Cluster Induction

Peter Jäckel

May 2018

Outline

1) The case for multi-dimensional conditional expectations of moderate accuracy

2 Scattered (cluster) data interpolation

3 Meshless induction

4 There's more to life than Crank-Nicolson

5 Implementation aspects and optimisations

6 Recent innovations and future work

1 Introduction

Introduction

 CVA, FVA, and other XVA calculations, even in their simplest form, all involve calculations that require some kind of conditional expectation of a contract's future (deflated) value v(t, x) for some t > 0 with x being a suitable (Markovian) state vector subject to some conditioning, e.g., conditional on being positive:

$$F(t) = \mathbf{E}\left[\left(v(t, \boldsymbol{x})\right)_{+} \cdot f(x)\right]$$
(2.1)

for some function $f(x) \ge 0$.

• This is formally a forward-conditional expectation:

$$\mathbf{E}\left[\left(v(t,\boldsymbol{x})\right)_{+}\cdot f(x)\right] = \mathbf{E}\left[v(t,\boldsymbol{x})\cdot\mathbf{1}_{\{v(t,\boldsymbol{x})>0\}}\cdot f(x)\right]$$
(2.2)



- In low dimensionalities, this can be computed directly with conventional finite differencing methods.
- In moderate or higher dimensions, it can, in principle, be computed with a Monte-Carlo simulation if at least some kind of Markovian approximation for the indicator function $h^{\star}(t, \boldsymbol{x}) \approx \mathbf{1}_{\{v(t, \boldsymbol{x}) > 0\}}$ can be found, i.e. $F^{\star} \approx F$ with

$$F^{\star}(t) = \operatorname{E}\left[v(t, \boldsymbol{x}) \cdot \boldsymbol{h}^{\star}(t, \boldsymbol{x}) \cdot f(\boldsymbol{x})\right]$$
(2.3)

 Note that any small error ε in the domain in x on which the approximation for h^{*}(t, x) is 1 only enters the calculation in second order in ε. To see this, consider that h^{*}(t, x) = 1 for x ∈ ϵ⁽⁺⁾ where ϵ⁽⁺⁾ is a small subdomain on which an accurate calculation would tell us that v(t, x) < 0, but otherwise h^{*}(t, x) is telling us the truth.



In other words, $h^{\star}(t, x)$ is giving us some *false positives* and the calculation (2.3) will add up all of the required positive contributions, but also some erroneous negative contributions from $\epsilon^{(+)}$.

The net result is $F^{\star} < F$.

Peter Jäckel **Cluster Induction** May 2018 5 / 117

1 Introduction

• Now consider the reverse situation that $h^{\star}(t, x) = 0$ for $x \in \epsilon^{(-)}$ where $\epsilon^{(-)}$ is a small subdomain on which an accurate calculation would tell us that v(t, x) > 0, but otherwise $h^{\star}(t, x)$ is telling us the truth.



In other words, $h^{\star}(t, x)$ is giving us some *false negatives* and the calculation (2.3) will not add up all of the required positive contributions since some positive contributions from $\epsilon^{(-)}$ will be missing.

The net result is $F^{\star} < F$.

- We see that any error e in the domain where h[★](t, x) is positive leads to F[★] < F.
- Thus, $\epsilon = \{\}$ is a maximum for F^* , i.e.,

$$\left. \frac{\partial F^{\star}}{\partial \epsilon} \right|_{\epsilon = \{\}} = 0.$$
 (2.4)

• Thus, a local expansion of F^* in ϵ leads with a second order term.



The conventional methodology to establish an approximate indicator function h^{*}(t, x) is to use a regression of path values of futures payments onto a chosen functional form, e.g., typically, a multivariate polynomial to approximate the value function itself, e.g.,

$$\boldsymbol{v}^{\star}(t,\boldsymbol{x}) = \sum_{i,j} \alpha_{ijkl} \cdot x_i^k x_j^l$$
(2.5)

$$\alpha_{ijk\ell} = \arg\min_{\alpha_{ijkl}} \sum_{q=1}^{N} \left(\boldsymbol{v}^{\star}(t, \boldsymbol{x}_q) - \boldsymbol{v}_q(t) \right)^2$$
(2.6)

$$v_q(t) = (\text{deflated}) \text{ value of payments on path } q \text{ after t}$$
 (2.7)

$$h^{\star}(t, x) = \mathbf{1}_{\{v^{\star}(t, x) > 0\}}$$
(2.8)

 Then, in a second simulation, h^{*}(t, x) is used to trigger the contribution of each path to approximate F^{*} by virtue of the tower law:

$$F^{\star}(t) = \operatorname{E}\left[\boldsymbol{h}^{\star}(t, \boldsymbol{x}) \cdot f(\boldsymbol{x}) \cdot \operatorname{E}\left[\boldsymbol{v}(T, \boldsymbol{x}) | \mathcal{F}_{t}\right]\right]$$
(2.9)

- When the valuation of v(t, x) itself involves callability, it may in turn require another layer of regression approximation(s).
- In (very) low dimensions, this works well.
- However, as the dimensionality increases, the technical aspects of the regressions to approximate the value function increase considerably, not least because:-
 - the number n of coefficients grows fast with the dimensionality d:

$$n = \left(\begin{array}{c} d+k\\k\end{array}\right)$$

with k being the maximum total power of the polynomial.

- the shape of v(t, x) may be poorly matched by the used basis function set.
- the regression effort grows like $\mathcal{O}(\max(m, n) \cdot n^2)$ for m sample paths.



Let's recapitulate:-

- Finite differencing methods suffer the curse of dimensionality due to their requirement of a lattice of points.
- Monte Carlo regressions (e.g., onto polynomials) become awkward in high dimensions.

An important aspect of this is the difficulty to choose the right basis function set, or a *suitable interpolation candidate*.

Engineers have long grappled with the issue of (multivariate) interpolation to fit non-trivial topologies over scattered nodes.

In 1979, R. Franke produced a 380(!) page technical report comparing 29(!) different surface interpolation methods [Fra79].

The most impressive method in these tests is the **multiquadric method** of Hardy [Har71]. It is consistently best or near best in terms of accuracy, and always results in visually pleasant surfaces.

In 1984, Micchelli [Mic84] proved that the linear system to compute the interpolation coefficients for the "MultiQuadric Surface" (MQS) method [as it was then known] is always regular (in full arithmetic precision).

This was an important breakthrough since the Mairhuber-Curtis theorem [Mai56, Cur58] declared that, in two or more dimensions, a generic interpolant *leads to a singular system* for infinitely many configurations of interpolation nodes.

2 Scattered (cluster) data interpolation

The crucial factor that leads to the linear system being always regular is that

the basis functions must depend on the interpolation data!

$$\boldsymbol{v}^{\star}(t,\boldsymbol{x}) = \sum_{s=1}^{m} \boldsymbol{c}_{s} \cdot \boldsymbol{\psi}_{s}(\boldsymbol{x})$$
(3.1)

REPEAT:

Not just the interpolation coefficients, but also the basis must depend on the data!

This, and the radial nature of the multiquadric hyperboloids, led subsequently to a flurry of developments of

Radial Basis Functions.

The beauty of all this is that the input data can be arbitrarily scattered!

2 Scattered (cluster) data interpolation

Radial Basis Functions

Radial Basis Functions

A radial basis function decomposition for a function $f({\bm x}): \mathbb{R}^d \to \mathbb{R}$ is typically defined by

$$f^{\star}(\boldsymbol{x}) = \sum_{j=1}^{m} \lambda_j \cdot \phi(|\boldsymbol{x} - \boldsymbol{x}_j|; \epsilon_j)$$
(3.2)

with $\phi(r;\epsilon)$ being one of:-

| Multiquadric [Har71] | $\sqrt{1+(\epsilon r)^2}$ |
|------------------------------|--|
| Inverse Multiquadric [Har71] | $1/\sqrt{1+(\epsilon r)^2}$ |
| Gaussian | $e^{-(\epsilon r)^2}$ |
| Thin Plate Spline | $(\epsilon r)^2 \ln(\epsilon r)$ |
| C ⁰ Matérn | $e^{- \epsilon r }$ |
| C^2 Matérn | $e^{-\sqrt{3} \epsilon r } \cdot (1 + \sqrt{3} \epsilon r)$ |
| C^4 Matérn | $e^{-\sqrt{5} \epsilon r } \cdot (1 + \sqrt{5} \epsilon r + 5(\epsilon r)^2/3)$ |

```
Peter Jäckel
```

Cluster Induction

```
13 / 117
```

2 Scattered (cluster) data interpolation

Radial Basis Functions

Anyone coming from a background of conventional finite differencing or finite elements, might be forgiven to expect something like this to be the result of the decomposition (seen as a 1D interpolation here):



