Degenerate Diffusions in Population Genetics

IPAM Workshop on Mathematical Oncology

Charles L. Epstein

CCM

February 24, 2026

FLATIRON INSTITUTE
Center for Computational Mathematics

Introduction

I want to thank the organizers for inviting me to speak in meeting.

The work I will describe was done jointly with Rafe Mazzeo, and Jon Wilkening.

FLATIRON INSTITUTE
Center for Computational Mathematics

Outline

FLATIRON INSTITUTE
Center for Computational Mathematics

Population Genetics

In population genetics we study how the distribution of variants in a reproducing population evolves over time. There are typically four important effects:

1. The randomness in the number of offspring a given individual, or pair, has in a given generation.
2. Mutation from one type to another type.
3. Differences in "fitness" among the different types.
4. Migration in and out of a given environment.

FLATIRON INSTITUTE
Center for Computational Mathematics

Finite populations

In this model we assume a fixed population size N, with each individual of either the ancestral type A or the derived type a. There are 2 alleles.
In the simplest case the two types have the same fitness and there is no mutation. Since $n_a + n_A = N$, it is enough to keep track of $X^{(N)}(t)$, the number of type a in generation t. This is a Markov chain model, so we need the conditional probabilities

$$P(X^{(N)}(t+1) = j | X^{(N)}(t) = i).$$

In the standard Wright-Fisher model we assume that the mating is "random:" (an urn with replacement) the number of a in the next generation is given by the binomial sampling formula:

$$\text{Prob}(X^{(N)}(t+1) = j | X^{(N)}(t) = i) = \binom{N}{j} \left(\frac{i}{N}\right)^j \left(1 - \frac{i}{N}\right)^{N-j} \tag{1}$$

FLATIRON INSTITUTE
Center for Computational Mathematics

Mutation and Selection

To incorporate mutation and selection, we change the odds. If a and A have relative fitness $(1+s):1$, and the rate at which $a \to A$ is μ_1 and the rate at which $A \to a$ is μ_2, then we let:

$$p_i = \frac{i(1+s)(1-\mu_1)}{i(1+s)+N-i} + \frac{(N-i)\mu_2}{i(1+s)+N-i},$$

and alter the transition matrix to be

$$\text{Prob}(X^{(N)}(t+1) = j | X^{(N)}(t) = i) = \binom{N}{j} (p_i)^j (1-p_i)^{N-j} \tag{2}$$

There are many generalizations with more types and multiple populations with migration etc., but today we focus on the simplest case.

FLATIRON INSTITUTE
Center for Computational Mathematics

Kimura Diffusion Operator, I

The main topic of this talk concerns the operators that arise as limits of these sorts of processes as the population N tends to infinity. In the 1d-case, the rescaled processes

$$X_N(t) = \frac{1}{N} X^{(N)}([tN]), \tag{3}$$

converge to a continuous time stochastic process, $X(t)$, parametrized on the interval $[0, 1]$. The "backward" Kolmogorov operator is the second order differential operator:

$$Lf(x) = \frac{x(1-x)}{2}\partial_x^2 f + \sigma x(1-x)\partial_x f + m_2(1-x)\partial_x f - m_1 x \partial_x f. \tag{4}$$

Where $\sigma = Ns$, $m_1 = N\mu_1$ and $m_2 = N\mu_2$ are assumed fixed, as $N \to \infty$; this is not a biologically meaningful assumption, it just allows us to define the limiting process. In applications we choose a value for N; it is not the (often huge) census population size, but rather an "effective" population size.

Kimura Diffusion Operator, II

The forward operator is the formal adjoint

$$L^t f(x) = \partial_x^2 \left(\frac{x(1-x)}{2} f \right) - \sigma \partial_x [x(1-x)f] - m_2 \partial_x [(1-x)f] + m_1 \partial_x [xf]. \quad (5)$$

I call these Kimura diffusion operators.

If $p_t(x,y)$ denotes the fundamental solution, so that

$$\partial_t p_t(x,y) = L_x p_t(x,t) = L_y^t p_t(x,y), \text{ with } \lim_{t \to 0^+} p_t(x,y) = \delta(x-y). \quad (6)$$

For an allele at initial frequency $X(0) = y$, we have

$$\text{Prob}(a < X(t) < b) = \int_a^b p_t(x,y) dx. \quad (7)$$

FLATIRON INSTITUTE
Center for Computational Mathematics

The Different Terms

In this limit the second order term, $\frac{x(1-x)}{2}\partial_x^2 f$, is related to the randomness in the number of offspring; whereas mutation and selection become deterministic forces, represented by the vector field,

$$\sigma x(1-x)\partial_x f + m_2(1-x)\partial_x f - m_1 x\partial_x f = b(x)\partial_x f.$$

What makes it difficult to study this operator is the fact that the coefficient of the second order term vanishes at the boundary of $[0,1]$.

The paths of the process satisfy the SDE:

$$dX = \sqrt{2X(1-X)}dW + b(X)dt. \tag{8}$$

The linear vanishing of the coefficient of ∂_x^2 implies that, if $b(0) = 0$, or $b(1) = 0$, then paths reach $\{0\}$, or $\{1\}$ in finite time, with probability 1. The conditions $b(0) \geq 0$, $b(1) \leq 0$, prevent paths from leaving $[0,1]$.

FLATIRON INSTITUTE
Center for Computational Mathematics

Outline

FLATIRON INSTITUTE
Center for Computational Mathematics

The Forward Equation

The basic application of the forward equation is to understand how a distribution of variants in an ensemble of populations evolves over time. The forward equation naturally acts on Borel measures on $[0,1]$. In fact the domain on which the semi-group is strongly continuous as $t \to 0$, consists absolutely continuous measures, with the possible addition of δ-measures at $0, 1$.
Given $w_0(x) = h(x)dx$, then the solution to $\partial_t w(x,t) = L^t(w(x,t))$, describes how the mass of this measure is redistributed. If $\alpha\beta \neq 0$, then there is a unique stationary measure that is a solution to $L^t(w_{\text{stat}}) = 0$; all solutions converge to this measure as $t \to \infty$.
If one of α, or β vanishes then the null-space of L^t consists of $\delta(x)dx$, or $\delta(1-x)dx$. If both vanish, then both δ-measures are in the null-space of L^t. These are absorbing states, and all initial populations reach one or the other state in finite time, with probability 1.

