Quadrature by Expansion

Andreas Klöckner

Computer Science · UIUC

November 12, 2013

Thanks

- Leslie Greengard (NYU)
- Alex Barnett (Dartmouth)
- Michael O'Neil (NYU)
- Zydrunas Gimbutas (NIST)

Outline

Outline

1. Introduction
    - Method vs Machine
    - BIEs for PDEs
    - Existing Approaches to Quadrature
2. Developing QBX
3. Method design
4. Experimental results
5. Conclusions

Outline

1. Introduction
    - Method vs Machine
    - BIEs for PDEs
    - Existing Approaches to Quadrature
2. Developing QBX
3. Method design
4. Experimental results
5. Conclusions

CPU Chip Real Estate

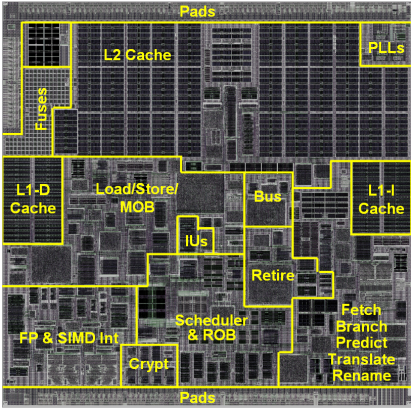

Die floorplan: VIA Isaiah (2008)

Improving matters for number crunching

Latency-constrained machine

Improving matters for number crunching

✗ ~~Latency-constrained machine~~

Improving matters for number crunching

✘ ~~Latency-constrained machine~~

✓ Throughput-constrained machine

Recent Processor Architecture

- Massively concurrent
  - ILP, SMT for (memory) latency
  - parallel by SIMD
  - parallel across cores

- Rarely flop-limited
- Limited by data motion

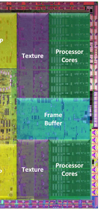

T200
08)

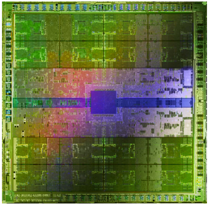

Nv Fermi
(2010)

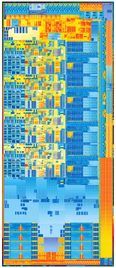

Intel IVB
(2012)

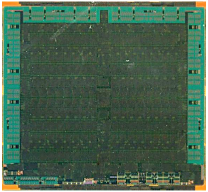

AMD Tahiti
(2012)

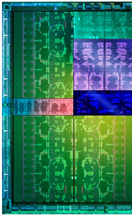

Nv GK1
(2012

Computational wish list for PDE schemes

A method and its algorithm should...

- admit massive concurrency
- provide a match for hierarchical concurrency
- have a high computation/access ratio ('intensity')
- use minimal storage

✓

Application-driven wish list

- Unstructured geometries
    - Adaptable to many engineering problems
    - Compatible with adaptive discretization
- General-purpose (as much as possible—in PDE, BCs, . . .)
- Robust
- Well-conditioned
- High order
    - Also in geometry representation

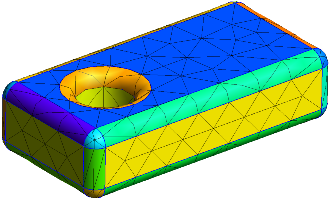

Why high order?

Order p: Error bounded as

$$\|u_h - u\| \leq Ch^p$$

Why high order?

Order p: Error bounded as

$$\|u_h - u\| \leq Ch^p$$

Thought experiment:

First order	Fifth order
1,000 DoFs $\approx$ 1,000 triangles Error: 0.1	1,000 DoFs $\approx$ 66 triangles Error: 0.1

Why high order?

Order p: Error bounded as

$$\|u_h - u\| \leq Ch^p$$

Thought experiment:

First order	Fifth order
1,000 DoFs $\approx$ 1,000 triangles Error: 0.1	1,000 DoFs $\approx$ 66 triangles Error: 0.1
Error: 0.01 @ **100,000 DoFs** $\approx$ **100,000 triangles**	Error: 0.01 @ **1,800 DoFs** $\approx$ **120 triangles**

Why high order?

Order p: Error bounded as

$$\|u_h -$$

Want $p \geq 3$ available.