May 2018

May 2018

So you may be forgiven to be somewhat suprised that the so praised *multi-quadric* basis is the complete opposite of a what a locally confined function would look like:



Scattered data interpolation over n cluster points is (usually) accomplished by centering n basis functions in the cluster data locations themselves, and solving the linear system

$$\Phi \cdot \boldsymbol{\lambda} = \boldsymbol{f} \tag{3.3}$$

with

$$(\Phi)_{ij} = \phi(|\boldsymbol{x}_i - \boldsymbol{x}_j|; \epsilon_j)$$
(3.4)

Multiquadric Φ matrix ($\epsilon = 2.6$) for 15 Gaussian Sobol' nodes in 1D:



2 Scattered (cluster) data interpolation

Radial Basis Functions

Multiquadric Φ^{-1} matrix for 15 Gaussian Sobol' nodes in 1D:



The spectrum of the multiquadric Φ matrix:



Peter Jäckel

2 Scattered (cluster) data interpolation

Radial Basis Functions

Cluster Induction

The multiquadric spectrum

May 2018



The first three harmonic modes of Φ and their eigenvalues:

19 / 117



Peter Jäckel

2 Scattered (cluster) data interpolation

Cluster Induction

Radial Basis Functions

The multiquadric spectrum



And all other modes of Φ in between:

One of the amazing properties of multiquadrics (and inverse multiquadrics) is that, under certain technical conditions on f(x), in one dimension, RBF-interpolation on a regular grid with node distance h converges to f(x)

exponentially.

That's

$$|f(x) - f^*(x)| \sim \mathcal{O}\left(e^{-c/h}\right)$$
(3.5)

which is obviously much faster than any power convergence where

for any p.

$$|f(x) - f^*(x)| \sim \mathcal{O}(h^p)$$
 (3.6)

This result was derived more than twenty years after Hardy picked them out as the best practical choice for scattered data interpolation!

What's more, "Madych and Nelson showed that for the space of conditionally positive definite functions to which MQ belongs, a semi-norm exists and is minimized by such functions" [Kan90b].

These findings finally justified why in practice they tend to be the best choice (and this also holds for higher dimensions).

Peter Jäckel
 Cluster Induction
 May 2018
 23 / 117

 2 Scattered (cluster) data interpolation
 Radial Basis Functions
 Facts worth knowing about RBFs

 RBF-interpolation converges to polynomial interpolation as
$$\epsilon \rightarrow 0$$
.
 $\epsilon = 10$:
 0.

 6
 0.1
 0.0
 0.1

 6
 0.05
 0.05
 0.05

 7
 0.05
 0.05
 0.05

 8
 0.05
 0.05
 0.05

 9
 0.05
 0.05
 0.05

 9
 0.05
 0.05
 0.05

 9
 0.05
 0.05
 0.05

 9
 0.05
 0.05
 0.05

 9
 0.05
 0.05
 0.05

 9
 0.05
 0.05
 0.05

 9
 0.05
 0.05
 0.05

 9
 0.05
 0.05
 0.05
 0.05

 9
 0.05
 0.05
 0.05
 0.05

 9
 0.05
 0.05
 0.05
 0.05

 9
 0.05
 0.05
 0.05
 0.05

Reciprocal condition number $\kappa(\Phi)^{-1} = 0.0152$.

⊥ -0.2 10



| Peter Jäckel | Cluster Induction | May 2018 | 25 / 117 |
|--|------------------------|-------------------|---------------|
| | | | |
| | | | |
| 2 Scattered (cluster) data interpolation | Radial Basis Functions | Facts worth knowi | ng about RBFs |





| Peter Jäckel | Cluster Induction | May 2018 | 27 / 117 |
|--------------|-------------------|----------|----------|
| | | | |

```
2 Scattered (cluster) data interpolation
```

Radial Basis Functions

Facts worth knowing about RBFs





| Peter Jäckel | Cluster Induction | May 2018 | 29 / 117 |
|--|------------------------|-------------------|---------------|
| | | | |
| | | | |
| 2 Scattered (cluster) data interpolation | Radial Basis Functions | Facts worth known | ng about RBFs |

- The limit of the *shape parameter* $\epsilon \to 0$ is also referred to as the *flat limit*.
- We saw how the interpolation becomes ragged as we approach this limit.
- Once $\kappa(\Phi)^{-1}$ approaches or drops below the machine's floating point resolution, the solution incurs more and more truncation-induced noise.
- Once again, *Subtractive Cancellation* raises its ugly head...



The Moore-Penrose (SVD) solution must be used.
 Use Lapack's "Divide-and-Conquer" algorithm (DGELSD).

Meshless induction

Obviously, it is possible to use Radial Basis Functions for the purpose of a Monte Carlo regression, to evaluate any callability or future conditional value, such as for the XVA valuation in (2.2).

However, that's not what this presentation is about.

Instead, we try to go one step further.

Wouldn't it be nice if we could not only interpolate over scattered data, but directly use such a *cluster of scattered nodes* as base locations for a backward induction of a partial differential equation?

Kansa [Kan90a] demonstrated that, apart from serving well to interpolate a value surface over scattered data, radial basis function decompositions are also useful to approximate partial derivatives of the value surface.

Armed with this result on the partial derivatives, Kansa [Kan90b] then proceeded as follows.

Given an advection-diffusion (forward) PDE

$$\partial_t f + u \cdot \partial_x f - D \cdot \partial_x^2 f = 0 \tag{4.1}$$

substitute the decomposition

Meshless induction

$$f(x) = \boldsymbol{\phi}(x)^{\top} \cdot \boldsymbol{\lambda} \tag{4.2}$$

for f, where $\phi(x)$ is a vector of radial basis functions

centered in an arbitrarily scattered cluster of vertices.

Changing notation to a generic spatial differential operator \mathcal{L} in d dimensions, we obtain

$$\boldsymbol{\phi}(\boldsymbol{x})^{\top} \cdot \dot{\boldsymbol{\lambda}}(t) + \left(\mathcal{L} \cdot \boldsymbol{\phi}\left(\boldsymbol{x}\right)\right)^{\top} \cdot \boldsymbol{\lambda}(t) = 0.$$
(4.3)

for the backward Kolmogorov (pricing) equation.

Note that \mathcal{L} applies **analytically** to the individual basis functions $\phi_i(x)$!

Denoting the result of the analytical evaluation of ${\cal L}$ applied to $\phi({m x})$ as

$$\mathcal{L}\phi(\boldsymbol{x}) := \mathcal{L} \cdot \phi(\boldsymbol{x})$$
 (4.4)

we now have

$$\boldsymbol{\phi}(\boldsymbol{x})^{\top} \cdot \dot{\boldsymbol{\lambda}}(t) + \mathcal{L}\boldsymbol{\phi}(\boldsymbol{x})^{\top} \cdot \boldsymbol{\lambda}(t) = 0, \qquad (4.5)$$

i.e., a

"clustered" ordinary differential equation in $\lambda(t)$.

Note that x takes on the role of a parameter vector!

| Peter Jäckel | Cluster Induction | May 2018 | 33 / 117 |
|--------------|-------------------|----------|----------|
| | | | |
| | | | |

3 Meshless induction

In order to meet Micchelli's non-singularity criterion, we ask for the cluster ODE to hold for x to be in any of the cluster vertices $\{x_i\} \forall i = 1, ..., n$ where the radial basis functions are centered.

Defining the matrices Φ and $\mathcal{L}\Phi$ whose elements are

$$(\Phi)_{ij} = \phi(|\boldsymbol{x}_i - \boldsymbol{x}_j|; \epsilon_j)$$
(4.6)

$$(\mathcal{L}\Phi)_{ij} = \mathcal{L} \cdot \phi(|\boldsymbol{x} - \boldsymbol{x}_j|; \epsilon_j)|_{\boldsymbol{x} = \boldsymbol{x}_i} , \qquad (4.7)$$

we obtain the cluster ODE system

$$\Phi \cdot \dot{\boldsymbol{\lambda}}(t) + \mathcal{L}\Phi \cdot \boldsymbol{\lambda}(t) = 0.$$
(4.8)

REMARK. It is in principle possible to have more or fewer basis functions than locations at which we ask for the original cluster ODE (4.5) to hold, or to choose locations different from the basis centres. Either way, we will then demand for the system (4.8) to be best matched in a least squares sense, naturally leading to the use of Singular Value Decomposition for its solution.

3 Meshless induction

The tedious bit

Crunching out the analytics...

