James' CAIMS Lecture Notes

James H. Davenport
J.H.Davenport@bath.ac.uk

June 2009
(While on Sabbatical at the University of Waterloo)

Contents

Chapter 1

2 June 2009: See No Evil, Hear No Evil: Banks, Universities and Risky Investments — Peter Forsyth

See the slides at `http://www.iqfi.uwaterloo.ca/Seminars/SeminarFiles/forsyth_june2slides.pdf`. The talk was advertised with the following abstract.

> It is commonplace to make the assumption that, for the long-term investor, a portfolio containing risky assets is an optimal strategy. This type of thinking has also been prevalent in the management of university endowments. For example, an endowed professorship at a university contains an implicit guarantee of a certain level of spending, each year. However, in any given year, if returns on invested capital are insufficient, the university must cover the deficit.
>
> We determine the no-arbitrage value of guaranteeing a level of spending funded by an endowment that is invested in risky assets and which has a reserve account. For typical parameter values, the implied value of the guarantee is quite large, i.e., about 25% of the endowment capital. In other words, these endowment policies are extremely risky. Why did Universities expose themselves to this type of risk? This can be traced to the competitive market for attracting donors. There is a parallel here with the behaviour of big banks, which exchanged short-term gains for long-term risks.

The speaker had been called by the person handling the School of Computer

Science's endowment accounts, which had been invested in stock/bond markets. At the time of the "dotcom bust", the School suddenly was told it had no money to meet its commitments, particularly to research students. The speaker got a call, whose samitised version was as follows.

> Peter, don't you know something about finance? Can you tell me what is going on here?

When a donor sets up a scheme, the department is told the "real return", which used to be 5%. The department then commits a "promised cash flow", typically salary and on-costs, which will inflate on an assumed "academic inflation rate", generally higher than normal inflation[1]. The details of the spending rules differ from institution to institution: he will describe a "typical" one.

Riskless investments do *not* yield 4–5%, whereas riskless bonds typically yield 2%.

In investment return is greater than promised cash flow, the excess is added to the reserve account, and if that exceeds a given amount, the capital account is increased.

If the endowment return is negative, then cash is returned to the capital account from the reserve account. *If possible*, the promised cash flow is made from the reserve account.

Dybvig, *Financial Analysts Journal* (1999) quote.

> [There is] a significant probability of a shortfall. To assert otherwise is to state that the fund is certain that stocks will go up and that going long stocks and short in the riskless asset is, in effect, a riskless arbitrage. Such cheerful optimism may be an appealing personality trait, but it is not a healthy attitude for an investment manager.

1.1 no-arbitrage value of this risk

The spending rule specifies valuation rates t_i (typicaly annually) spaced by Δt. Let R_i is thevalue of the reserve fund at t_i, I_i the inflation in $[t_i, t_{i+1}]$, $C_r =$ cap on reserve fund.

Define the 'real gain' by $RG_{i+1} = S_{i+1} - S_i I_i$.

- If $RG_{i+1} < 0$ then the reserve fund R_i is drawn down to preserve real capital, and if (after this) $R_i = 0$, there are no disbursements.
- If $RG_{i+1} \geq 0$ or $R_i > 0$, then attempt to spend S_i.
- If this can be done, then top up the reserve, or if full $(C_r S_i)$ increase capital fund.

His first cut at spending rules had 19 "if" statements, but now has it down to 10. He doubts whether this has been done **by** Senates etc.

[1] JHD has issues here: see section 1.4.

Table 1.1: Notation

r = risk free interest rate,

σ = volatility,

dZ = increment of a Wiener process

$$dq = \begin{cases} 0 & \text{with probability } 1 - \lambda dt \\ 1 & \text{with probability } \lambda dt \end{cases},$$

λ = mean arrival rate of Poisson jumps; $S \to JS$

κ = expected value of a jump.

Let D_{i+1} be the disbursement in year $i + 1$, E_{i+1} the actual promised cash flow. $G_{i+1} = \max(E_{i+1} - D_{i+1}, 0)$ is the cost to the department. So universities are taking on a risk to attract endowments.

$$\frac{dS}{S} = (r - \lambda\kappa)dt + \sigma dZ + (J - I),$$

which is a jump stochastic process. The no-arbitrage value is a *path-dependent contingent claim* on this equation. He can solve with a partial integro-differential equation. Assume our cash flow has to last for 20 years (typical life of an endowment appointment — we can not re-appoint for a while and let the endowment build up at the end of such a period). We then know that the risk value is 0 at the end of this period, and we integrate backwards to find the present value of the risk. The valuation rules mean applying "jump conditions" at every valuation date. d'Halluin, Forsyth, Labahn, *Numerische Math.* (2004) solve this problem. PF is taking 15% as reserve cap (high, but increasing it beyond this doesn't help much, see slide 31). Assume both inflation rates to be 2% — optimistic in the case of general inflation, but *very* optimistic to assume the two are equal: essentially assuming no merit awards. Assume σ (volatility) 10%, risk-free interest rate of 4%. Initally, no jumps. Then the no-arbitrage value of this risk tends to limit at 30% of the initial capital. With a higher σ, things get worse, say 45%. Also worse[2] if there are rules, apparently very common in the U.S., that no disbursements can take place in future years until the capital account is above water. In this case, $\sigma > 0$ is *always* worse than $\sigma = 0$.