Assumption: Solution sufficiently smooth

Ideally: p chosen by user

Thought experiment:

First order	Fifth order
1,000 DoFs $\approx$ 1,000 triangles Error: 0.1	1,000 DoFs $\approx$ 66 triangles Error: 0.1
Error: 0.01 @ **100,000 DoFs** $\approx$ **100,000 triangles**	Error: 0.01 @ **1,800 DoFs** $\approx$ **120 triangles**

Why high order?

Want $p \geq 3$ available.

Assumption: Solution sufficiently smooth

Ideally: p chosen by user

Order p: Error bounded as

$\|u_h$ -

Also beneficial for modern machines:

- More accuracy with fewer degrees of freedom
  - Fewer degrees of freedom needed for given accuracy
- More operations on less data
  - Exploit architecturally 'free' flops

th order

00 DoFs ≈ 66 triangles

ror: 0.1

ror: 0.01 @

300 DoFs

120 triangles

Outline

Fundamental Solutions

Laplace Equation

$\triangle u = 0$

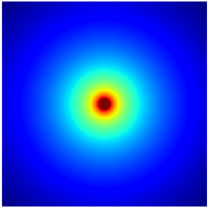

Monopole

Fundamental Solutions

Laplace Equation

$\triangle u = 0$

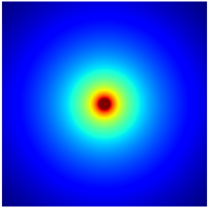

Monopole

Helmholtz Equation

$\triangle u + k^2 u = 0$

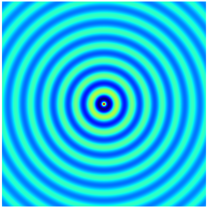

Monopole

Fundamental Solutions

Laplace Equation

$\triangle u = 0$

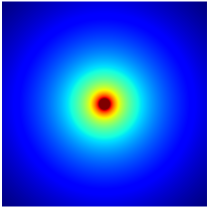

Monopole

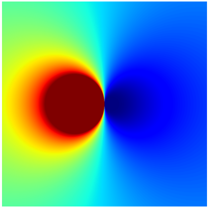

Dipole

Helmholtz Equation

$\triangle u + k^2 u = 0$

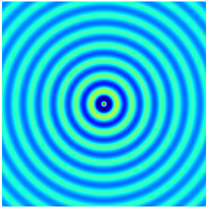

Monopole

Fundamental Solutions

Laplace Equation

$\triangle u = 0$

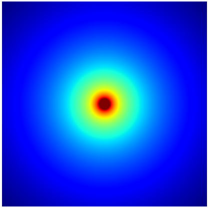

Monopole

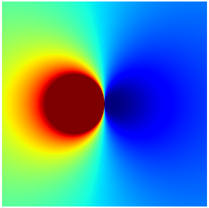

Dipole

Helmholtz Equation

$\triangle u + k^2 u = 0$

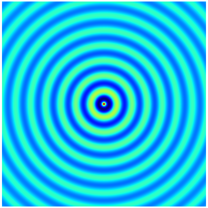

Monopole

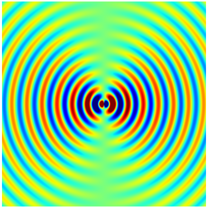

Dipole

Fundamental Solutions

Laplace Equation

$\triangle u = 0$

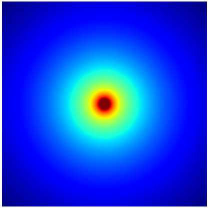

Monopole

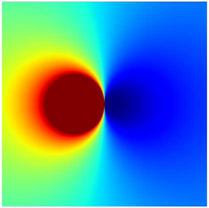

Dipole

Helmholtz Equation

Can take this arbitrarily far:

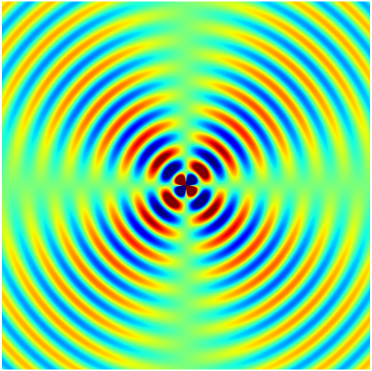

Quadrupole, ...

Building Solutions

Main question for numerical solution of PDEs:

How is the solution represented?