FLATIRON INSTITUTE
Center for Computational Mathematics

The Allele Frequency Spectrum

Using these models one can estimate statistics of quantities that are observable in real applications, or infer quantities that are not directly observable. Of particular interest is the allele or site frequency spectrum.

Suppose that one has sequenced n chromosomes, from a much larger population, and found k sites where there is variation. It is assumed that a site is either unmutated or has exactly 2 alleles (infinite sites model). Further mutation and back-mutation are assumed not to occure.
The AFS is then the sequence of numbers $\{f_1, \ldots, f_{n-1}\}$, where f_j is the number of sites where an allele appears in j samples. These numbers are between 1 and k and are often represented as frequencies; in the large n limit, these frequencies satisfy a Kimura diffusion equation.

Their statistics are affected by changes in the population size, mutation rates, selection, etc. Using various probabilistic models their values can be predicted, which can in turn be used to predict the likelihood of observations under various assumptions.

FLATIRON INSTITUTE
Center for Computational Mathematics

The Allele Frequency Spectrum Example

	SNP 1	SNP 2	SNP 3	SNP 4	SNP 5	SNP 6	SNP 7	SNP 8
Sample 1	0	1	0	0	0	0	1	0
Sample 2	1	0	1	0	0	0	1	0
Sample 3	0	1	1	0	0	1	0	0
Sample 4	0	0	0	0	1	0	1	1
Sample 5	0	0	1	0	0	0	1	0
Sample 6	0	0	0	1	0	1	1	0
Total	1	2	3	1	1	2	5	1

From:

https://en.wikipedia.org/wiki/Allele_frequency_spectrum

In this example there are 6 samples and 8 "segregating" sites, loci with variations. The AFS is therefore

$$(f_1, f_2, f_3, f_4, f_5) = (4, 2, 1, 0, 1) \tag{9}$$

FLATIRON INSTITUTE
Center for Computational Mathematics

Usage of Diffusion Models

Many problems in population genetics have been addressed using these diffusion models by Kimura, Crow, Ohta, Ewens, and more recently Patterson, Evens, Shvets, etc. Most of these applications have used explicit solutions and clever probabilistic arguments.

The usage of the forward model (as opposed to the coalescent, which I won't define) seems to have fallen out of favor, and that seems to be connected to the difficulty of finding accurate numerical solutions to the equation $\partial_t u = Lu$.

I will now explain the sources of this difficulty and outline approaches Wilkening and I have introduced to overcome them.

FLATIRON INSTITUTE
Center for Computational Mathematics

Outline

FLATIRON INSTITUTE
Center for Computational Mathematics

Boundary Conditions

The fact that $x(1-x)$, the coefficient of the ∂_x^2-term, vanishes at the boundaries of the interval has a significant impact on the types of boundary conditions that can be specified. This problem was carefully addressed by Feller in a classic 1952 paper, [1], where he found all closed operators on $\mathcal{C}^0([0,1])$ with formal symbol L that generate positivity preserving semi-groups. In fact he considers the operators

$$Lu = x(1-x)\partial_x^2 u + [\alpha x - \beta(1-x) + x(1-x)s(x)]\partial_x u + c(x)u, \qquad (10)$$

where $s(x), c(x)$ are bounded and continuous. The constants α, β are non-negative, which means the vector field, $[\alpha x - \beta(1-x) + x(1-x)s(x)]\partial_x$, is “inward" pointing. This is important for applications as it means that paths of process stay in $[0,1]$ with probability 1.
We won’t go through this theory, but consider some useful, special cases.

FLATIRON INSTITUTE
Center for Computational Mathematics

The Graph Closure

The simplest choice is the $\mathcal{C}^0$-graph closure of the formal operator L acting on $\mathcal{C}^2([0,1])$: the domain of $\overline{L}$, is the closure of $\mathcal{C}^2([0,1])$ in the topology defined by the graph norm

$$\|f\|_{G(L)} = \|f\|_{\mathcal{C}^0([0,1])} + \|Lf\|_{\mathcal{C}^0([0,1])}.$$

The operator $\overline{L}$ defines the strongly continuous, positivity preserving semi-group on $\mathcal{C}^0([0,1])$, which is the most useful for applications.

The domain of this operator, $\mathrm{Dom}(\overline{L}) \subset \mathcal{C}^0([0,1]) \cap \mathcal{C}^2((0,1))$, is described by regularity conditions, which depend on α and β.
In all cases $x(1-x)\partial_x^2 f(x)$ extends continuously to 0 and 1, where it vanishes.

If α (or β) is positive, then $f \in \mathrm{Dom}(\overline{L})$, provided $\lim_{x\to 0+} \partial_x f(x)$ exists (or $\lim_{x\to 1-} \partial_x f(x)$ exists). We call the solution to the diffusion equation

$$\partial_t u(x,t) = Lu(x,t), \quad u(x,0) = u_0(x) \in \mathcal{C}^0([0,1]), \tag{11}$$

given by this semigroup the *regular solution*. In the population genetics literature this is often called the "no-flux" solution.

The Regular Solution

The key thing to understand is that the solution to $\partial_t u = Lu$ with $u(x,0) = f(x)$ is *uniquely determined* by its **regularity** at the boundary. This indicates why it is not really possible to accurately solve these problems using a finite difference method: **the boundary condition is not a pointwise condition**.
For the regular solution we cannot specify either Dirichlet or Neumann data at $\partial[0,1]$. Unlike the case of a non-degenerate operator (e.g. ∂_x^2), regularity alone implies uniqueness.