In his, very optimistic, worked example, the liability is still 25% of the capital, so "mark to market"[3] would place this on the balance sheet as a debit.

[2] Slide 28 versus slide 24.

[3] Which universities seem universally to be exempt from. PF noted that the "solution" to the U.S. banking problem seems to be replacing "mark to market" by "mark to make believe", which certainly makes thinks *look* better.

Could add jumps from Anderson *et al.* papers (2000) to fit S&P 500 data, or use no jumps, but adjust volatility to $\sigma = 0.28$. Results are very different. So the conventional wisdom, that the two are equivalent, does not hold in this case: presumably because of the path dependency.[4]

Why is the base case so bad? We could try changing the spending rules so that the promised cash flow comes first, i.e. before replacing the capital, and this indeed reduces the deficit from 33% to 15% (slide 33), at the cost, of course, of this policy being less likely to maintain the value of the endowment.

1.2 Why are we in this situation?

Chuck Prince quote: "as long as the music's playing, you've got to dance".

In fact, universities are taking naked (unhedged) puts in the stock market.

Inadequate "mark to market", either in banking or in university endowments.

Performance indicators are based on the short term, and do not take account of long-term risk, and decision makers "have no long-term skin in the game".

The conventional wisdom, that the stock marker is good in the long term, is (possibly) OK for the long-term investor, for whom the rate of random return is irrelevant. Essentially, if you are pulling out cash, a bad return at the start of the period is *far* worse than one at the end.

1.3 Conclusions

University spending rules and a complicated contingent liability, and the conservative estimate is 25%.

The risk/reward ratio for decision makers is wrong. See Chen, Forsyth and Vetzal: Insurance – Mathematics and Economics.

JHD asked whether starting with a non-zero reserve would help — yes, but not as much as one would think.

1.4 Lunch with PF — 9 June

JHD raised the question of the "academic inflation rate", which was assumed to be higher than the general inflation rate, to allow for 'merit' or 'incremental' pay rises, over the 'cost-of-living' ones. There is the same issue at Bath, where the Director of Finance looks at the year-on-year rises for *those in post*, and does not allow for the fact that retiring professors should, in general, be replaced by junior lecturers, with a consequent substantial saving. In "steady state", he argued that the pay bill should[5] more closely approximate general inflation.

PF agreed, but pointed out that an *individual* endowment, which is how these things tend to be accounted, was not in steady state. It would be possible,

[4] JHD wonders: "Is this a more general result?"

[5] "Should" in the sense of "is likely to" rather than "ought to" — we both agreed that academics are, by most measures, significantly underpaid.

JHD argued, to take "academic inflation" into account, and, assuming a 20-year life span, to plan on the average spend rate over that period, rather than the initial spend rate. Hence one would expect to build up a reserve in the first few years. PF said he might model this.

The real drawback was designing a mechanism that was financially sound and which could be "sold" to donors. JHD agreed, but since universities *should be* planning for the long term, and be seen to be doing so, it ought to be possible to cnvince donors of this. Admittedly, one wouldn't be spending the whole endowment from day 1, but the case could be distibnguished from that of not filling the post, since the person was in post, and the expenditure *was* committed.

We discussed footnote 4. It is indeed a general rule that the equivalence of jumps and volatility does not apply to *path-dependent* claims. This is in fact the main problem underpinning the disasters behind various endowment funds: ManuLife in Canada and Equitable in the U.K. being the most egregious examples. See Chen, Forsyth & Vetzal *Insurance: Mathematics and Economics* (2008). This estimates that insurance companies are "undercharging" by a factor of about 3.

At the end, we discussed the fact that he was speaking at CAIMS in a couple of days time, and he admitted the Waterloo talk has been by way of a rehearsal.

Chapter 2

11 June

2.1 See No Evil, Hear No Evil: Banks, Universities and Risky Investments — Peter Forsyth

Since this was a repeat of that described in chapter 1, the notes are merged.

2.2 Generating Efficient Numerical Evaluation Routines for Bivariate Functions via Tensor Product Series — Keith Geddes

History: approximations for univariates were one of the early tasks for digital computers.

$$s_n(x, y) = \sum_{i=1}^{n} g_i(x)h_i(y)$$

is our meaning of "Tensor Product Series". We say it is *natural* if each factor can be derived from the original by a finite set of linear operators.