Our choice here:
Sums of fundamental solutions

$$\tilde{u}(x) = \sum_{i=1}^{N} G(|x - y_i|)\sigma_i$$

- Is the solution reachable in this way?
    - Uniqueness?
- Linearity → must satisfy PDE
- Boundary conditions: not necessarily

Summing Fundamental Solutions

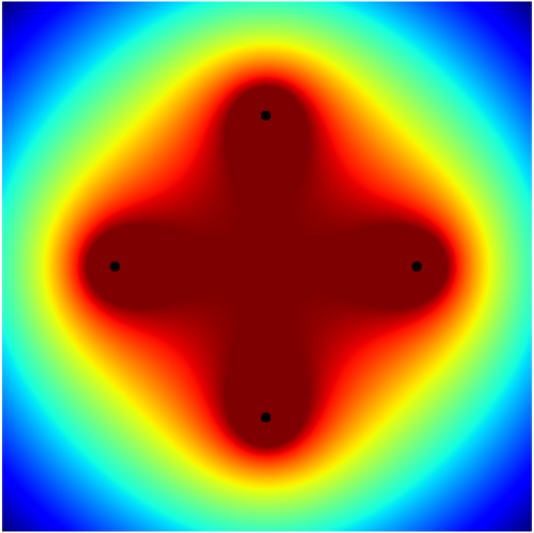

Summing Fundamental Solutions

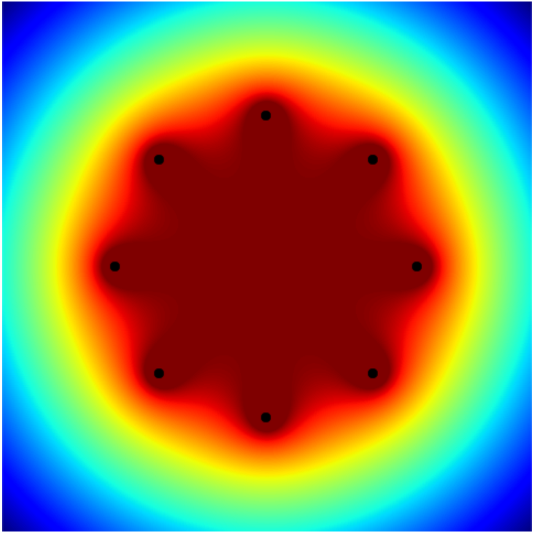

Summing Fundamental Solutions

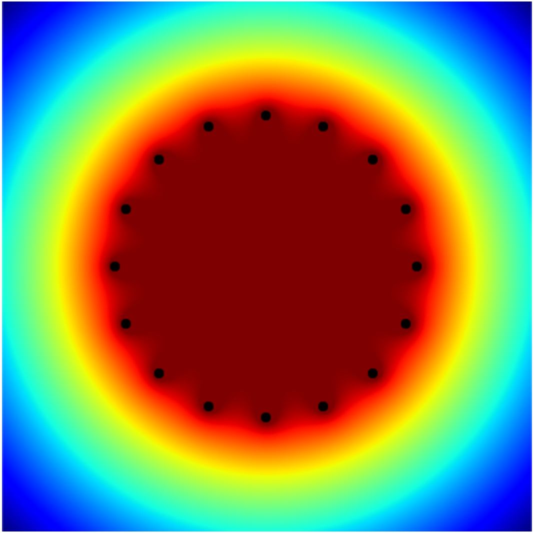

Summing Fundamental Solutions

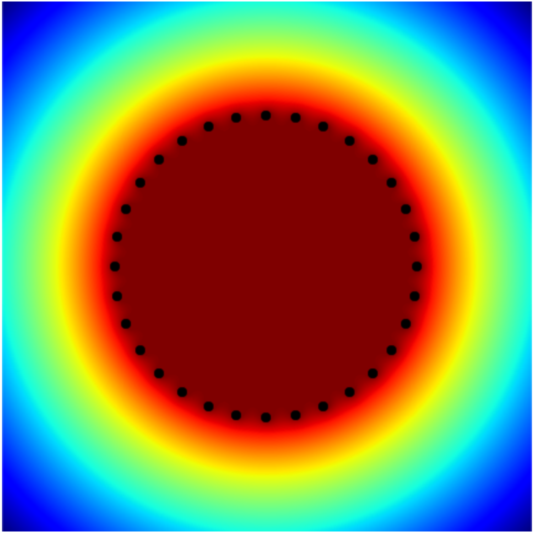

Summing Fundamental Solutions

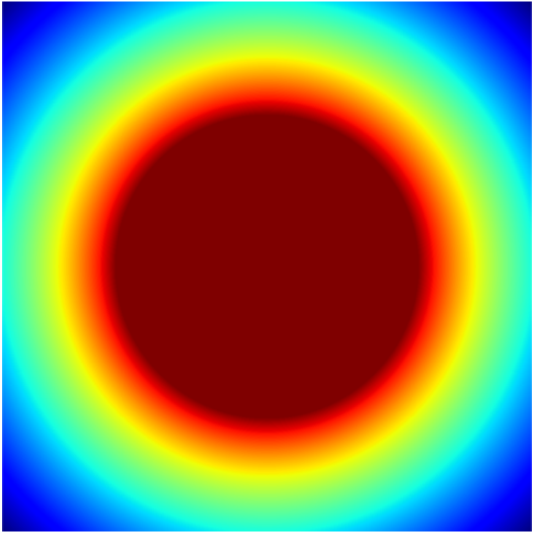

Layer Potentials

$$(S_k\sigma)(x) := \int_\Gamma G_k(x-y)\sigma(y)ds_y$$

$$(S'_k\sigma)(x) := n\cdot\nabla_x PV\int_\Gamma G_k(x-y)\sigma(y)ds_y$$

$$(D_k\sigma)(x) := PV\int_\Gamma n\cdot\nabla_y G_k(x-y)\sigma(y)ds_y$$

$$(D'_k\sigma)(x) := n\cdot\nabla_x f.p.\int_\Gamma n\cdot\nabla_y G_k(x-y)\sigma(y)ds_y$$

- Operators–map function σ on Γ to…
  - …function on $\mathbb{R}^n$
  - …function on Γ (in particular)
- S'' (and higher) analogously
- Called *layer potentials*
- G_k is the Helmholtz kernel ($k = 0 \to$ Laplace)

Example: Integral Equations for Maxwell's

Want to solve the

- time-harmonic
- linear