The regular solution has the following remarkable property: If $f \in \mathcal{C}^m([0,1])$, then the regular solution $u \in \mathcal{C}^m([0,1] \times [0,\infty)) \cap \mathcal{C}^\infty([0,1] \times (0,\infty))$.

The conditions specified at the two endpoints are largely independent.

FLATIRON INSTITUTE
Center for Computational Mathematics

Dirichlet Conditions

If $0 \leq \alpha < 1$ (or $0 \leq \beta < 1$), then we can also define a closed semigroup with the homogeneous Dirichlet boundary conditions

$$\lim_{x \to 0^+} u(x,t) = 0, \text{ (or } \lim_{x \to 1^-} u(x,t) = 0).$$

A solution of $\partial_t u = Lu$ satisfying this condition necessarily has a singularity at the boundary. If $0 < \alpha < 1$, then $u(x,t) \sim x^{1-\alpha}$, (or $u(x,t) \sim (1-x)^{1-\beta}$).

If $\alpha = 0$, then we typically get an $x \log x$ singularity (or if $\beta = 0$ an $(1-x)\log(1-x)$ singularity).
If we want to impose inhomogeneous Dirichlet conditions:
$u(0,t) = h_0(t), u(1,t) = h_1(t)$, then we'll write the solution as

$$u(x,t) = (1-x)h_0(t) + xh_1(t) + u_0(x,t),$$

where $u_0(x,t)$ satisfies an inhomogeneous diffusion equation with homogeneous Dirichlet conditions. This means that $u_0(x,t)$ will not be smooth at the boundary.

FLATIRON INSTITUTE
Center for Computational Mathematics

Motivating Example

Our work is motivated by an example found in the paper of Evans, Shvets and Slatkin, wherein Kimura diffusion is used to estimate quantities connected to the AFS with a varying population. It asks for a solution to the problem:

$$\partial_t u = \frac{1}{2\rho(t)} x(1-x)\partial_x^2 u - Sx(1-x)\partial_x u, \tag{12}$$

with

$$\lim_{t\to 0} u(x,t) = u_0(x) \text{ and } \lim_{x\to 0^+} u(x,t) = \theta\rho(t),\ \lim_{x\to 1^-} u(x,t) = 0. \tag{13}$$

The allele frequency spectrum is actually $f(x,t) = \frac{u(x,t)}{x(1-x)}$; θ is related to the per site mutation rate. This is challenging for a wide variety of reasons: the imposition of a Dirichlet condition forces the solution to be singular, and the population $\rho(t)$ can vary over a very large range for the time frame of interest. I will discuss numerical methods (motivated by analysis of this problem) that we hope will convince people to try using diffusion methods again.

Asymptotic Behavior

Let $L_{\alpha,\beta} = x(1-x)\partial_x^2 + [\alpha(1-x) - \beta x]\partial_x]$.

It is useful to consider how the asymptotic behavior of the regular solution depends on α, β. The operator $L_{0,0} + x(1-x)s(x)\partial_x$ has a 2-dimensional null-space spanned by $1, v_0(x)$, where $v_0(0) = 0$, and $\partial_x v_0(x) > 0$, for $x > 0$. For a regular solution $u_t(0,1) = u_t(1,t) = 0$, and therefore $u(x,t)$ converges to $a + bu_0(x)$, where $a = u_0(0)$, and $u_0(1) = a + bv_0(1)$.

If α and β are non-zero, then the operator $L_{\alpha,\beta} + x(1-x)s(x)\partial_x$ has a null-spaced spanned constant functions. The PDE can be re-written as $w(x)u_t = [x(1-x)w(x)u_x]_x$, where $w(x) = x^{\alpha-1}(1-x)^{\beta-1}e^{S(x)}$. Hence the limiting constant value equals

$$\frac{\int_0^1 w(x)u_0(x)dx}{\int_0^1 w(x)dx}. \tag{14}$$

Similar considerations apply if one of α or β in positive.

FLATIRON INSTITUTE
Center for Computational Mathematics

Outline

FLATIRON INSTITUTE
Center for Computational Mathematics

The Eigenfunction Approach

For many years it has been understood that the Kimura diffusion equation, with $L_{\alpha,\beta} = x(1-x)\partial_x^2 + (\alpha(1-x) - \beta x)\partial_x$ can be solved using eigenfunction expansions.

These operators preserve the polynomials of degree d, for every d and therefore have a basis of eigenfunctions $\{p_n^{\alpha,\beta}(x)\}$, that are polynomials, which are, in fact, classical Jacobi polynomials. If α, β are positive then $L_{\alpha,\beta} p_n^{\alpha,\beta} = -n(\alpha+\beta+n-1)p_n^{\alpha,\beta}$, where

$$p_n^{\alpha,\beta}(x) = c_{\alpha,\beta,n} P_n^{\alpha-1,\beta-1}(1-2x).$$

The cases where α or β vanish can be obtained by taking limits. In these cases the eigenfunctions vanish at the end point where α or β vanish.

FLATIRON INSTITUTE
Center for Computational Mathematics

Solving the PDE

The solution is given by the infinite sum

$$u(x,t) = \sum_{j=0}^{\infty} f_j e^{t\lambda_j(\alpha,\beta)} p_j^{\alpha,\beta}(x), \tag{15}$$

where the coefficients $\{f_j\}$ are determined by the initial data, $f(x)$. The smoother the initial data, the faster these coefficients vanish.
Finite partial sums give spectrally accurate approximations to $u(x,t)$, but need to be calculated carefully.