In a practical implementation, it is useful to define the cluster vertices in terms of some 'cluster coordinate system' that scales away any relative magnitude differences of the individual state space coordinates, i.e., the elements of the state vector $x \in \mathbb{R}^d$.

Consider then a given coordinate-scaling diagonal matrix Ω such that we have the transformation from the cluster coordinates $\boldsymbol{\xi}$ to state space coordinates \boldsymbol{x} via

$$\boldsymbol{x} = \Omega \cdot \boldsymbol{\xi} \tag{4.9}$$

Obviously, the same scaling applies to the cluster vertices, i.e., the basis function centres:

$$\boldsymbol{x}_i = \Omega \cdot \boldsymbol{\xi}_i \tag{4.10}$$

for i = 1, .., n for n cluster points.

| Peter Jäckel | Cluster Induction | May 2018 | 35 / 117 |
|----------------------|-------------------|----------|----------|
| | | | |
| | | | |
| | | | |
| 3 Meshless induction | I he tedious hit | | |

We then only ever evaluate all basis functions in terms of cluster coordinates, i.e. the *j*-th basis function as a function of \boldsymbol{x} evaluates according to

$$\phi_j(\boldsymbol{x}) := \phi(|\boldsymbol{\xi} - \boldsymbol{\xi}_j|, \epsilon_j) \tag{4.11}$$

with

$$\boldsymbol{\xi} = \frac{\boldsymbol{x}}{\Omega} \tag{4.12}$$

where $\phi(r; \epsilon)$ is one of the fundamental radial basis functions, e.g., a multiquadric (which is really a hyperbola and not a 'quadr[at]ic' function).

For the evaluation of the operator ${\cal L}$ applied to the radial basis functions, we note then

$$\frac{\partial}{\partial \boldsymbol{x}_{k}} = \frac{1}{\omega_{k}} \cdot \frac{\partial}{\partial \xi_{k}}$$
(4.13)

for k = 1, .., d.

For a generic advection-diffusion generator with advection coefficients that are linear in the state variables, we obtain

$$\mathcal{L} = \frac{1}{2} \sum_{k,l=1}^{d} \sigma_k \cdot \rho_{kl} \cdot \sigma_l \cdot \frac{\partial^2}{\partial x_k \partial x_l} + \sum_{k,l=1}^{d} x_k \cdot \varkappa_{kl} \cdot \frac{\partial}{\partial x_l}$$
(4.14)

$$= \frac{1}{2} \sum_{k,l=1}^{d} \rho'_{kl} \cdot \frac{\partial^2}{\partial \xi_k \partial \xi_l} + \sum_{k,l=1}^{d} \xi_k \cdot \varkappa'_{kl} \cdot \frac{\partial}{\partial \xi_l}$$
(4.15)

with

$$\rho_{kl}' := \frac{\sigma_k}{\omega_k} \cdot \rho_{kl} \cdot \frac{\sigma_l}{\omega_l}$$
(4.16)

$$\varkappa'_{kl} := \frac{\omega_k}{\omega_l} \cdot \varkappa_{kl}$$
 (4.17)

though we note that we will drop the primes in the following.

| Peter Jäckel | Cluster Induction | May 2018 | 37 / 117 |
|----------------------|-------------------|----------|----------|
| | | | |
| | | | |
| 3 Meshless induction | The tedious bit | | |

After a considerable amount of tedious calculations, we arrive at the following condensed expressions for \mathcal{L} applied to $\phi(|\Delta \boldsymbol{\xi}|, \epsilon)$, with $\Delta \boldsymbol{\xi} := \boldsymbol{\xi} - \boldsymbol{\xi}^*$, i.e., a radial basis function centered in some cluster node $\boldsymbol{\xi}^*$:-

$$\mathcal{L}\phi = \frac{1}{2}(\tilde{d} \cdot \alpha + \beta \cdot \gamma) + \alpha \cdot \chi_1 \tag{4.18}$$

$$\mathcal{L}^{2}\phi = \frac{1}{4} \cdot \left(\eta \cdot \gamma^{2} + 2 \cdot \zeta \cdot (\tilde{d} \cdot \gamma + 2 \cdot \theta) + \beta \cdot (\tilde{d}^{2} + 2 \cdot \iota) \right)$$

$$+ \frac{1}{2} \cdot \left(2 \cdot \alpha \cdot \tilde{d}_{2} + \beta \cdot \tilde{d} \cdot \chi_{1} + 2 \cdot \beta \cdot (\chi_{4} + \chi_{5}) + \zeta \cdot \gamma \cdot \chi_{1} \right)$$

$$+ \frac{1}{2} \cdot \left((\zeta \cdot \gamma + \beta \cdot \tilde{d}) \cdot \chi_{1} + 2 \cdot \beta \cdot \chi_{5} \right)$$

$$+ \alpha \cdot (\chi_{2} + \chi_{3}) + \beta \cdot \chi_{1}^{2}$$

$$(4.19)$$

where, with $(R)_{kl} := \rho_{kl}$, $(K)_{kl} := \varkappa_{kl}$, and $r := |\Delta \boldsymbol{\xi}|$,

$$\alpha = \frac{\phi'}{r} \qquad \gamma = \Delta \xi^{\top} \cdot R \cdot \Delta \xi$$

$$\beta = \frac{1}{r^2} \left(\phi'' - \frac{\phi'}{r} \right) \qquad \theta = \Delta \xi^{\top} \cdot R^2 \cdot \Delta \xi$$

$$\zeta = \frac{1}{r^3} \left[\phi''' - \frac{3}{r} (\phi'' - \frac{\phi'}{r}) \right] \qquad \chi_1 = \xi^{\top} \cdot K \cdot \Delta \xi$$

$$\eta = \frac{1}{r^4} \left[\phi'''' - 6\frac{\phi'''}{r} + \frac{15}{r^2} (\phi'' - \frac{\phi'}{r}) \right] \qquad \chi_2 = \xi^{\top} \cdot K \cdot K^{\top} \cdot \xi$$

$$\tilde{d} = \operatorname{Tr}(R) \qquad \chi_3 = \xi^{\top} \cdot K^2 \cdot \Delta \xi$$

$$\iota = \operatorname{Tr}(R^2) \qquad \chi_4 = \Delta \xi^{\top} \cdot R \cdot K \cdot \Delta \xi$$

$$\tilde{d}_2 = \operatorname{Tr}(R \cdot K) \qquad \chi_5 = \xi^{\top} \cdot K \cdot R \cdot \Delta \xi \qquad (4.20)$$

Only
$$\alpha$$
, β , ζ , and η depend on the type of RBF!

| Peter Jäckel | Cluster Induction | May 2018 | 39 / 117 |
|----------------------|-------------------|----------|----------|
| | | | |
| 3 Meshless induction | The tedious bit | | |
| | | | |

Further, we obtain:-

whenwhen
$$\phi = \sqrt{1 + (\epsilon r)^2}$$
: $\phi = e^{-(\epsilon r)^2}$: $\alpha = \frac{\epsilon^2}{\phi}$ $\alpha = -2 \cdot \epsilon^2 \cdot \phi$ $\beta = -\alpha \cdot \frac{\epsilon^2}{\phi^2}$ $\beta = -2 \cdot \epsilon^2 \cdot \alpha$ $\zeta = -3 \cdot \beta \cdot \frac{\epsilon^2}{\phi^2}$ $\zeta = -2 \cdot \epsilon^2 \cdot \beta$ $\eta = -5 \cdot \zeta \cdot \frac{\epsilon^2}{\phi^2}$ $\eta = -2 \cdot \epsilon^2 \cdot \zeta$

Before we proceed to numerical integration scheemes of our system of ODEs, we make a few observations.

If we assume the existence of Φ^{-1} , the system (4.8) can be written as

$$\dot{\boldsymbol{\lambda}}(t) + \left(\Phi^{-1} \cdot \mathcal{L}\Phi\right) \cdot \boldsymbol{\lambda}(t) = 0 \qquad (4.22)$$

with formal solution

$$\boldsymbol{\lambda}(t - \Delta t) = e^{\Delta t \cdot \left(\Phi^{-1} \cdot \mathcal{L}\Phi\right)} \cdot \boldsymbol{\lambda}(t) . \qquad (4.23)$$



Equally, if we use

$$\Phi \cdot \boldsymbol{\lambda} = \boldsymbol{f} \qquad \iff \qquad \boldsymbol{\lambda} = \Phi^{-1} \cdot \boldsymbol{f} \qquad (4.24)$$

and define $G := \mathcal{L}\Phi \cdot \Phi^{-1}$, we have

$$\dot{f}(t) + G \cdot f(t) = 0.$$
 (4.25)

Note that (4.25) is a system of ODEs directly for the value vector **f**! This means that we can in principle carry out a backward induction directly on a set of arbitrarily scattered cluster vertices, without ever computing any form of local (finite-differencing or otherwise) partial derivatives from the value function!