- isotropic

Maxwell's equations

- for the scattered field response
  (to a given incoming field)
- in an exterior domain
- with a metallic scatterer.

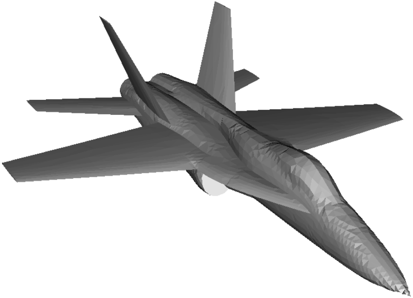

Integral Equations for Maxwell's

Use:

- Vector potential $H = \nabla \times A$ with $\triangle A + k^2 A = 0$
- Integral representation for vector potential

$$A(x) = (S_k J^s)(x) = \frac{1}{4\pi} \int_\Gamma \frac{e^{ik|x-x'|}}{|x-x'|} J^s(x') dx'$$

Integral Equations for Maxwell's

Then

- Continuity condition

$$\hat{n} \times \left(H^{+}_{\text{tot}} - H^{-}_{\text{tot}}\right) = J^{s}$$

- No field on int. of PEC: $H^{-}_{\text{tot}} = 0$
- Jump condition

together yield MFIE [Maue, ...]

$$\hat{n} \times H^{+}_{\text{inc}} = \frac{J^{s}}{2} - (n \times (PV)(\nabla \times S_k J^{s}))^{s}$$

Integral Equations for Maxwell's

Then

- Continuity condition

$$\hat{n} \times \left(H^+_{\text{tot}} - H^-_{\text{tot}}\right) = J^s$$

- No field on int. of PEC: $H^-_{\text{tot}} = 0$
- Jump condition

together yield MFIE [Maue, . . .]

$$\hat{n} \times H^+_{\text{inc}} = \underbrace{\frac{J^s}{2}}_{\text{Identity}} + \underbrace{\left(n \times (PV)\left(\nabla \times S_k J^s\right)\right)^s}_{\text{Compact}}$$