FLATIRON INSTITUTE
Center for Computational Mathematics

Difficulty of the Eigenfunction Approach

The principal difficulty is in evaluating the eigenfunctions as the degree grows. The stable and accurate way to do this employs the 3-term recurrence relations that these eigenfunctions always satisfy.

$$b_{n+1}(\alpha,\beta)p_{n+1}^{\alpha,\beta}(x) = (x - a_n(\alpha,\beta))p_n^{\alpha,\beta}(x) - b_n(\alpha,\beta)p_{n-1}^{\alpha,\beta}(x). \tag{16}$$

The coefficients, $\{(a_n(\alpha,\beta), b_n(\alpha,\beta))\}$, are readily computable.

Using the expressions for these polynomials in terms of monomials quickly leads to very inaccurate results due to large coefficients and the consequent catastrophic cancellations.
Eigenfunction expansions is a standard method, so we just show examples as to why one needs to be careful evaluating these eigenfunctions.

FLATIRON INSTITUTE
Center for Computational Mathematics

Catastrophic Cancellation

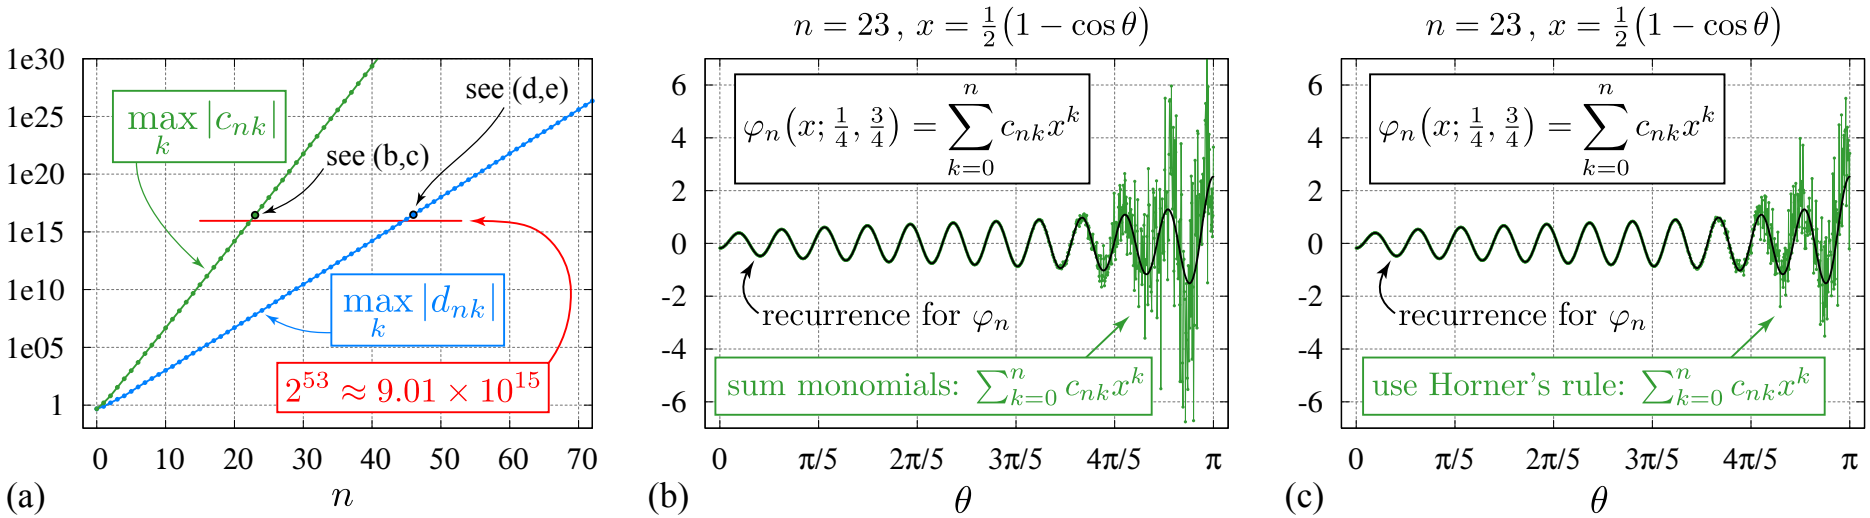

Whether evaluated directly or with Horner's rule, in floating point arithmetic, the monomial expansions suffer catastrophic cancellation of digits for x near 1 (see (b,c)). The black curves were evaluated via the recurrence above and remain accurate throughout the interval $0 \leq x \leq 1$. Things get much worse as n grows, though not for the 3-term recurrence.

FLATIRON INSTITUTE
Center for Computational Mathematics

Limitations of the Eigenfunction Approach

The eigenfunction approach is only directly applicable to the case $L = L_{\alpha,\beta}$.
If we are interested in solving more general equations:

$$\partial_t u = L_{\alpha,\beta} u + x(1-x)s(x)\partial_x u(x) + f(x,t), \quad u(x,0) = u_0(x), \tag{17}$$

then we need another method. It is important to solve the inhomogeneous problem as we are led to this if there is an inhomogeneous Dirichlet condition, as in the allele frequency example above.
The term, which we denote by $B = x(1-x)s(x)\partial_x$, is a lower order perturbation of $L_{\alpha,\beta}$, in that it does not change the domain of the regular operator.

FLATIRON INSTITUTE
Center for Computational Mathematics

Outline

FLATIRON INSTITUTE
Center for Computational Mathematics

Alternate Approaches

We introduce two methods. The first employs orthogonal polynomials defined by the weight function, $w(x) = x^{\alpha-1}(1-x)^{\beta-1}e^{S(x)}$, introduced above. If $s \neq 0$, then these are not eigenfunctions, but, by choosing sample points correctly, this discretization leads to an exact Galerkin approximation w.r.t. the polynomial bases. This is described in detail in [5], and I won't consider it further. Once again, it is important to use a 3-term recurrence to evaluate the polynomials.