Finally, unsurprisingly, the formal solution of (4.25) is

$$\boldsymbol{f}(t - \Delta t) = e^{\Delta t \cdot G} \cdot \boldsymbol{f}(t) . \qquad (4.26)$$

Remark: in practice, due to the need for SVD with cut-off, the use of G induces an unnecessary early projection onto the subspace of functions that can be represented by the basis.

Peter Jäckel

Some observations

 $\mathcal{L}\Phi$ matrix for 15 Gaussian Sobol' nodes in 1D (for a standard diffusion):



3 Meshless induction

Some observations



The generator matrix $G = \mathcal{L}\Phi \cdot \Phi^{-1}$:



This looks like the typical Viking ship we expect for a proper transition probability matrix!!!

| Peter Jäckel | Cluster Induction | May 2018 | 45 / 117 |
|----------------------|-------------------|--------------------------|----------|
| | | | |
| 3 Meshless induction | Some observations | The generator's spectrum | |

The spectrum of the multiquadric generator matrix $G = \mathcal{L}\Phi \cdot \Phi^{-1}$:





Peter Jäckel



^{4 The truth is implicit} There's more to life than Crank-Nicolson

Despite there being all the mentioned possible ways to transform the ODE

$$\Phi \cdot \dot{\boldsymbol{\lambda}}(t) + \mathcal{L} \Phi \cdot \boldsymbol{\lambda}(t) = 0 , \qquad (4.8)$$

numerical experience suggests it is best to proceed with it directly as it is! As a first order implicit scheme:

$$(\Phi - \Delta t \cdot \mathcal{L}\Phi) \cdot \boldsymbol{\lambda}(t - \Delta t) = \Phi \cdot \boldsymbol{\lambda}(t) .$$
(5.1)

Note that since $f(t) = \Phi \cdot \lambda(t)$, we can always evaluate the right hand side directly from the respectively given terminal boundary conditions, i.e., the (final) pay-off (where applicable):

$$(\Phi - \Delta t \cdot \mathcal{L}\Phi) \cdot \boldsymbol{\lambda}(T - \Delta t) = \boldsymbol{f}(T) .$$
 (5.2)

117

This avoids the projection of any pay-off value functions onto the RBF basis!

| Peter Jäckel | Cluster Induction | May 2018 | 5(|
|--------------|-------------------|----------|----|

This may not seem like much of an advantage since this projection, by its nature of being an exact interpolation, ought to return the exact pay-off values at the cluster vertices.

However, since some trades may be long-dated, the consistent valuation of all the legs of a portfolio may involve laying out a calculation for tens of years, with some of the trades in the portfolio being only short dated.

If we fix the cluster-to-state-space scaling matrix Ω to suit the longest trade, the short-dated components can incur valuation over a set of (state space) nodes that are effectively *many* standard deviations away from the money.



It is for this reason that it can be beneficial to use a different $\Omega_q := \Omega(t_q)$ for each time horizon t_q , q = 0, ..., m for m time steps from the valuation date to the last observation date.

Since we always use the same (de-scaled) cluster coordinate vertices $\{\xi_i\}$ for i = 1, ..., n, and since, typically, we choose

$$\Omega_{q-1} < \Omega_q \tag{5.3}$$

in all of its (diagonal) elements. This means that our cluster, in state space coordinates, typically *contracts* as we carry out the backward induction:

$$\boldsymbol{x}_{i}(t_{q}) = \Omega_{q} \cdot \boldsymbol{\xi}_{i} \qquad \Longrightarrow \qquad |\boldsymbol{x}_{i}(t_{q-1})| < |\boldsymbol{x}_{i}(t_{q})| .$$
 (5.4)

Note that this has no impact on the elements of the matrix Φ since it is fully evaluated simply on cluster coordinates, but it does mean that $\mathcal{L}\Phi$ varies from one time step to the next due to its dependence on Ω , i.e.,

$$\mathcal{L}\Phi_q := \mathcal{L}\Phi(t_q) . \tag{5.5}$$

4 The truth is implicit

Cluster contraction

Having rolled back to some time horizon t_q , i.e., having computed

$$\boldsymbol{\lambda}_q := \boldsymbol{\lambda}(t_q) , \qquad (5.6)$$

and taking into account additional payoff contributions happening at t_q denoted as $\pi_q(\boldsymbol{x}, t_q)$, the next standard fully implicit step backwards to t_{q-1} is then

$$(\Phi - \Delta t \cdot \mathcal{L}\Phi_{q-1}) \cdot \boldsymbol{\lambda}_{q-1} = \Phi_{\langle q-1,1 \rangle} \cdot \boldsymbol{\lambda}_{q} + \boldsymbol{\pi}_{q}$$
(5.7)

with the matrix $\Phi_{\langle q-1,1\rangle}$ defined by

$$\left(\Phi_{\langle q-1,1 \rangle} \right)_{ij} = \phi(|(\boldsymbol{x}_i(t_{q-1}) - \boldsymbol{x}_j(t_q))/\Omega_q|; \epsilon_j)$$

$$= \phi(|\frac{\Omega_{q-1}}{\Omega_q} \cdot \boldsymbol{\xi}_i - \boldsymbol{\xi}_j|; \epsilon_j)$$
(5.8)

and the elements of the vector π_q given by

$$(\boldsymbol{\pi}_{\boldsymbol{q}})_i = \boldsymbol{\pi}_{\boldsymbol{q}}(\boldsymbol{x}_i(t_{q-1}), t_q) = \boldsymbol{\pi}_{\boldsymbol{q}}(\Omega_{q-1} \cdot \boldsymbol{\xi}_i, t_q)$$
(5.9)

Note that we avoid the projection of the extra payment π_q onto the basis!

| Peter Jäckel | Cluster Induction | May 2018 | 53 / 117 |
|--|-------------------------|----------|-------------------|
| | | | |
| 4 The truth is implicit | Cluster contraction | | |
| To illustrate the basis fur | nction evaluation logic | in (5.8) | |
| x^{\uparrow} in state space coordinates the space coordinates x^{\uparrow} | ates: | | •••••• |
| | | | |
| | | * | |
| | | | _ |
| | _ | + | |
| · · · · · · · · · · · · · · · · · · · | | | |
| | | * | \xrightarrow{t} |
| | | | |
| | + | * | |
| + | | + | + |
| $\cdots \sim \Omega(t)$ | | * | |
| + $oldsymbol{x}_i(t)$ | | | |
| $\mathbf{x} \mathbf{x}_{i}(t_{q-1})$ projected t | o t_q | T | |
| 1 | t_{a-1} | t_a | T |



4 The truth is implicit

A word on stability

To analyze, say, the stability of the first order implicit scheme

$$(\Phi - \Delta t \cdot \mathcal{L}\Phi_{q-1}) \cdot \boldsymbol{\lambda}_{q-1} = \Phi \cdot \boldsymbol{\lambda}_q$$
 (5.10)

we investigate the spectrum of

$$(1 - \Delta t \cdot \tilde{G})^{-1} \tag{5.11}$$

where

$$\tilde{G}_{q-1} := \Phi^{-1} \cdot \mathcal{L}\Phi_{q-1}$$
(5.12)

which is the transpose of G_{q-1} when all ϵ_j are the same (and thus $\Phi = \Phi^{\top}$). In comparison, to analyze, say, the stability of the first order explicit scheme

$$\Phi \cdot \boldsymbol{\lambda}_{q-1} = (\Phi + \Delta t \cdot \mathcal{L} \Phi_q) \cdot \boldsymbol{\lambda}_q .$$
(5.13)

we need to see the spectrum of

 $(1 + \Delta t \cdot \tilde{G}_q) . \tag{5.14}$

Note that both schemes require the solution of a dense linear system!

| Peter Jäckel | Cluster Induction | May 2018 | 56 / 117 |
|--------------|-------------------|----------|----------|



Given that the (first order) *explicit* and *implicit* schemes *both* require the solution of a dense linear system, the advantage of the explicit scheme reduces to the fact that we can pre-factorise Φ once for all time steps.

However, caution should be applied for the first step backward from any payment date, i.e., from any point in time that adds a time-inhomogeneous source term.

Since financial contingent payment formulæ often contain non-differentiable terms, attempting a fully explicit step straight out of a payment date amounts to a projection of the payoff onto the radial basis, and that rarely works well!

All in all, implicit schemes appear to be generally the better choice.