Integral Equations for Maxw

Solve for J_s

Can now compute magnetic and electric (…) fields everywhere

Then

- Continuity condition

$$\hat{n} \times \left(H^{+}_{\text{tot}} - H^{-}_{\text{tot}}\right) = J^{s}$$

- No field on int. of PEC: $H^{-}_{\text{tot}} = 0$
- Jump condition

together yield MFIE [Maue, …]

$$\hat{n} \times H^{+}_{\text{inc}} = \underbrace{\frac{J^{s}}{2}}_{\text{Identity}} + \underbrace{\left(n \times (PV)\left(\nabla \times S_k J^{s}\right)\right)^{s}}_{\text{Compact}}$$

One issue with BIE methods

Need to compute

$$n \times (\nabla \times S_k J^s)(x) = \frac{1}{4\pi} n \times PV \int_\Gamma \nabla_x \times \frac{e^{ik|x-x'|}}{|x-x'|} J^s(x')dx'$$

for $x \in \Gamma$.

One issue with BIE methods

Need to compute

$$n \times (\nabla \times S_k J^s)(x) = \frac{1}{4\pi} n \times PV \int_\Gamma \underbrace{\nabla_x \times \frac{e^{ik|x-x'|}}{|x-x'|}}_{\textbf{Difficult when } x \approx x'} J^s(x') dx'$$

for $x \in \Gamma$.

One issue with BIE methoc

Two (main) approaches:

- Galerkin (MoM)
- Nyström

Need to compute

$$n \times (\nabla \times S_k J^s)(x) = \frac{1}{4\pi} n \times PV \int_\Gamma \nabla_x \times \frac{e^{ik|x-x'|}}{|x-x'|} J^s(x') dx'$$

for $x \in \Gamma$.

Difficult when $x \approx x'$

One issue with BIE metho

Two (main) approaches:

- Galerkin (MoM)
- **Nyström**

Need to compute

$$n \times (\nabla \times S_k J^s)(x) = \frac{1}{4\pi} n \times PV \int_\Gamma \nabla_x \times \frac{e^{ik|x-x'|}}{|x-x'|} J^s(x') dx'$$

for $x \in \Gamma$.

Difficult when $x \approx x'$

One issue with BIE metho

Two (main) approaches:

- Galerkin (MoM)
- **Nyström**

Substantial efficiency gains, especially at high order

Need to compute

$$n \times (\nabla \times S_k J^s)(x) = \frac{1}{4\pi} n \times PV \int_\Gamma \nabla_x \times \frac{e^{ik|x-x'|}}{|x-x'|} J^s(x') dx'$$

for $x \in \Gamma$.

Difficult when $x \approx x'$

Outline

1. Introduction
   - Method vs Machine
   - BIEs for PDEs
   - Existing Approaches to Quadrature
2. Developing QBX
3. Method design
4. Experimental results
5. Conclusions

A Sampling of Ideas

- Analytic/symbolic integration
  - Sometimes possible
  - Problematic because of geometry description
  - Numerical stability of resulting formulas?
- Adaptive integration
  - Fails because (many) singularities are not integrable
  - Expensive
- Change of variables/singularity subtraction/cancellation
  - Algebraic trickery weakens/removes singularity
  - Not general-purpose (across dimensions, kernels)

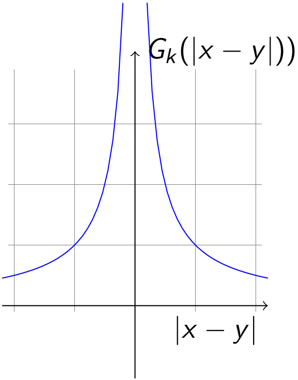

A Sampling of Ideas

- Analytic/symbolic integration
  - Sometimes possible
  - Problematic because of geometry description
  - Numerical stability of resulting formulas?
- Adaptive integration
  - Fails because (many) singularities are not integrable
  - Expensive
- Change of variables/singularity subtraction/cancellation
  - Algebraic trickery weakens/re singularity
  - Not general-purpose (across kernels)

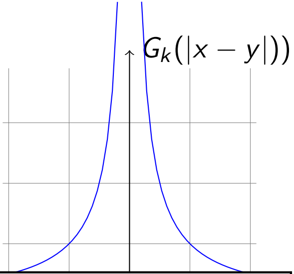

Primary drawbacks:

- Not generic in singularity
- Depend on curve (2D?) geometry

A Sampling of Ideas

- Analytic/symbolic integration
  - Sometimes possible
  - Problematic because of geometry description
  - Numerical stability of resulting formulas?

gularities are

rity

ns/re

ross

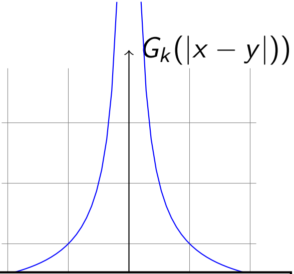

Many more have worked on the problem: Sidi, Strain, Helsing, Davis, Duffy, Graglia, Hackbusch, Khayat, Schwab, Ying, Beale, Goodman, Haroldson, Lowengrub, Alpert, Rokhlin, Gimbutas, Bruno, Zorin, . . .

