by the prevalence of irrational entries in rotation matrices. *FairSquare* uses the *Z3* SMT solver to generate hyperrectangles, and *Z3* operates best over the rational numbers. When irrational numbers are given to *Z3*, it rounds them to rational numbers. This is problematic, as it (a) introduces numerical error, undermining the proof; and (b) yields numbers with many digits, which causes *Z3* to become very slow.

In these notes, we present a way to find reflection matrices that have rational entries. The method yields a rich set of reflections, enabling any desired reflection to be approximated without undermining *FairSquare*'s correctness.

Householder Matrices

Householder matrices are a class of reflection matrices of the following form:

$$H = I - 2\hat{n}\hat{n}^{\mathrm{T}}$$

where $\hat{n}$ is a unit vector. Intuitively, a Householder matrix reflects a vector across the hyperplane whose normal vector is $\hat{n}$.

Since a Householder matrix is a reflection, its determinant is -1. Hence, the volume of a region transformed by a Householder matrix is the same as the volume of the original region. Householder matrices also have the interesting property of being their own inverses.

Using a Householder Matrix

Suppose we are computing the volume of a region in $\mathbb{R}^d$. Ordinarily, *FairSquare* would compute this volume by generating hyperrectangles that are aligned with the axes. However, suppose we wish to generate hyperrectangles that are aligned with a certain hyperplane, defined by normal vector v. As discussed elsewhere, this can be done by either generating hyperrectangles that are aligned as desired, or by transforming the region of integration such that v is in alignment with one of the axes. We adopt the latter convention in the work that follows.

We construct a Householder matrix that reflects the vector v into alignment with the standard basis vector $\hat{e}_k$. What we seek is a reflection about the hyperplane normal to $\frac{v}{\|v\|} - \hat{e}_k$. The corresponding Householder matrix is given by

$$H = I - 2\frac{\left(\frac{v}{\|v\|} - \hat{e}_k\right)\left(\frac{v}{\|v\|} - \hat{e}_k\right)^{\mathrm{T}}}{\|\frac{v}{\|v\|} - \hat{e}_k\|^2}$$

$$= I - 2\frac{(v - \|v\|\hat{e}_k)\,(v - \|v\|\hat{e}_k)^{\mathrm{T}}}{(v - \|v\|\hat{e}_k)^{\mathrm{T}}\,(v - \|v\|\hat{e}_k)}$$

It turns out that the entries of H are given by the following:

$$H_{i,j} = \frac{1}{\|v\|} \cdot \begin{cases} v_k, & i = j = k \\ \|v\| - \frac{v_i^2}{\|v\| - v_k}, & i = j \neq k \\ v_i, & i \neq k, j = k \\ v_j, & i = k, j \neq k \\ -v_i v_j, & i \neq k, j \neq k, i \neq j \end{cases}$$

At this point we are equipped to ask the question: under what circumstances will H have rational entries? The answer is that H will have rational entries whenever v has rational entries, and its norm $\|v\|$ is also rational.

This condition is satisfied if we require the d entries of v, along with its norm, to form a Pythagorean $(d+1)$-tuple.

Generating Pythagorean n-Tuples

Integers $a_1, \ldots, a_n$ form a Pythagorean n-tuple if they satisfy the generalized Pythagorean formula:

$$a_1^2 + \cdots a_{n-1}^2 = a_n^2.$$

These can be generated from integers $b_1, \ldots, b_{n-1}$ as follows:

$$\begin{aligned} a_1 &= b_1^2 - (b_2^2 + b_3^2 + \cdots + b_{n-1}^2) \\ a_i &= 2b_1 b_i \qquad\qquad \forall i \in \{2, \ldots, n-1\} \\ a_n &= b_1^2 + b_2^2 + \cdots + b_{n-1}^2 \end{aligned}$$

It can be easily verified that these satisfy the generalized Pythagorean formula.

In practice, this scheme can be used to generate a rich set of Pythagorean n-tuples. The number of unique n-tuples generated for varying values of n are tabulated below; "limit" refers to a limit placed on the size of any integer in the tuples.

Unique Pythagorean n-Tuples Generated			
	limit: 100	limit: 1,000	limit: 10,000
$n = 3$	15	113	1,061
$n = 4$	32	939	28,305
$n = 5$	30	2,741	250,505

Note that each of these tuples are truly unique; they have been sorted, consist only of nonnegative integers, and have been divided through by their greatest common divisors (i.e. vectors in the same direction but of differing length have been removed). By permitting negative integers, the numbers of valid tuples would be multiplied by 2^n. Furthermore, allowing entries of the tuples corresponding to the LHS of the generalized Pythagorean formula would multiply the numbers of valid tuples by $(n-1)!$.

So there is no shortage of Pythagorean n-tuples whose entries have a reasonable number of digits.

Using Pythagorean n-Tuples

With a rich set of Pythagorean tuples at our disposal, I propose the following:

- Given a vector v that we wish to align with the axes in $\mathbb{R}^d$, we replace it with a rational vector r_v that best approximates it, drawing from those vectors generated as Pythagorean $(d+1)$-tuples. There are probably smart ways to do this with spatial data structures.

- We construct a Householder matrix from r_v; the result is is *exactly* a reflection matrix, and will have rational entries with reasonable numbers of digits.

- We reflect each of the atoms (defining the region of integration) by this householder matrix, and then proceed to compute the volume of the transformed formula (region).

- Ideally, we will enjoy benefits from transforming the region of integration without needing to worry about numerical issues. We will be able to move on to bigger questions, like how to *select* the reflections we want (optimization problem).