FLATIRON INSTITUTE
Center for Computational Mathematics

Duhamel's formula

For our second approach we replace the PDE with an integral equation known as Duhamel's formula

$$u(x,t) = [e^{tL_{\alpha,\beta}} u_0](x) + \int_0^t e^{(t-s)L_{\alpha,\beta}} [B(s)u(\cdot,s) + f(\cdot,s)]ds. \qquad (18)$$

Here $e^{tL_{\alpha,\beta}}$ is the regular solution operator for $\partial_t - L_{\alpha,\beta}$. Note that here the perturbation term, $B(s)$, is a lower order operator that can depend on time.

Starting with this formula one is discretizing the solution operator for the PDE, rather than the PDE itself. This is inherently well conditioned, whereas, in finite difference, or Galerkin approximations, discreasing the mesh spacing, or increasing the order, **increases** the condition number of the linear systems that need to be solved.

FLATIRON INSTITUTE
Center for Computational Mathematics

Evaluating the Exponential

The main challenge is to find a way to approximate the kernel of $e^{tL_{\alpha,\beta}}$. A "simple" approach is to use the eigenfunctions to write

$$e^{tL_{\alpha,\beta}} = \sum_{j=0}^{N} e^{t\lambda_j(\alpha,\beta)} \varphi_j^{\alpha,\beta}(x) \varphi_j^{\alpha,\beta}(y), \tag{19}$$

where $\{\varphi_j^{\alpha,\beta}\}$ are appropriately normalized eigenfunctions of $L_{\alpha,\beta}$.
In principle, this is a rapidly convergent series, however, since we're only going 1 time step $(t-s)$ will remain small. To get accurate results requires a somewhat indirect approach.

FLATIRON INSTITUTE
Center for Computational Mathematics

The Time Stepper

We split the interval $[0,h]$ into subintervals with end-points $\{t_j = c_j h\}$, where the $\{c_j\}$ are Gauss-Lobatto quadrature nodes. Given $u_0(x)$ we seek functions $\{u_1(x),\ldots,u_\nu(x)\}$ so that, if we use polynomial interpolation in time to define $[B(s)u(s)+f(s)]$ for $s\in[0,h]$, then the resulting $u(x,c_r h) = u_r(x)$. This is equivalent to requiring

$$u_r = e^{c_r h L_{\alpha,\beta}} u_0 + h\sum_{j=0}^{\nu} a_{rj}(hA)[B(c_j h)u_j + f(c_j h)],\ \text{for } r = 1,\ldots,\nu, \tag{20}$$

where

$$a_{rj}(z) = \int_0^{c_r} e^{(c_r - s)z} l_j(s)\,ds, \tag{21}$$

for the Lagrange interpolating polynomials:

$$l_j(s) = \prod_{k\neq j} \frac{s - c_k}{c_j - c_k}. \tag{22}$$

FLATIRON INSTITUTE
Center for Computational Mathematics

The System of Equations

We can rewrite the system of equations as

$$\left[\begin{pmatrix} \mathrm{Id} & & \\ & \ddots & \\ & & \mathrm{Id} \end{pmatrix} - h \begin{pmatrix} a_{11}(hA) & \cdots & a_{1\nu}(hA) \\ \vdots & \ddots & \vdots \\ a_{\nu 1}(hA) & \cdots & a_{\nu\nu}(hA) \end{pmatrix} \times \right.$$
$$\left. \begin{pmatrix} B(c_1 h) & & \\ & \ddots & \\ & & B(c_\nu h) \end{pmatrix} \right] \begin{pmatrix} u_1 \\ \vdots \\ u_\nu \end{pmatrix} = \begin{pmatrix} F_0 \\ \vdots \\ F_\nu \end{pmatrix}, \quad (23)$$

or $[\mathbb{I} - h\mathbb{A}\mathbb{B}]\vec{u} = \vec{F}$, for short.

It is not difficult to show that $\|\mathbb{A}\mathbb{B}\| \leq C\,\mathrm{Leb}_\nu$, where Leb_ν is the Lebesgue number of the nodes $(c_0, \ldots, c_\nu)$. With Gauss-Lobatto nodes it is less than $\frac{2}{\pi}\ln\nu + 1$. The main work is in evaluating $\mathbb{A}$.