Since the Crank-Nicolson scheme contains an explicit component, we equally (generally) don't favour it for higher order convergence in Δt .

Instead, the following schemes have been found to be of merit:-

• BDF2. This is a second-order fully implicit *stiff* scheme [CH52, Gea71] originally for the forward ODE

$$\dot{y} = f(y) \tag{5.15}$$

given by

$$y(t_{n+2}) = \frac{4}{3}y(t_{n+1}) - \frac{1}{3}y(t_n) + \frac{2}{3} \cdot \tau \cdot f(y(t_{n+2}))$$
 (5.16)

with time steps of length τ . In this context here, this becomes

$$(\Phi - \frac{2}{3}\Delta t \cdot \mathcal{L}\Phi_q) \cdot \boldsymbol{\lambda}_q = \frac{4}{3} \cdot \Phi \cdot \boldsymbol{\lambda}_{q+1} - \frac{1}{3} \cdot \underbrace{\Phi \cdot \boldsymbol{\lambda}_{q+2}}_{\boldsymbol{f}(t_{q+2})}.$$
(5.17)

This scheme requires a start-up.

When started with Crank-Nicolson, it is also referred to as TR-BDF2.

| Peter Jäckel | Cluster Induction | May 2018 | 59 / 117 |
|--------------|-------------------|----------|----------|
| | | | |
| | | | |

4 The truth is implicit Second order schemes

• Padé(0,2) is based on the (0,2)-Padé expansion $e^t \approx \frac{1}{1-t+t^2/2}$. It is a second-order fully implicit scheme, sometimes (though rarely) also referred to as *second order backwards Euler* or simply as

second order implicit.

In this context here, we have

$$(\Phi - \Delta t \cdot \mathcal{L}\Phi_q + \frac{1}{2}\Delta t^2 \cdot \mathcal{L}^2\Phi_q) \cdot \boldsymbol{\lambda}_q = \underbrace{\Phi \cdot \boldsymbol{\lambda}_{q+1}}_{\boldsymbol{f}(t_{q+1})}.$$
(5.18)

NOTE that $\mathcal{L}^2 \Phi$ is the matrix given by the analytical evaluation of $\mathcal{L}^2 \cdot \phi$, i.e., its elements are

$$\left(\mathcal{L}^{2}\Phi\right)_{ij} = \mathcal{L}^{2} \cdot \phi(|\boldsymbol{x} - \boldsymbol{x}_{j}|)|_{\boldsymbol{x} = \boldsymbol{x}_{i}}$$
 (5.19)

The matrix $\mathcal{L}^2 \Phi$ is **not** the same as $(\mathcal{L} \Phi)^2$!

Details for the calculation of $\mathcal{L}^2\Phi$ were given in equation (4.19) on page 39.

The second order implicit method also serves well to start-up the BDF2 scheme.

4 The truth is implicit

Some results

We look at a 1Y ATM vanilla FX option with 20% volatility.

We show its price convergence to the exact value as a function of:-

- the number of time steps n_T for fixed n_{cluster} , and
- the number of cluster nodes n_{cluster} fixed n_T for,

priced with Black-Scholes-Merton model and a cross-currency Hull-White (aka 'Linear Gaussian Markov') model with normal rates volatilites of 200bps and mean reversion speeds of 5% for USD and 10% else.

The cluster nodes are all from a Sobol' sequence mapped to normal variates.







| Peter Jäckel | Cluster Induction | May 2018 | 63 / 117 |
|--------------|-------------------|----------|----------|
| | | | |
| | | | |

4 The truth is implicit

```
Some results
```

The detachment effect of the long range growth mode of multiquadrics:

option value of 1Y call ($\sigma = 20\%$, K = 1.12)



Peter Jäckel

The decay effect of Gaussian RBFs (1Y ATM call):

Analytical 6 --- Gaussian RBF cluster induction 3 2 1 $15 \ln(S/K)$ -1.5 -0.5 0.5 -1 -Þ



| Peter Jäckel | Cluster Induction | May 2018 | 65 / 117 |
|-------------------------|-------------------|----------|----------|
| | | | |
| | | | |
| 4 The truth is implicit | Some results | | |



| er Jäckel | Cluster Induction | M |
|-----------|-------------------|---|
| | | |

ay 2018







Peter Jäckel



Some results

Comparison with exact analytical values (in gray):



| Peter Jäckel | Cluster Induction | May 2018 | 69 / 117 |
|--------------|-------------------|----------|----------|
| | | | |

4 The truth is implicit

Some results



 $\log_{10}(|\mathrm{rel.error}|)$



Comparison with exact analytical values (in gray):



4 The truth is implicit

Some results



ckel Cluster Induction May 2018



 Peter Jäckel
 Cluster Induction
 May 2018
 75 / 117

 4 The truth is implicit
 Some results

1Y composite option with Linear Gaussian Markov model and $n_T = 16$.



| Jäckel | Cluster Induction |
|--------|-------------------|
| lackel | Cluster Induction |

Peter

Implementation and optimisations

Like with Libor market models, the *real* work starts with the implementation:-

• How do we choose the shape parameters ϵ_j ?

There is some research on this subject of variable practical applicability.

• How do we generate the cluster?

There is little in the literature.

- Generic optimisations, e.g.:
 - Pre-factorise matrices where possible and multiple (non-simultaneous) solutions with different right-hand-sides are required.
 - Invoke Lapack's "solve-with-SVD" routine instead of "give-me-the-full-SVD" when you don't actually need the SVD itself.

| | IVIAY 2018 | 77 / 117 |
|---------------------------------------|-------------------------------------|-------------------------------------|
| | | |
| out do us shooss the shops powerster? | | |
| 0 | w do we choose the shape parameter? | w do we choose the shape parameter? |

- When ϵ is too large, the approximation diverges: the RBFs are then too localised to genuinely *interpolate* between cluster nodes.
- When ϵ is too small, the linear system starts to suffer from subtractive cancellation and becomes numerically unsolvable. The system's matrix condition number diverges.
- The sweet spot tends to be when the reciprocal matrix condition number *approaches* machine precision.

5 Implementation and optimisations How do we choose the shape parameter?

On the influence of the shape parameter on the accuracy, e.g., see:-

- figures 2 and 3 in [CS09].
- figures 11 and 13 in [Mil14].
- figures 2.3 and 7.4 in [FZ07].
- figures 5.1, 5.2, 5.3, 7.1, 7.2, and 7.3 in [FM12].
- figures 22, 26, and 27 in [Fas11].
- figure 3.1 b) and 4.2 in [FF15b].
- figure 2 in [SL16].
- figure 5.2 in [Wan14].
- figures 2, 4 and 6 in [LF03].
- figures 4-8, 10-12, and 14-17 in [Mon11] (my favourite).



The optimum shape parameter ϵ tends to be near the point where the kernel matrix Φ becomes singular, i.e., where $\kappa(\Phi)^{-1}\gtrsim \texttt{DBL}_\texttt{EPSILON}.$

5 Implementation and optimisations How do we choose the shape parameter?

- Various methods have been suggested for (semi-)automatic choice algorithms for ϵ such as 'Leave-one-out-cross-validation' (LOOCV), but, for our industrial application purposes, none of them appear to be anywhere near robust enough for commercial use.
- I found that choosing each \(\ell_j\), to some extent, with a global cap and floor, based on the reciprocal distance to its near(est) neighbour(s) gives a significant improvement.
- Interestingly, there is disappointingly little discussion of this observation of mine in the literature, (arguably) the sole noteworthy exception being "*The Runge phenomenon and spatially variable shape parameters in RBF interpolation*" [FZ07]. Quote:

The conditioning of the A-matrix is likely to be greatly improved when using spatially variable ϵ_k .

• The conditioning of the A-matrix is not the only benefit, but I am glad I found *at least some literature* in support of my findings.

An unexplored idea:-

• Instead of merely demanding that the equation

$$\boldsymbol{\phi}(\boldsymbol{x})^{\top} \cdot \dot{\boldsymbol{\lambda}}(t) + \mathcal{L}\boldsymbol{\phi}(\boldsymbol{x})^{\top} \cdot \boldsymbol{\lambda}(t) = 0 , \qquad (4.5)$$

holds in all of the radial basis function centres, to arrive at the ODE

$$\Phi \cdot \dot{\boldsymbol{\lambda}}(t) + \mathcal{L} \Phi \cdot \boldsymbol{\lambda}(t) = 0, \qquad (4.8)$$

with square matrix Φ ,

• we *oversample* equation (4.5), i.e., ask for it to hold (as best as possible) *in an additional set of* (somehow, arbitrarily) *chosen locations*.