Define

$$\mathcal{Y}_{(a,b)}f(x, y) = \frac{f(x, b) - f(a, x)}{f(a, b)}.$$

We proceed by analogy with Newton iteration. At each step, we find[1] the point (a, b) in the region of interest for which $f-$series is maximal, and the next term is ??

A consequence of this is that the interpolation agrees with the original on all horizontal and vertical lines *through* the (a, b) points chosen.

It is not immediate that these series converge, though it is intuitively appealing. There are proofs for many cases of *symmetric* functions.

[1] In fact, we don't do a complete search.

The series appears to have exponential growth in size, but in fact it is only evaluation as n points, so there is a compact and efficient matrix–vector representation (though care needs to be takento ensure stability).

Motivation/Example: Maple found that, even for hardware floats, Bessel[2] was too slow. Motivating case: $\int_0^4 \int_1^{10} Y(\nu, z) d\nu dz$, which takes 26 seconds, as opposed to 0.1 *if* the numerical library understands the integrand.

In general, having got the tensor product series, the individual summands, now univariate, are approximated by Chebyshev or Chebyshev–Padé.

Taking $T_\nu(z)$ over $[0, 5] \times [1/4, 5]$ (recall that $(0, 0)$ is a singularity), we end up 15 regions, in each of which we have such a tensor series.

For $J_\nu(z)$ on $[0, 1] \times [1/1000, 1]$ allowing degree 20 can be done with one region,and is generated in 18 seconds, whereas if we insist on 10 terms, we need 22 regions and a 95 second generation time. In both cases, the maximum error is about $4 \cdot 10^{-15}$.

JacobiSN on $[0, 6] \times [0, 6]$ shows an improvement of about 180 in `evalhf` mode. In cases where the function is very expensive to approximate, he can see 1800.

2.3 Reconstruction algorithms in computerised tomography — Nargol Rezvani

Basically interactions source–object–detector. Invented in 1972, but today's goals, which also include reduced dosage[3] etc., include her goals of faster reconstruction.

FBP (Filter Back Projection) is the common method of reconstructing tomography. Traditionally we assume straight lines (no refraction), monochromatic X-ray source, Beer's Law:

$$\frac{dI}{ds} = -\mu I \Rightarrow \int_i \mu ds = -\ln\left(\frac{I}{I_0}\right).$$

The Radon transform provides the mathematical model for the measured attention

$$[R\mu(x, y)](t, \omega) := \int_{I_{(t,\omega)}} \mu(x, y) ds.$$

Central Slice Theorem (d'après Epstein2003)

$$\int_{-\infty}^{+\infty} [R\mu(x, y)](t, \omega) e^{-itr} dt = \hat{\mu}(r\omega).$$

But the use of this depends on the geometry, and these days we tend to use 2D detectors for speed of scanning.

[2]Note that Bessel in Maple has already had quite a lot of work done on it, so this is not a straw man.

[3]One way of actieving this, of course, is more accurate reconstruction, so the goals interact.

OSCaR — an open-source MatLab implementation of the Feldkamp–Davis–Kress algorithm. See her home page at Toronto.

FBP has limitations, notably dealing with movement and the monochromatic assumption. The artefacts are known as “beam hardening”

Replacement — Algebraic Reconstruction Technique (ART). Used initially, but soon replaced by FBP as it was (then) too computationally expensive. We formulate the reconstruction problem as a series of linear equations. Typically use a pixel basis, and then need to solve a linear system where the matrix is the Radon transform of the pixel basis functions. The linear system might be big, e.g. 16K×19K. Use Kaczmarz’s method (1937!). Normally 3–4 iterations are used, and the stopping criterion is an experienced operator’s “looks OK”. Sometimes 2 is “enough”. In practice, relaxation is also incorporated into the iterations. Conceptually simpler, but lacks the computational speed and the acceptance of FBP.

Statistical techniques (SAR) are sometimes used to account for polychromatic effects. Assume that μ now depends on the X-ray energy as well as (x, y, z), but assume some independence ??.

This has been solved by penalised least-squares methods.

Open problems:

- How to handle the attenuation coefficients.
- reducing the number of coefficients.
- hybrid ART/SIRT?
- More sophisticated generation of the attenuation matrix in ART.
- Extension to 4D (i.e. time) tomography.

Chapter 3

12 June

The day started with the CAIMS Doctoral awards.