Primary drawbacks:

- Not generic in singularity
- Depend on curve (2D?) geometry

Kernel Regularization

Singularity makes integration troublesome: *Get rid of it!*

$$\frac{\cdots}{\sqrt{(x-y)^2}} \quad \rightarrow \quad \frac{\cdots}{\sqrt{(x-y)^2+\varepsilon^2}}$$

Use Richardson extrapolation to recover limit as $\varepsilon \to 0$.

(May also use geometric motivation: limit along line towards singular point.)

Primary drawbacks:

- Low-order accurate
- Need to make ε smaller (i.e. kernel more singular) to get better accuracy

Outline

1 Introduction

2 Developing QBX
- Some intuition
- Insight through theory
- Other Potentials

3 Method design

4 Experimental results

5 Conclusions

Outline

Using just the trapezoidal rule

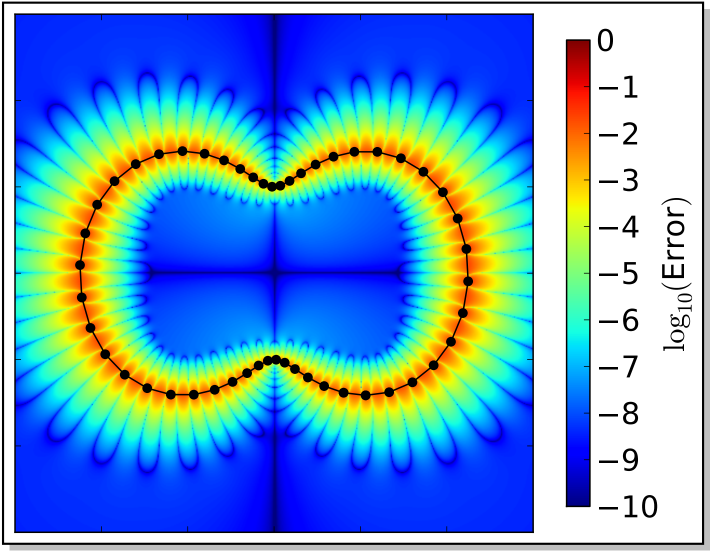

Using just the trapezoidal rule

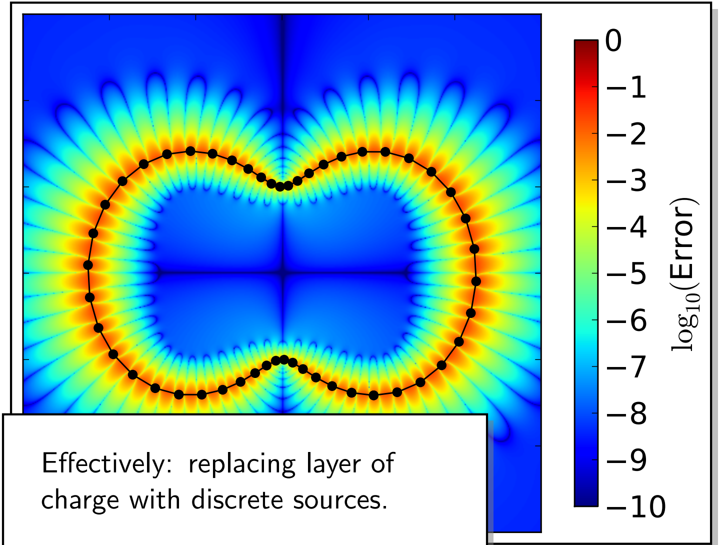

Effectively: replacing layer of charge with discrete sources.