5 Implementation and optimisations How do we choose the shape parameter?

• Applying this to whichever chosen specific time-integration scheme, this results in an overdetermined system of the form

$$A(\epsilon) \cdot \boldsymbol{\lambda} = \boldsymbol{r} . \tag{6.1}$$

• Next, we demand that the RMS error of this overdetermined system is minimized over both λ and ϵ , which gives us the non-linear system

$$A^{\top} \cdot A \cdot \boldsymbol{\lambda} - A^{\top} \cdot \boldsymbol{r} = 0$$

$$\boldsymbol{\lambda}^{\top} \cdot A^{\top} \cdot A' \cdot \boldsymbol{\lambda} - \boldsymbol{r}^{\top} \cdot A' \cdot \boldsymbol{\lambda} = 0$$
(6.2)

where

$$A' = \frac{\mathrm{d}A(\epsilon)}{\mathrm{d}\epsilon}$$

| Peter Jäckel | Cluster Induction | May 2018 | 83 / 117 |
|------------------------------------|---|----------|----------|
| | | | |
| E Implementation and antimications | How do we choose the choose personator? | | |

- Solving this non-linear system is obviously some numerical effort.
- And it may require safeguarding against its Jacobian becoming singular!
- One may not want to do this for each and every time step.
- One could attempt to solve it at every payment date, to capture the contribution of each payoff,
- and freeze ϵ over subsequent (regular) time steps where we'd only use the regular, i.e., square, matrix Φ .

To be tested...

How do we generate the cluster?

- Sadly, most research is done in just 1 dimension ! (!!!)
- There is little in the literature for more than 2 dimensions, up to which researchers tend to use fairly regular point distributions.
- A transfer of ideas from finite element methods is appealing.
- Only that little is out there for more than 3 dimensions, and effective (adaptive) tesselations for FEM tend to be specialised for 2 dimensions.

Quote [FF15a]:

In 2-D: Quick to go from quasi-uniform nodes to well-balanced Delaunay triangularization.

In 3-D: Finding good tetrahedral sets can even become a dominant cost.

| Peter Jäckel | Cluster Induction | May 2018 | 85 / 117 |
|------------------------------------|--------------------|----------------------------|----------|
| | | | |
| 5 Implementation and optimisations | Cluster generation | Prof. Sobol' to the rescue | |
| | | | |

A pragmatic approach:

• In *d* dimensions, draw Sobol' vectors and transform them from uniform to Gaussian coordinates via the inverse cumulative normal function.

Gaussian Sobol' vectors

- We use these as our standardised ξ -coordinate cluster nodes.
- If you wish, combine with further cluster improvement techniques...

In analogy to conventional (explicit) trees and (multivariate) quadrature methods [BW04, Jäc05], we may wish to avoid having any nodes at ξ -coordinates that in a one-dimensional standard normal distribution we would describe as lying in the tail beyond a certain number of standard deviations, say n standard deviations.

Since a one-dimensional normal distribution has two tails, this gives us that we demand

$$|\boldsymbol{\xi}| \leq r^* \tag{6.3}$$

with

$$1 - \chi^2(r^{*2}; d) = 2 \cdot \Phi(-n)$$
(6.4)

where $\Phi(\cdot)$ is here the normal distribution and $\chi^2(\cdot; d)$ the distribution of the squares of d independent standard normal variates¹.

 \Rightarrow We simply drop all drawn Gaussian Sobol' vectors that lie outside r^* .

¹The inverse χ^2 function can be evaluated as a special case of the inverse Gamma distribution function for which there is an extremely efficient implementation [DM87]. Peter Jäckel Cluster Induction May 2018 87 / 117

5 Implementation and optimisations

Cluster generation

Cluster pruning

The pruning radius r^* as a function of dimensionality dand Gaussian standard deviations n.



Latin-Hypercubing of the (perhaps pruned) cluster appears to help (in general):-

- For each dimension j from 1 to d, individually,
- sort the whole cluster by the *j*-th dimension (excluding the central node),
- and replace all the elements in the *j*-th dimension by standardised, (probability-)equidistant values (excluding the origin).

This used to be popular with pseudo-random Monte Carlo simulations before Sobol'-sampling superceded conventional Monte Carlo methods².

²Sobol' sampling has an approximate, asymptotic, version of the Latin hypercube built into it which obviates explicit Latin-Hypercubing.





Cluster generation



| Peter Jäckel | Cluster Induction | May 2018 | 91 / 117 |
|--------------|-------------------|----------|----------|
| | | | |



As we saw, even with Latin-Hypercubing we may still have an unpleasantly irregular distribution of the cluster nodes.

One way to regularise the cluster further is *Elastic Cluster Relaxation*:-

- Identify the *outer* nodes. These are the vertices of the convex hull.
 Consider these nodes, plus the origin (the *spot node*), as fixed.
- Create a suitable point-to-point set of connections to form an (irregular) web over the cluster nodes.
- Interpret the edges of the web as perfectly elastic springs, with all nodes (except those that were fixed) being only defined as vertices of the web by its topology.

• Relax the web!



- In practice, the trickiest part of this procedure is step 2.
- A good candidate for the web generation is the *Delaunay triangulation*, which happens to be the "dual" of the *Voronoi tesselation*.
- The good news is that there is published software out there for the calculation of both the convex hull and the Delaunay triangulation in arbitrary dimensions (in principle) [BDH96].
- The bad news is that the computational effort for the Delaunay triangulation grows like

 4^{d} .

- In a numerical experiment for 127 nodes, I was able to obtain the Delaunay triangulation up to 9 dimensions (though it took hours towards the end).
- In 10 dimensions, I gave up after one week.

The relaxation result is given by the minimisation of the total elastic spring energy

$$\mathcal{E} = \frac{1}{2} \sum_{(i,j)\in\mathcal{T}} |\xi_i - \xi_j|^2$$
 (6.5)

where \mathcal{T} is the set of point index tuples that comprises the Delaunay triangulation over the variation of all the mobile vertices' coordinates.

The resulting linear system for the movable nodes separates into d identical individual positive definite³ linear systems, one for each dimension.

Let's see what this does for the ugly cluster shown on page 92 !

In fact, we'll do it twice since a repetition of the Delaunay triangulation after the first relaxation is not guaranteed to result in the same topology.

³This is easy to see since the energy can be written as

$$\frac{1}{2} \left| D \cdot \Xi \right|^2 = \frac{1}{2} \Xi^\top \cdot \left(D^\top \cdot D \right) \cdot \Xi$$
(6.6)

where D is a matrix whose rows each represent the distance calculation for one triangulation edge and Ξ is the matrix of the cluster coordinates.







Cluster generation

Elastic Cluster Relaxation



Peter Jäckel

Cluster Induction

May 2018

97 / 117

5 Implementation and optimisations

Cluster generation

Elastic Cluster Relaxation



Peter Jäckel



Cluster generation

Elastic Cluster Relaxation



Consider that each point $oldsymbol{x} \in \mathbb{R}^d$ is represented by its nearest cluster node

$$\boldsymbol{x} \rightarrow \boldsymbol{\xi}_{\hat{k}(\boldsymbol{x};\mathcal{C})}$$
 (6.7)

with

$$\hat{k}(\boldsymbol{x}; \mathcal{C}) := \arg \min_{k \in \{1, \dots, N\}} |\boldsymbol{x} - \boldsymbol{\xi}_k|$$
(6.8)

where N is the number of nodes in the given cluster $C = \{\xi_1, ..., \xi_N\}$.

Then, we call the expectation

$$D(\mathcal{C}) := \mathrm{E}\left[\left|\boldsymbol{x} - \boldsymbol{\xi}_{\hat{k}(\boldsymbol{x};\mathcal{C})}\right|^2\right]$$
 (6.9)

the *average distortion* [PD51] of C's representation of \mathbb{R}^d .

Since we operate with x to represent independent asset factors, we take the expectation under a multi-variate standard normal distribution for x (with zero correlation).

An image says more than a thousand words to explain the name "*distortion*":



A classical result in statistics is that the distortion $D(\mathcal{C})$ is minimal when

$$\boldsymbol{\xi}_{k} \cdot \mathrm{E}\left[\mathbf{1}_{\left\{\hat{k}(\boldsymbol{x};\mathcal{C})=k\right\}}\right] = \mathrm{E}\left[\boldsymbol{x} \cdot \mathbf{1}_{\left\{\hat{k}(\boldsymbol{x};\mathcal{C})=k\right\}}\right]$$
(6.10)

for all $k \in \{1, .., N\}$.

This is to say that ξ_k must be equal to the first moment of x conditioned on the domain of all points in \mathbb{R}^d whose nearest cluster node is ξ_k .