3.1 Modeling group formation and activity patterns in self-organizing animal aggregations — R. Eftimie (McMaster)

Motivational slides — "finding Nemo", and more generally schools of fish, or flocks of birds (starlings), typically very tight, or Serengeti animals foraging, much looser. Also bacteria. Common property is that this is self-organising, with no "leader". Genetic motivations for flocking: communicating abut food, defence innumbers against pedators, finding mates. Is there a common biological mechanism?

- Lagrangian (individual-based) models can generate swarms and ripples, but there are not many analytic techniques.
- Eulerian techniques (PDEs) can generate ripples, stationary aggregations and vortices, and have access to all the PDE toolkit.

Both kinds have assumed three forces: repulsion (personal space), alignment and attraction. Her research: a "non-local hyperbolic model with constant speed". She actually models 1D. $u^+(x,t)$ = density of individuals moving right, similarly $u^-(x,t)$. Let $\lambda^+(U^+,u^-)$ be the rate at which individuals moving right turn round.

Let $\lambda^+(U^+,u^-) = \lambda_1 + \lambda_2 f$ where λ_1 is the random term, and f models the interactions, based on various kernels for the three forces. In her model, f will depend on the signals received from other animals. So for example, alignment decisions can depend only on signals

If the λ are Lipschitz coninuous then there is a unique mild solution. Periodic initial conditions on a finite domain yield periodic solutions. In the appropriate limit, this tends to a well-known parabolic model, but one which does not model pulses. If she assumes asymmetric (+/−) sugnal reception, she can get travelling pulses.

She has numerical results, with periodic boundary conditions (matches experiments with ants). This can model ripples, and resembles the bacteria behaviours. It also models classic results such as staionary pulses, travelling pulses and travelling trains. She also gets new (observed in real life, but not previously modeled) patterns, stationary breathers (expanding/contracting) groups, travelling breathers, feathers (?), zigzags (as seen when flocks of birds suddenly change direction), as well as a semi-zigzag, for which she has no natural analogue. Different patterns tend to come from different assumed signal reception behaviour. Bifurcation behaviour observed in nature can also be seen in her models.

In conclusion, complexity of animal pattersn can be explained by assuming different reception mechanisms.

Q 2D?

A Not possible, partly because of time, but also it is now necessary to model angle, rather than simply +/−.

3.2 Computing on surfaces with the Closest Point Method — Colin McDonald

Many applications of PDEs on surfaces: biological, materials, computer graphics/images, etc. Closest point models, to every x we associate the closest point of the surface $cp(x)$, do not assume orientation, hence Klein bottles are possible. Need an intrinsic gradient ∇_S and a Laplace–Beltrami operator Δ_S. Principle: the surface gradient of v agrees at the surface with the standard gradient of u, i.e.

$$\nabla_S(v(x)) = \nabla u(x) = \nabla[v(cp(x))].$$

The explicit method is formualtes with Euler, but Runge-Kutta and multigrid are also applicable. In fact one computes a tight band round the surface, but, unlike other embedding methods, we do not need to impose additional boundary conditions at the edges of the band.

Can model, e.g., forest fires on a mountain surface using existing MatLab code with a one-line modification.

Implicit methods can also be used, especially if the system is stiff. However, this is not quite trivial, and we need to adjust the matrix product to ensure that (? the gradient of) u is still normal to the surface.

Again, it only takes a one-line addition to model biharmonic problems.

Can use this to find Laplace–Beltrami eigenvalues on a Möbius strip.

Q But these solutions[1] are non-physical: is the problem that the curvature is not incorporated?

[1] Subsequent discussion revealed that the questioner's main concern was the heat equation, where the *volume*, not the surface, heats up.

A The forest fire model does incorporate the fact that convection is vertical, not normal[2].

3.2.1 Subsequent discussion

JHD encountered the speaker in the lunch queue the next day, and pursued the topic. The speaker admitted that he now understood the question better. The point is that his methodology/software *does* solve PDE on the surface as posed. The PDE on the surface might not be the restriction of the PDE from space to the surface, though. From his point of view, that's the user's problem: an attitude which JHD is inclined to agree with.

It subsequently occurred to JHD that the point about forest fires, as raised in the original question is related to the fact that heat transfer to the surrounding air *will* depend critically on the curvature. Possibly this is related to the King's Cross fire, where extensive supercomputer simulations revealed a previously-unknown physical effect relating to the formation of super-heated air channels. Clearly here it would not be sufficient to model the PDE on the surface.

3.3 Rayleigh-Quotient Algorithms for Nonsymmetric Matrix Pencils – Amir Amiraslani; Peter Lancaster (Calgary)

Consider the pencil $A - \lambda B$, not assuming symmetry. Assume semi-simple, that the pencil is regular, and that B is nonsingular. Let the eigenevlaues (possibly repeated) be λ_i, with right eigenvectors x_i and left y_i. The proof method has an unconventional structure.