FLATIRON INSTITUTE
Center for Computational Mathematics

Constant Population Example

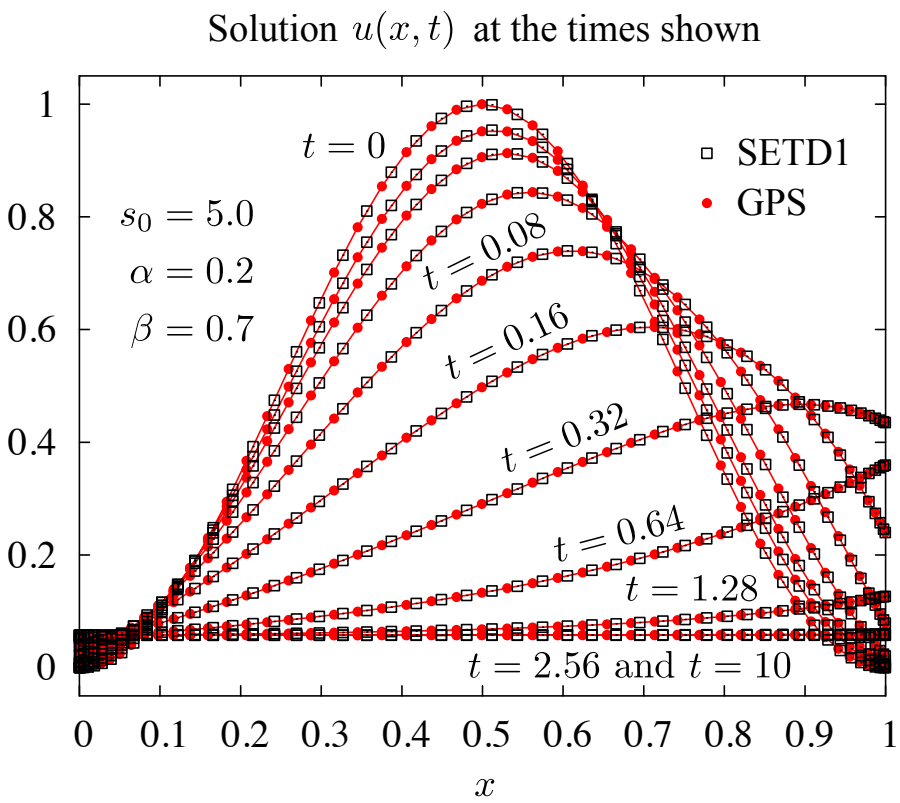

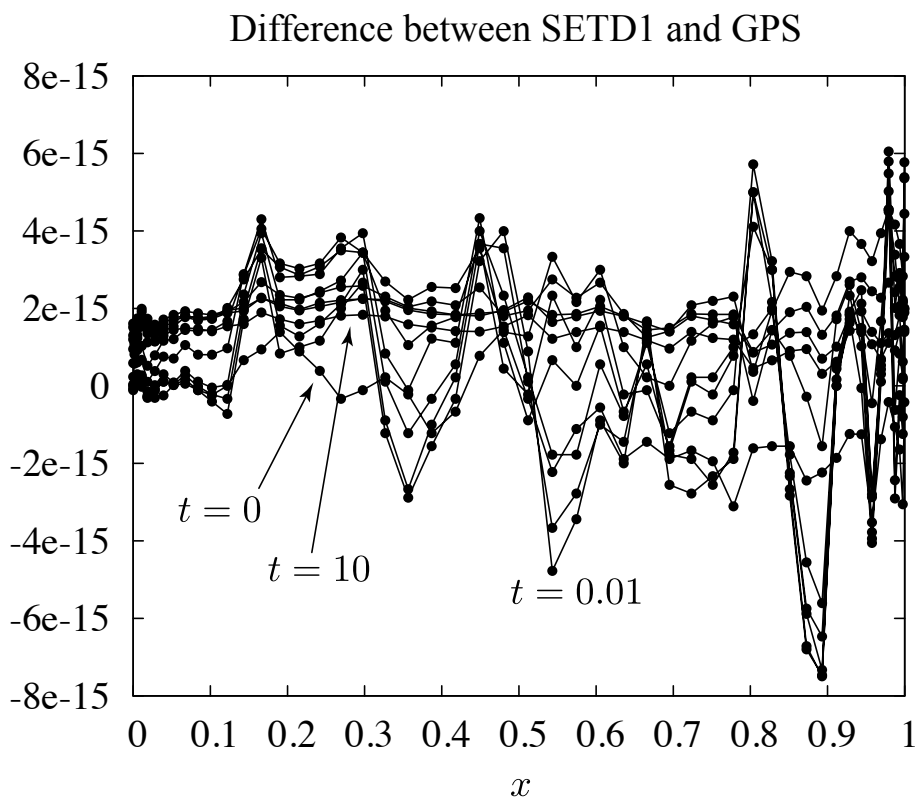

Solutions to $\partial_t u = L_{.2,.7}u - 5x(1-x)\partial_x u$ with $u_0(x) = \sin^2(\pi x)$. The left plot shows the solutions using two different methods of computing $e^{tL_{.2,.7}}$, the right plot shows the differences, which are quite small.