We call the cluster C^* that satisfies (6.10) the *Minimum Distortion Cluster* though it is better known as the

Centroidal Voronoi Tesselation⁴.

⁴Some authors in mathematical finance have unfortunately referred to such a choice of nodes simply as a "*quantization*" which is not in line with the rest of mathematics, physics, and engineering, where a *quantization* only indicates the representation of a continuum by a subset of discrete values, without any statement about the rule that led to the choice of nodes.

Due to the considerable geometrical complications, there is in practice no competitive analytical method to compute C^* when d is more than two(-ish).

A simple algorithm attributed to Lloyd [Llo57] is to iterate

$$\boldsymbol{\xi}_{k}^{(n+1)} := \frac{\mathrm{E}\left[\boldsymbol{x}\cdot\boldsymbol{1}_{\left\{\hat{k}(\boldsymbol{x};\mathcal{C}^{(n)})=k\right\}}\right]}{\mathrm{E}\left[\boldsymbol{1}_{\left\{\hat{k}(\boldsymbol{x};\mathcal{C}^{(n)})=k\right\}}\right]} = \mathrm{E}\left[\boldsymbol{x}\left|\hat{k}(\boldsymbol{x};\mathcal{C}^{(n)})=k\right]\right].$$
(6.11)

This algorithm is also known as *Voronoi iteration* or *Voronoi relaxation*.

Again, in more than two (or so) dimensions, it is not practical to compute the conditional expectations in (6.11) (semi-)analytically whence we resort to a good old-fashioned Sobol'-Monte-Carlo evaluation⁵.

⁵Once more, unfortunately, some authors have referred to the numerical evaluation of those conditional expectations by means of a sampling method as "*stochastic*" gradient descent methods despite the fact that none of the above has anything to do with *stochasticity* or *randomness* or the concept of any [*stochastic*] process in time.

5 Implementation and optimisations

Cluster generation

Minimum Distortion Cluster

The algorithm in a nutshell:

- Precompute a d-dimensional sampling set of Gaussian Sobol'-Monte-Carlo draws, say M =32767 points, *x_i* for *i* = 1, ..., M.
- Start with a *d*-dimensional initial cluster comprised by Gaussian Sobol'-Monte-Carlo draws, say, of size N = 255.

Then, in each iteration (of, say, 127), zero-out a workspace of N vectors $\Xi_k \in \mathbb{R}^d$, and zero-out N counter variables m_k for k = 1, ..., N and:-

• Loop over all of the M sampling points \boldsymbol{x}_i for i = 1, ..., M to find the nearest (previous) cluster node $\boldsymbol{\xi}_{\hat{k}}^{(n)}$ and, having identified this nearest cluster node index $\hat{k} = \hat{k}(\boldsymbol{x}_i, \mathcal{C}^{(n)})$, set

$$\Xi_{\hat{k}} += x_i$$
$$m_{\hat{k}} += 1$$

• Upon completing the loop over the sampling points (i = 1, .., M), set

$$\boldsymbol{\xi}_k^{(n+1)} := \boldsymbol{\Xi}_k / m_k$$

for k = 1, .., N.

Notes:-

- The outer iteration (6.11) *may not converge (!)*. This is due to the discrete underlying sampling set used to evaluate the conditional expectation.
- The above mentioned non-convergence is nothing to worry about! It simply means that the algorithm eventually just cycles over a *discrete* set of equally good estimates for C^* for the given size N, dimensionality d, and Gaussian-Sobol' sampling set.
- The higher the dimensionality *d*, alas, the more outer iterations you will wish to use to attain a satisfactory cluster.
- Cache all computed clusters in memory for this run-time session since more computations are likely to want the same! Caching it mitigates the possibly considerable time it can take to create this cluster.

- A good small cluster ($N \le 255$) wins over a large bad cluster.
- You may want to grow the cluster N size gently with increasing dimensionality d.
- In higher dimensions, such optimum-representation cluster computation techniques are part of a range of *machine learning algorithms*, e.g., *"k-Means Clustering"*.
- You may find the oft-praised "Anderson acceleration" technique of little practical use.
- See Wikipedia [Llo18] for references to contemporary acceleration techniques and to links as to how these methods are also used in finite element calculations.



After 127 iterations, we arrive at:



Here is an alternative (different algorithm, \sim 3 times faster):







Peter Jäckel

A recent development is to evaluate the spatial differential operator $\mathcal L$ in

$$\partial_t f + \mathcal{L} f = 0 \tag{7.1}$$

on each cluster node *not directly*, i.e., analytically, on radial basis functions, but *as a weighted average of the function value over a judiciously chosen nearby subcluster*.

This was inspired by Bellman's *Differential Quadrature* [BKC72, BM96].

In a nutshell:

• A standard (spatial) finite-differencing method approximates, for instance, the term $\frac{\partial^2 f}{\partial x^2}$ by a weighted sum of the function value at the central node and its two nearby neighbours in the *x*-direction on a regular lattice:

$$\frac{\partial^2 f}{\partial x^2} \approx w_1 \cdot f(x - \Delta x) + w_2 \cdot f(x) + w_3 \cdot f(x + \Delta x)$$
(7.2)



- The weights are derived from a local Taylor expansion.
- With the Taylor expansion being a polynomial expansion, this is tantamount to demanding that the approximation is exact for the three test functions $f_1(x) = 1$, $f_2(x) = x$, and $f_3(x) = x^2$.
- This means the weights are derived from the linear system

$$\begin{pmatrix} f_{1}(x - \Delta x) & f_{1}(x) & f_{1}(x + \Delta x) \\ f_{2}(x - \Delta x) & f_{2}(x) & f_{2}(x + \Delta x) \\ f_{3}(x - \Delta x) & f_{3}(x) & f_{3}(x + \Delta x) \end{pmatrix} \cdot \begin{pmatrix} w_{1} \\ w_{2} \\ w_{3} \end{pmatrix} = \begin{pmatrix} f_{1}''(x) \\ f_{2}''(x) \\ f_{3}''(x) \end{pmatrix}$$
(7.3)

- Bellman's Differential Quadrature is the extension of this concept to all the nodes in the lattice line using a complete set of polynomial test functions to the maximum order attainable.
- As usual, Chebyshev polynomials turn out to be very useful.

RBF-DQ [Tol00, WS02, TS03] is the translation of this idea to:-

- focus on the full differential operator $\mathcal{L}f$ monolithically without splitting it into individual derivative terms.
- make the test functions radial basis functions centered in the nodes.

Local RBF-DQ [SDY03] then:-

- evaluates the differential operator $\mathcal{L}f$ on any one node only by weighting the function values over judiciously selected nearby *supporting nodes*.
- use only the RBFs centered in the *supporting nodes* as test functions.



Alas,

- people have started calling this localised derivative collocation technique by the name "*RBF-FD*".
- This was meant to help conjure up the notion of a local stencil for the evaluation of $\mathcal{L}f$.
- But as we clearly see, this method has nothing to do with the calculation of *"finite differences"*.
- "RBF-FD" is a misnomer just like the abhorrent term "quasi-random".

Caveat emptor.-

- All that was previously said about radial basis functions still applies!
- This includes the need for a suitably sized shape parameter,
- the need to choose among the many RBF types, and so on.
- In addition, one needs to choose the local stencil for each and every node!
- And even in 2 dimensions this may be as large as 37 nearby nodes [FF15a].

However, there is a lot of promising research on Local RBF-DQ methods!

And by virtue of its *locally dispersed support nodes* for *Lf*, it *may* even permit the use of **explicit** time integration schemes!

It remains to be seen...

THE END



Gaussian Φ^{-1} matrix:



A Spectral analysis of Gaussian RBFs

The spectrum of the Gaussian Φ matrix:





The first three harmonic modes of Φ and their eigenvalues:

Peter Jäckel

Cluster Induction

May 2018

121 / 117

A Spectral analysis of Gaussian RBFs



The Nyquist mode of Φ :



And all other modes of Φ in between:

Peter Jäckel

Cluster Induction

123 / 117

A Spectral analysis of Gaussian RBFs

 $\mathcal{L}\Phi$ for a standard diffusion:



May 2018



A Viking ship in rough water ...