1. (At least) Second-order convergence of eigenvalues.
2. (At least) Second-order convergence of eigenvectors.
3. Third-order convergence of eigenvalues.
4. Third-order convergence of eigenvectors.

Numerical data shown seemed to show cubic convergence (though it doesn't take many cubic steps to hit machine precision effects!).

Each eigenvalue λ_p is the centre of an open disc D such that, if we ever land in D, the iteration converges, not merely to *any* egenvalue, but in fact to λ_p.

[2] JHD does not think that this was a real answer to the question, and the speaker seemed to admit this.

3.4 Algebraic Systems Biology — Brandylin Stigler (SMU)

Mathematics has been very effective in explaining physical phenomena, but rarer in biology, though cites disease models, and predator-prey. Molecular biology has been very reductionist, but we can now try to do systems biology: understand the qantitative eatures of a multicomponent biological system.

Example: lactose metabolism. There are feedback loops with lactose concentrations and glucose concentrations affecting the transcription of genes for lactose-processing proteins. The system seems to be bistable. There are differential equation models, but she wishes to do a Finite-State simulation.

Can also look at the inputs into a large graph, say 100 nodes, and use monomial ideals to determine a minimal "wiring diagram" (formal term: Polynomial Dymanic System - PDS). See her Gröbner-fan paper in ISSAC 2007.

There is still a dichotomy between continuous and discrete models. Example: stress in yeast. Have 13 genes identified, so including external perturbation gives is a 14×14 system. The use of PDSs let us reduce the number of parameters by 38%. The model predicted an interaction GPX2/KN7, which was validated in the laboratory.

Conclusion: PDS provides a *compact* representation of the model space, permitting a reduction in parameter space, and consequent improvements in complexity.

3.5 Symbolic-Numeric Computation of Distances between curved objects Avoiding Footprints Determination — Lalo Gonzalez-Vega

Low-degree objects typically ellipses/ellipsoids.

Example: relative position of two ellipses, say $X^TAX = 0$ and $X^TBX = 0$. Then the key object is $f(\lambda) \equiv \lambda^3 + a\lambda^2 + b\lambda + c$. Have 3/4 (depending on $\text{sign}(a)$) polynomial constraints on a, b and c for separation. See section 3.5.1

Special configurations imply strong factorizations of the polynomials involved: how to make use of this algorithmically? One important point is that the minimum distance varies continuously, even though the footpoint can change discontinuously.

3.5.1 Subsequent discussion

In fact f is cubic for ellipses, and quartic for ellipsoids. The condition for ellipses is that$f$ should have one (two for ellpisoids) distinct positive real roots. In the case of the cubic, the situation clearly depends on $\text{sign}(a)$. The case of ellipsis is discussed in Computer-Aided Geometric Design about 3 years ago ([EGVdR06]); ellipsoids is to appear Computer-Aided Design ([GVM09]).

LLV has a paper [GVPB09] on Vermielle's degree 4 polynomial (univariate, but parametric) describing a change of coordinates. Borkowski's polynomial does a similar job. Again the problem is to characterise the root behaviour o fthe polynomials. There is a 'bad region', on which this transformation does not work. The paper is in J. Geodesy. condition is the union of two semi-algebraic sets.

3.6 Mathematical challenges in the modelling of biological invasions — Mark A. Lewis, Alberta

Fisher's 1930s studies of population spread. Apparently muskrats were introduced to Prague in 1905, and $\sqrt{\text{area}}$ colonised is almost linear with time. If $f(u) \leq f'(0)u$ then the spreading speed is determined. Experimental spread rates tend to correlate well with the diffusion model, with one exception (cereal leaf bee?). Reid (1899) noted that trees recolonised behind the ice age at about 100m/year (oak data from North America, is most visible, but Reid used oak in the U.K.). Conversely, Skellan calculates that to fit the spread data, acorns must on average distribute 0.83km, again unlikely.

Claim: rare long-distance dispersal events typically change the distribution from Gaussian to leptokurtic. He has an integro-difference model, which seems to fit the data, also Red Maple data.

Multispecies competition paradox (data are red/grey squirrels in the U.K.). Apparently in Alberta, the red squirrels are in Edmonton, and the grey ones in Calgary :-). Again, Fisher's linearisation does *not* fit the data.Convert the competeive system u/v into a coperative system p/q with $p = u$, $q = 1 - v$.

Chapter 4

13 June

4.1 Symmetry-Breaking, Synchrony-Breaking and Coupled Cell Models — Martin Golubitsky

A simple case is $\dot{x}_1 = f_1(x_1, x_2)$; $\dot{x}_2 = f_2(x_1, x_2)$; which can be represented as

$$1 \rightleftharpoons 2. \tag{4.1}$$

We can call this the network architecture of the system.