FLATIRON INSTITUTE
Center for Computational Mathematics

Time Varying Population Example

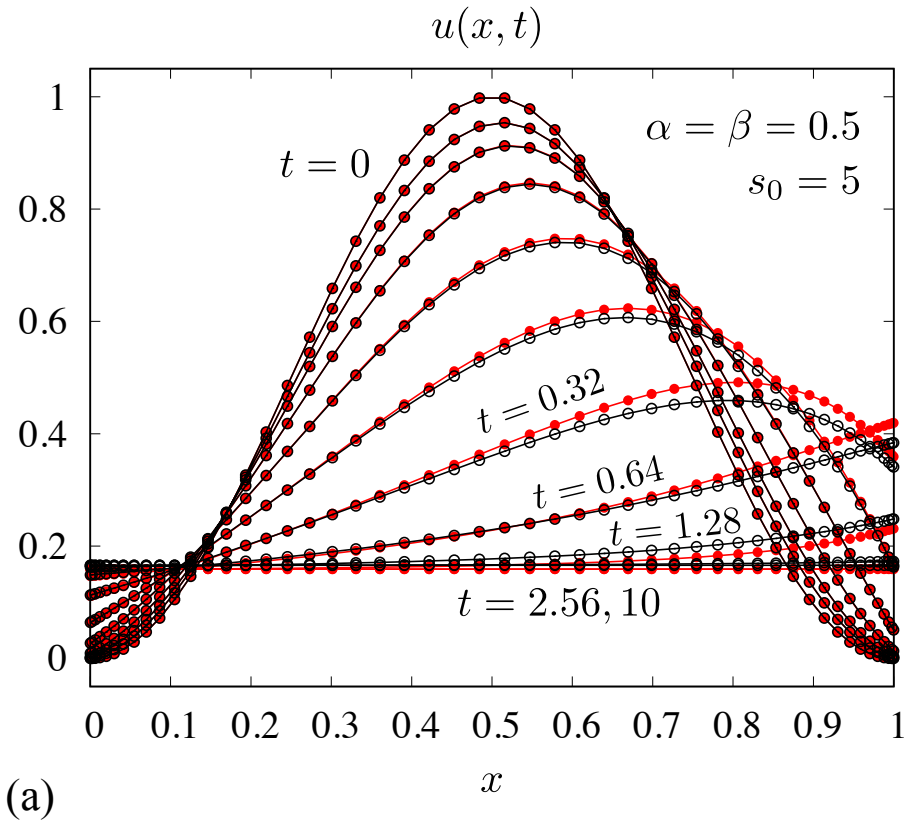

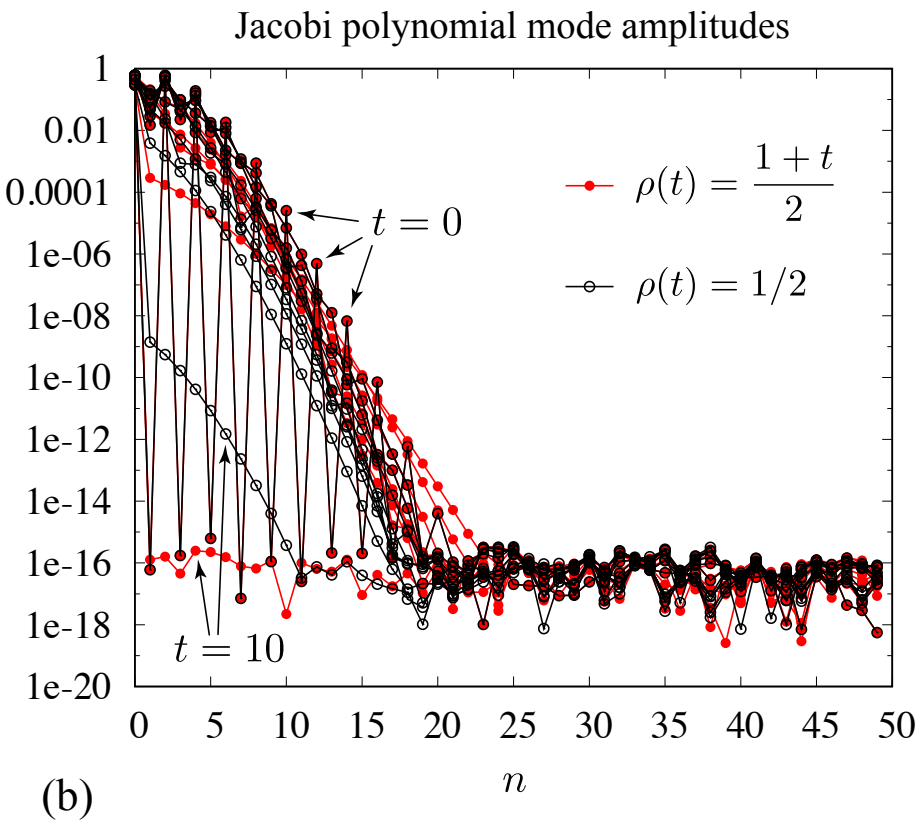

Open black markers are for $\rho(t) = 1/2$, red markers to $\rho(t) = (1+t)/2$.

Here we compare the results with a constant population and a time dependent population. With $\alpha = \beta = 0.5$, $s_0 = 5$, and $\rho(t) = (1+t)/2$, here we solve

$$u_t = \frac{x(1-x)}{2\rho(t)} u_{xx} + \alpha(1-x)u_x - \beta x u_x - s_0 x(1-x)u_x, \tag{24}$$

Outline

FLATIRON INSTITUTE
Center for Computational Mathematics

Application to the Allele Frequency Spectrum

At the beginning of the lecture we described an application of the Kimura Diffusion to study the evolution of the allele frequency spectrum with a varying population.
The problem we'd like to solve is

$$\partial_t u = \frac{x(1-x)}{\rho(t)}\partial_x^2 u, \text{ with } u(0,1) = \theta\rho(t),\ u(1,t) = 0, \tag{25}$$

and the initial condition $u(x,0) = \theta(1-x)$. Recall that the AFS is $u(x,t)/x(1-x)$.

The population model of particular interest has $\rho(t) = 1$, for $t \leq 0$, and $\rho(t) = e^{Rt}$ for $t > 0$, with $R = 40$. The solution to this equation is singular along the $x = 0$ boundary, with a leading singularity of the form $a(t)x\log x$. The methods introduced earlier that compute $e^{(t-s)L_{0,0,}}$ using eignefunction expansions are not able to adequately resolve this singularity.

FLATIRON INSTITUTE
Center for Computational Mathematics

Final Remarks

To resolve this singularity, we express the exponential $e^{(t-s)L_{0,0},}$ as a function of the resolvent $(\mathrm{Id} - hL_{0,0})^{-1}$. The resolvent can be explicitly expressed in terms of hypergeometric functions. This allows the usage of adaptive coordinates near $x = 0$.

Nonetheless, for sufficiently small h the dynamic range of the resolvent is $1 : 10^{13,000}$. My collaborator Jon Wilkening was undaunted by these extraordinary challenges, and was able to compute solutions to the problem above out to $t = .3$, where $\rho(t) = 162,754$.

I don't have time to give any of the details of this analysis, but will show the sort of results Jon obtains. Let me know if you'd like to know more about how this is actually done. Hopefully software will be freely available online so anyone can do it.