The spectrum of the Gaussian generator matrix $G = \mathcal{L}\Phi \cdot \Phi^{-1}$:



Pete<mark>r Jäckel</mark>

Cluster Induction

May 2018

127 / 117

A Spectral analysis of Gaussian RBFs

The generator's spectrum







Peter Jäckel

Cluster Induction

129 / 117

May 2018

A Spectral analysis of Gaussian RBFs

The generator's spectrum



And all other modes of G in between:

And finally, the first order scheme stabilities with $\Delta t = 0.25$:



Peter Jäckel **Cluster Induction** May 2018 131 / 117

References

- [BDH96] C.B. Barber, D.P. Dobkin, and H. Huhdanpaa. The Quickhull algorithm for convex hulls. ACM Transactions on Mathematical Software, 22(4):469-483, 1996. www.qhull.org/.
- [BKC72] R. Bellman, B.G. Kashef, and J. Casti. Differential guadrature: A technique for the rapid solution of nonlinear partial differential equations. Journal of Computational Physics, 10(1):40-52, 1972.
- [BL88] D.S. Broomhead and D. Lowe. Multivariable Functional Interpolation and Adaptive Networks. Complex Systems, 2(1):321-355, 1988.
- [BM96] C.W. Bert and M. Malik. Differential quadrature method in computational mechanics: a review. Applied Mechanics Reviews, 49:1-28, 1996.
- [BW04] R. Baule and M. Wilkens. Lean Trees — A General Approach for Improving Performance of Lattice Models for Option Pricing. Review of Derivatives Research, 7(1):53-72, 2004. www.fernuniversität-hagen.de/bwlbuf/uploads/allgemein/sonstiges/lean_trees.pdf.
- [CH52] C.F. Curtiss and J.O. Hirschfelder. Integration of Stiff Equations. Proceedings of the National Academy of Sciences of the United States of America, 38(3):235-243, 1952.
- [CS09] M.E. Chenoweth and S.A. Sarra. A numerical study of generalized multiquadric radial basis function interpolation. SIAM Undergraduate Research Online, 2:58–70, 2009.
- [Cur58] P.C.J. Curtis. n-parameter families and best approximation. Pacific J. Math., 93:1013-1027, 1958.

- [DM87] A. R. DiDonato and A. H. Morris. Incomplete gamma function ratios and their inverse. ACM TOMS, 13:318-319, 1987. http://www.netlib.org/toms/654.
- [Fas11] G.E. Fasshauer. Positive definite kernels: past, present and future. Dolomite Research Notes on Approximation, 4:21-63, 2011. num.math.uni-goettingen.de/meshless/abstracts/fasshauer.pdf.
- [FF15a] B. Fornberg and N. Flyer. Radial Basis Function Generated Finite Differences (RBF-FD): New Computational Opportunities for Solving PDEs, 2015. amath.colorado.edu/faculty/fornberg/Docs/2015%200xford.pdf.
- [FF15b] B. Fornberg and N. Flyer. Solving PDEs with radial basis functions. Acta Numerica, 24:215-258, 2015. amath.colorado.edu/faculty/fornberg/Docs/2015_FF_Solving%20PDEs%20with%20RBFs_Acta_Numerica.pdf.
- [FM12] G.E. Fasshauer and M.J. McCourt. Stable Evaluation of Gaussian Radial Basis Function Interpolants. SIAM J. Sci. Comput., 34(2):737-762, March 2012. math.iit.edu/~fass/StableGaussianRBFs.pdf.

[Fra79] Richard Franke. A critical comparison of some methods for interpolation of scattered data. Technical report, Monterey, California: Naval Postgraduate School., 1979. www.dtic.mil/dtic/tr/fulltext/u2/a081688.pdf.

[FZ07] B. Fornberg and J. Zuev. The Runge phenomenon and spatially variable shape parameters in RBF interpolation. Computers & Mathematics with Applications, 54(3):379–398, 2007.

| Peter Jäckel | Cluster Induction | May 2018 | 133 / 117 |
|--------------|-------------------|----------|-----------|
| | | | |

References

| [Gea71] | C.W. Gear. |
|---------|--|
| | Numerical initial value problems in Ordinary Differential Equations. |
| | Prentice Hall, 1971. |

[Har71] R.L. Hardy. Multiquadric equations of topography and other irregular surfaces. Journal of Geophysical Research, 76(8):1905–1915, 1971.

[Jäc05] P. Jäckel.

A note on multivariate Gauss-Hermite quadrature.

www.jaeckel.org/ANoteOnMultivariateGaussHermiteQuadrature.pdf, May 2005.

[Jäc13] P. Jäckel.

Finite differencing schemes as Padé approximants, April 2013.

www.jaeckel.org/FiniteDifferencingSchemesAsPad%C3%A9Approximants.pdf.

[Kan90a] E.J. Kansa.

Multiquadrics - A scattered data approximation scheme with applications to computational fluid-dynamics—I surface approximations and partial derivative estimates. *Computers & Mathematics with Applications*, 19(8–9):127–145, 1990.

[Kan90b] E.J. Kansa.

Multiquadrics - A scattered data approximation scheme with applications to computational fluid-dynamics—II solutions to parabolic, hyperbolic and elliptic partial differential equations. *Computers & Mathematics with Applications*, 19(8–9):147–161, 1990.

[LF03] E. Larsson and B. Fornberg.

A numerical study of some radial basis function based solution methods for elliptic PDEs. Computers & Mathematics with Applications, 46(5):891-902, 2003. amath.colorado.edu/faculty/fornberg/Docs/el_bf.pdf.

| [LI057] | SP. Lloyd. |
|---------|---|
| | Least Squares Quantization in PCM. |
| | Technical report, Bell Laboratories, 1957. |
| | Unpublished Bell Laboratories Technical Note. Portions presented at the Institute of Mathematical Statistics Meeting Atlantic City New Jersey September 1957. Published in the March 1982 special issue on quantization of the IEEE Transactions on Information Theory. |
| [Llo18] | Llovd's algorithm. |
| | Llovd's algorithm — Wikipedia. 2018. |
| | [Online; accessed 02-March-2018]. |
| [Mai56] | J.C. Mairhuber. |
| | On Haar's theorem concerning Chebyshev approximation problems having unique solutions. <i>Proc. Amer. Math. Soc</i> , 7:609–615, 1956. |
| [Mic84] | C.A. Micchelli. |
| | Interpolation of scattered data: distance matrices and conditionally positive definite functions. In <i>Approximation theory and spline functions</i> , pages 143–145. Springer, 1984. |
| [Mic86] | C.A. Micchelli. |
| _ | Intermediation of constrained data. Distance metalogical conditionally metalogical fortations |

Interpolation of scattered data: Distance matrices and conditionally positive definite functions. *Constructive Approximation*, 2(1):11–22, 1986.

[Mil14] S. Milovanović.

Radial Basis Functions generated Finite Differences (RBF-FD) for Solving High-Dimensional PDEs in Finance.

Uppsala University, June 2014. people.kth.se/~ejhall/cirkus/talks/milovanovic.pdf.

[Mon11] M. Mongillo.

Choosing basis functions and shape parameters for radial basis function methods. SIAM Undergraduate Research Online, 4:190-209, 2011. 72.32.205.185/students/siuro/vol4/S01084.pdf.

| Peter Jäckel | Cluster Induction | May 2018 | 135 / 117 |
|--------------|-------------------|----------|-----------|

References

[PD51] PF. Panter and W. Dite. Quantization distortion in pulse-count modulation with nonuniform spacing of levels. Proc. I.R.E., 39, 1 1951.

[SDY03] C. Shu, H. Ding, and K.S. Yeo. Local radial basis function-based differential quadrature method and its application to solve two-dimensional incompressible Navier-Stokes equations. *Computer Methods in Applied Mechanics and Engineering*, 192:941–954, 02 2003.

[SL16] V. Shcherbakov and E. Larsson. Radial basis function partition of unity methods for pricing vanilla basket options. Computers & Mathematics with Applications, 71(1):185 - 200, 2016. www.it.uu.se/research/publications/reports/2015-001/2015-001-nc.pdf.

[Smi86] G. Smith. Numerical Solution of Partial Differential Equations: Finite Difference Methods. Oxford University Press, January 1986.

[Tol00] A.I. Tolstykh. On using RBF-based differencing formulas for unstructured and mixed structured–unstructured grid calculations. In Proc. 16th IMACS World Congress, pages 4606—4624, 2000.

[TS03] A.I. Tolstykh and D.A. Shirobokov. On using radial basis functions in a "finite difference mode" with applications to elasticity problems. *Computational Mechanics*, 33(1):68–79, 2003.

[Wan14] Xi Wang. Finite differences based on radial basis functions to price options, 2014. www.diva-portal.org/smash/get/diva2:787451/FULLTEXT01.pdf.

[WS02] Y.L. Wu and C. Shu.

Development of RBF-DQ method for derivative approximation and its application to simulate natural convection in concentric annuli.

Computational Mechanics, 29(6):477-485, 2002.

Peter Jäckel

Cluster Induction

May 2018