Symmetries γ might fix $x(t)$ *pointwise*, $\gamma x(t) = x(t)$, called K-symmetries, or *globally*, $\gamma\{x(t)\} = \{x(t)\}$, called H-symmetries. In fact (since the groups are finite) H/K is cyclic.

Generically, $\ker J$ is an *abolutely irreducible* representation of Γ. Reduction implies that there is a unique steady-state bifurcation theory for each absolutely irreducible representation — not necessarily easy, but we know what we're trying to do.

A Hopf bifurcation hasthe origin as an equilibrium for all λ, and goes from from a spiral sink to a spiral source as $\lambda \nearrow 0$. The amplitude growth is $\lambda^{\frac{1}{2}}$. For equivariant Hopf, the eigenvalues of J are $\pm\omega i$.

In (4.1) the eigenvalues of J are $\alpha \pm \beta$, and $\alpha + \beta$ being critical gives us synchronous bifucrations. For a three-cell unidirectional ring, $\Gamma = \mathbf{Z}_3$, and we get discrete rotating waves. For a three-cell bidirectional ring, $\Gamma = \S_3$, and we getthe orevious pus in-phase periodic solutoins and out-of-phase periodic solutions, and two cells half-period out of phase, with the third going at twice the frequency (the most interesting case).

A seven-chain cell, but with 3 feeding back into 1 (so every cell has precisely one input) — Ian Stewart's example — we can get

$$\{x_1 = x_4 = x_7; x_2 = x_5; x_3 = x_6\} \tag{4.2}$$

is in fact invariant, but without symmetry. There is an equivariant colouring: every cell of a given colour, as given by (4.2), receives input from cells of a given colour. This sort of "network symmetry" is key. For example, we can take a 2-D lattice with NSEW connections, alternately coloured black and white, so every cell has two black and two white inputs. If we take a 45° diagonal axis, and swap black and white, the colouring is *still equivariant*: every cell has two black and two white inputs.

Take a bidirectional 3-ring, and colour two cells red: an equivariant colouring. Then we can form a quotient network by idetifying the two red cells (note we need a multigraph, and have self-coupling from red to itself). *In generality*, every admissible vector field on the full graph restricts to one on the quotient, and *vice versa*. Now the infinite lattice above has a two-element (black/white) quotient network ....

Example of an asymmetric network whose quotient is the three-cell bidirectional network. Then the interesting case above *can* be lifted to the asymmetric network.

4.1.1 Synchrony-breaking bifurcations

On a egular three-cell network, there are 34 different allocations of valencies (?): 14 are transitive, but the rest have a feed-forward behaviour (common on neuro-science). Consider

$$1 \to 2 \to 3 \tag{4.3}$$

with 1 feeding into itself as well. The matrix is $\begin{pmatrix} \alpha+\beta & 0 & 0 \\ \beta & \alpha & 0 \\ 0 & \beta & \alpha \end{pmatrix}$. This network sopports solution with x_1 is equilibium aand x_2, x_3 time periodic.

Theorem: There is a codimension one steady-state bifurcation correspoding to each absolutely irreducible subspace.

Theorem: There is a codimension one Hopf bifurcation correspoding to each irreducible subspace.

For networks, take any Jordan form N whose eigenvalues lie on the imaginary axis. There is a regular[1] network with a codimension one bifucation with $dF|E^C = \rho N$.

For the 34 3-networks, we have the following classification.

20 networks have real simple eigenvalues.

5 have simplex complex (no zero) eigenvalues

3 have two synchrony-breaking eigenvectors

5 are nilpotent

1 is double with synchrony-preserving eigenvectors.

Note that this split is *not* the same as the feed-forward split.

Conjecture: the number of regular networks grows superexponentially, but the number of eigenspace types grows much more slowly.

[1] All nodes have the same number of inputs.

4.2 chebfun computing — Nick Trefethen (Oxford, U.K.)

Built on MatLab, both actually and conceptually.

1. Floating-point arithmetic is (nowadays) a well-worked out approach.

* One computes exactly, then rounds.

2. Would like to do the same with functions

3. Numerical linear algebra is well worked-out, but the vectors come from functions, and the matrices from operators.

4. Can we represent these more directly?

A 'chebfun' is a series of 'funcs', each of which is a Chebyshev series on a region. Many MatLab functions are overloaded, so that `sum` computes an integral, for example. chebfuns are *not* as good as best approximations, but "it's not worth it" in terms of time, and in practice the fact that 'best' minimises the ∞-norm in not really what one wants.

$\backslash$ can be overloaded for differential operators! The big goal though would be 2D-chebfuns, but see Geddes' talk (section 2.2). *However* in 1D everything is an interval, but in 2D not everything in a rectangle, so it is not straightforward.