FLATIRON INSTITUTE
Center for Computational Mathematics

A Computational Example

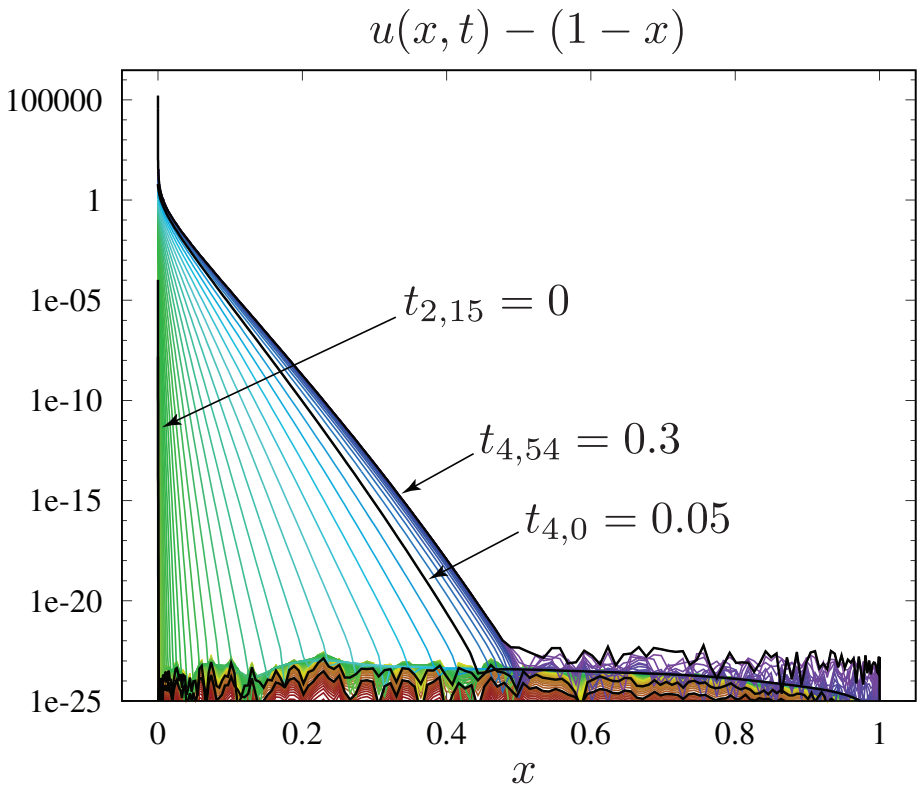

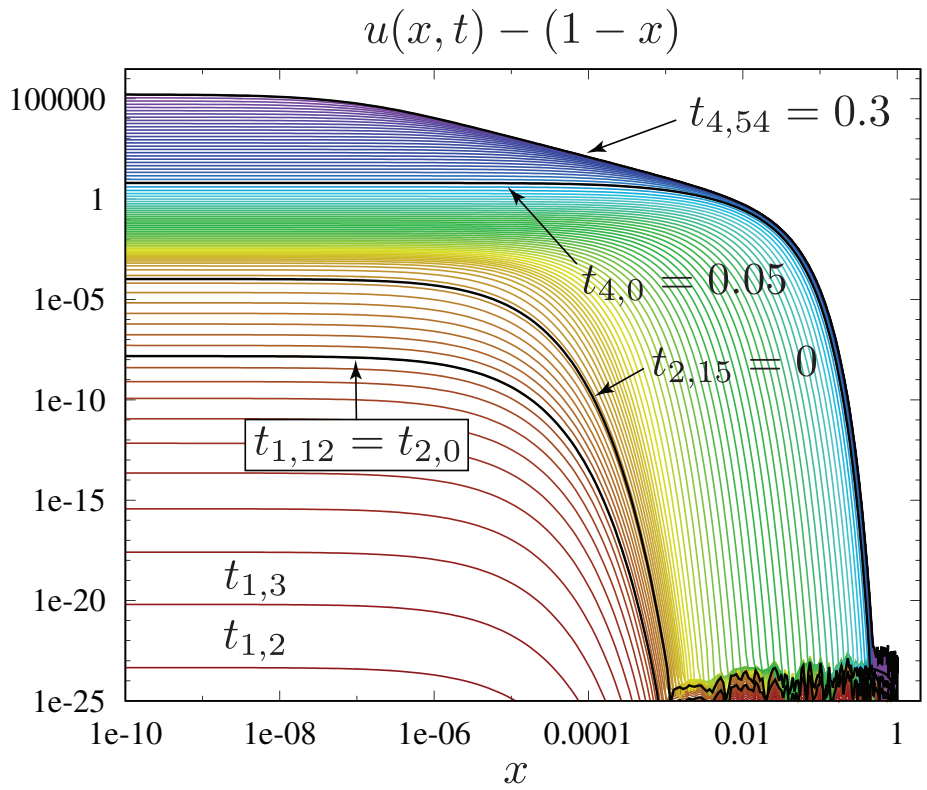

These plots show the difference between the initial data $(1-x)$ and the solution, $u(x,t)$, with a linear x-scale on the left and a log-scale on the right, which resolves the boundary layer.

Time is ordered from red-to-yellow-to-green-to-blue. Recall that $u(0,t)) = \rho(t)$, and note that the y-axis is logarithmic, going from 10^{-25} to 10^{5}.

Thanks!

Thanks for your attention!

And thanks to my sponsors the NSF, DARPA, ARO and the Flatiron Institute.

FLATIRON INSTITUTE
Center for Computational Mathematics

Outline

FLATIRON INSTITUTE
Center for Computational Mathematics

References

1. W. FELLER, *The parabolic differential equations and the associated semi-groups of transformations*, Ann. of Math., 55 (1952), 468–519.
2. M. KIMURA, *Solution of a process of random genetic drift with a continuous model*, PNAS, 41(1955), 144-150.
3. C. L. EPSTEIN AND R. MAZZEO, *Wright-Fisher diffusion in one dimension*, SIAM J. Math. Anal., 42 (2010), 568–608.
4. C. L. EPSTEIN AND J. WILKENING, *Eigenfunctions and the Dirichlet Problem for the Classical Kimura Diffusion Operator*, SIAM J. App. Math., 77(2017), 51–81.
5. C. L. EPSTEIN AND J. WILKENING, *Numerical Solvers for Kimura Diffusion Equations, I: Spectral Spatial and Temporal Discretizations via Exponential Time-Differencing*, preprint...soon to appear in arXiv.
6. C. L. EPSTEIN AND J. WILKENING, *Numerical Solvers for Kimura Diffusion Equations, II: The Resolvent Method*, preprint...soon to appear in arXiv.

FLATIRON INSTITUTE
Center for Computational Mathematics