Chapter 5

14 June

5.1 Genetic Strategies for Controlling Mosquito-Borne Diseases — Alun Lloyd (NCSU)

Joint work with many people in Entomology. Main focuses is Dengue, which is a mosquito-borne, but viral, infection. There are in fact four dengue serotypes, so one can get it four times. *Most* cases are mild, and people (including one of his own co-workers) often dont realise they have had it, just writing it down to 'flu. About 1% of cases lead to dengue haemorrhagic fever, which is fatal in 20% of untreated cases. Having immunity to one serotype makes one *more* susceptible to the others, and increases the likelihood of dengue haemorrhagic fever. This makes a vaccine very problematic.

Female mosquitoes need blood to lay eggs. After the mosquito bites an infected person, it takes 7–14 days (temperature dependent) to incubate in the mosquito before the mosquito is infectious. Note that an adult female mosquito has a life-span of about 21 days, so the incubation period is a significant fraction.

$$R_o = ma^2bcD_HD_MP \tag{5.1}$$

m = mosquitoes/person

a = bite rate (squared because of two bites)

b, c transmission probabilities

D_H, D_M duration of infection

Malaria mosquitoes bite at night, so bed-nets are very effective here, but dengue mosquitoes bite by day. Malaria pills exist, but have side effects: there is no dengue equivalent. Hence dengue control is problematic, which leads to mosquito reduction programs.

Sterile male approach — one factory in Guatemala produces $2 \cdot 10^9$ sterile fruitflies () per week (using canadian radiation sources). This approach *does* work: the screwworm fly has been eradicated from the U.S. and Central America.

But tends not to work on mosquitoes: one problem is that lab-based mosquitoes are not very attractive to wild females. Also programmes have failed for sociological reasons (El Salvador, India).

Another approach is population replacement. If we could shorten the lifespan, we would make a major impact on transmission. One guy at Colorado State has engineered a variant of the mosquito that cannot transmit *one* of the four stereotypes. *But* this GM tends to make the mosquito less fertile, hence it cannot be a long-term solution.

There is a bacterium (Wolbachia) which is maternally transmitted, and halves the life-span of the mosquito. Models show that there is a 27% "invasion threshhold".

MEDEA: Maternal-effect domiant embryonic arrest. Offspring from a Medea female that do not inherit the gene die. Hence the gene 'beats Mendel'. Bruce Hay (CalTech) has a synthetic Medea in drosophilia. Science 316(2007) p. 597. They are having problems with the mosquito, though. The Medea gene does spread, again subject to an invasion threshhold.

Hence need to model multiple processes at multiple scales. Will mosquitoes carrying the desirable gene spread? Mosquito and human populations are highly structured and heterogeneous — will 'our' mosquitoes spread thrugh the town?

PDE Models for the Spatial Spread of Wolbachia (Heredity 43(1979) p. 341). In the California Central valley, Wolbachia travels at 100km/year — occasional long-move events[1]. Age structure is highly important, and the effectiveness of the release process is dependent on this: releasing young adults is 10 times as effective as releasing eggs.

Has an "all-singing" model[2], goes down to distrbution of mosquito weights, rainfall, temperature, human society (houses, water containers[3]). It turns out that some of the biology is unknown or almost so (there are *seven* dat points on mosquito larva populations in water containers). Unfortunateky, it turns out that much of the heterogeneity is important. Rare long-range movement is a key determinant — see section 3.6. Method of release is also key — dispersed release tends to do better, but is much harder to do in practice. Timed release is also more effective, but again harder to do.

Conclusion: if the molecular biologists *could* produce a GM mosquito with the right properties, then we can model the dispersal methods, and recommend the technically best solutions, but the Gates Foundation projects are also engaging the people to gain public acceptance. It is also imprtance to use these to model risk.

[1]JHD notes that this is the same point as made in §3.6: the speed is dominated by the rare events.

[2]Hand-coded C++, very slow.

[3]These mosquitoes, unlike malaria ones, prefer clean water.

5.2 Superconvergent Interpolants for Efficient Error Estimation in 1D Time-Dependent PDE Solvers

Typical test problems: Burger's equation, Catalytic equations.

"Methd-of-LInes" software. Type 1 has a discretised spatial mesh, which approximates the PDE by a family of ODEs, and then passes this to DAE software. A second class fixes the *number* of mesh points, but lets them move. Example D03PPF (Berzins). A third class works from a user-provided spatial error tolerance, which uses an adapive mesh, HPNEW and HPDASSL are examples.

His work is BACOL/ BACOLR. They ahve spatial discretization, approximating the solution as B-splines with time-dependent coefficients. $U_s(x,t)$ is required to satisfy the PDE at the mesh points. The Jacobian of eacj DAE is "almost block diagonal". If we are asked to compute U or order p, we also compute $\overline{U}$ of order $p+1$ and compaute thetwo to get error estimates. Do spatial remeshing aiming to do equidistribution of the error between mesh points. BACOL uses DASSL for the DAEs, while BACOLR uses RADAU5, but modified to use COLROW. The error estimate is

$$E_s(t) = \sqrt{\int_a^b \left(\frac{U_s - \overline{U}_s}{ATOL + RTOL|U_s|}\right)} \tag{5.2}$$

so the extra cost of computing $\overline{U}$ is substantial. At the mesh points,our error is $O(h^{2k})$, but at other points $O(h^{k+2})$: in fact $P(x)O(h^{k+2})$. If we evaluate at the zeros of (known) P, we get $O(h^{k+3})$, i.e. superconvergence. Therefore we can use this to replace $\overline{U}$. Use Hermite–Birkhoff interpolants, using two points from *outside* the current interval to ensure that we have enough independent data points. The error ends up dependent on the square of the ratio of the mesh sizes. SCI error estimates lead to the introduction of (a few) extra mesh points to reduce this effect, automatically (and better than his manual attempts).

5.3 Anisotropic Mesh Generation

We have a set U of functions satisfying Dirichlet boundary conditions, and we use a Galerkin formulation of the variational problem. This leads to adapted FEM: most work is isotropic.

Q Is this split analogous to the split between the inherent condition number of the problem, and the condition of the soluton.

A Yes.

5.4 High-Order Solution of the Grating Diffraction Problem for Singular Perfect Conductors — Mike Haslam

Scattering from gratings with corrugation heights large with respect to the period remains a chanllenging problem, not least with respect to quality control problems in chip manufacturing. This requires a fast solution to the inverse problem. We are working on a direct solution of the forward problem, which could be used as part of an inverse solver.

$$\sin\theta_n = \sin\theta + \frac{n\lambda}{P} \tag{5.3}$$

so there are a ininte number of propagating modes, which can be desribed by an integral equation:

$$\frac{\partial u_{inc}(r)}{\partial n(r)} = \frac{1}{2}\phi(r) \pm \int_S \frac{\partial G(r,r')}{partialn(r)}\phi(r')dS' \tag{5.4}$$

The Green's function $G_P(x-x', z-z') = \frac{1}{4}\sum_{m=-\infty}^{\infty}$ of an expression in slowly-converging Hankel functions. It possesses a logarithmic singularity at every integer multiple of $(x-x')$.

We use high-order integration with the trapezoidal rule on *periodic* functions, which has amazing convergence, e.g. for $\int \exp(\cos^2 x)$ we get an error of 10^{-23} when we integrate over the whole period ($N = 8192$), rather than 10^{-7} when we integrate over part of the period.

Need very accurate tabulation of $\int_{-1}^{1} \ln|y|T_n(y)$ and friends. The closed-form expressions are numerically unstable, but he has asymptotic expansions.

He can work with features of the size of 40 periods, whereas the Rayleigh approach limits at 15% of a period. Comments that this is all based on Chebyshev.

5.5 Multiresolution approach to accelerate the finite volume reservoir technique – E. Lorin

We want to solve Hyperbolic Systems of Conservation Laws (HSCL) by accurate Finite Volume Schemes.

$$\partial_t u + \partial_x f(u) = 0 \tag{5.5}$$

Hyperbolic if $A(u) = df(u)$ is diagonalizable over $\mathbf{R}$. We have a finite volume flux scheme, which is a generalization of the order-1 upwind scheme. But order 1 FVS are very diffusive. Higher-order methods need to introduce a flux limiter, which resolves problems near singularities, but is computationally complex and has hard fundamental properties.

J. Sci. Comput. 2007, introduces a *reservoir* technique, which is simple and not very diffusive. It is expensive, and can be accelerated by multiresolution analysis. The cost of the Riemann solver is $O(N^2)$ (though it would appear larger: the point is that not every cells needs solution every time step).

Other than initially, he never works on the finest mesh, which makes his scheme "fully adaptative", as opposed to previous "semi-adaptative" methods. See A. Cohen *et al.* (Math. Comp. 2007).

Bibliography

[EGVdR06] F. Etayo, L. Gonzalez-Vega, and N. del Rio. A new approach to characterizing the relative position of two ellipses depending on one parameter. *Computer Aided Geometric Design*, 23:324–350, 2006.

[GVM09] L. Gonzalez-Vega and E. Mainar. Solving the separation problem for two ellipsoids involving only the evaluation of six polynomials. *Preprint*, 2009.

[GVPB09] L. Gonzalez-Vega and I. Polo-Blanco. A symbolic analysis of Vermeille and Borkowski polynomials for transforming 3D Cartesian to geodetic coordinates. *To appear in J. Geodesy*, 2009.