Using just the tr

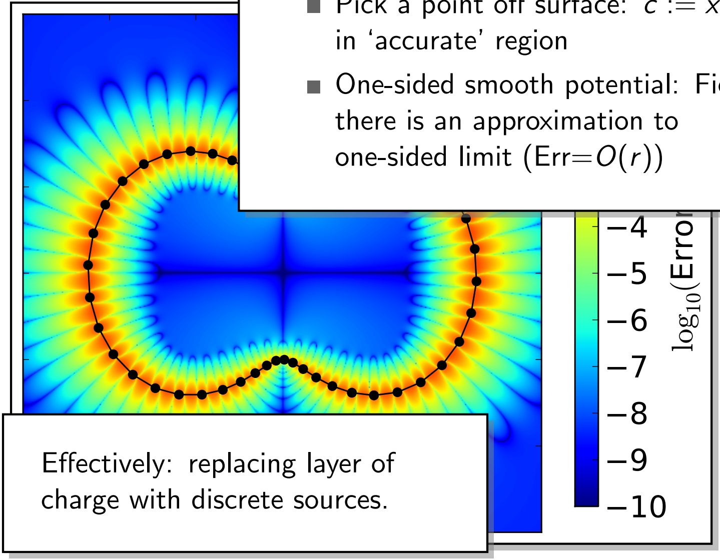

Idea:

- Pick a point off surface: $c := x + \hat{n}r$ in ‘accurate’ region
- One-sided smooth potential: Field value there is an approximation to one-sided limit (Err$=O(r)$)

Effectively: replacing layer of charge with discrete sources.

Using just the tr

Idea:

- Pick a point off surface: $c := x + \hat{n}r$ in ‘accurate’ region
- One-sided smooth potential: Field value there is an approximation to one-sided limit (Err$=O(r)$)

But: Can do much better!

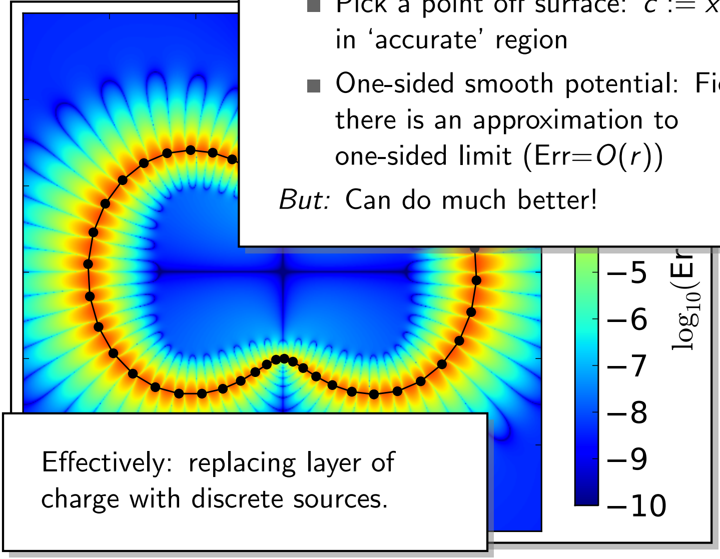

Effectively: replacing layer of charge with discrete sources.

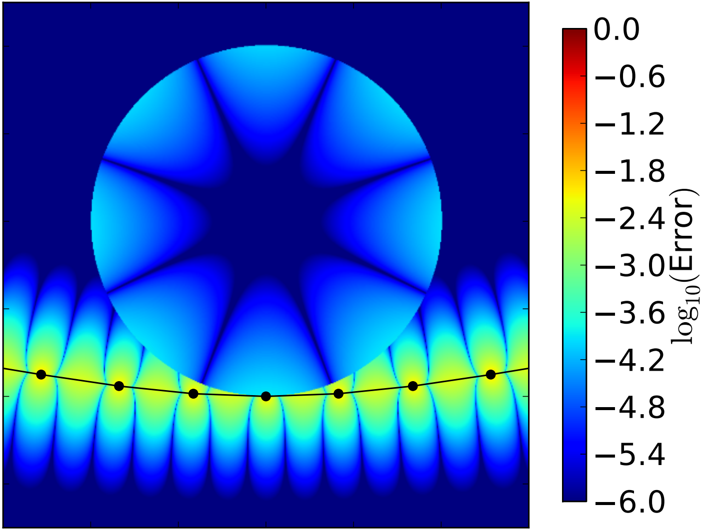
0.0
−0.6
−1.2
−1.8
−2.4
−3.0
−3.6
−4.2
−4.8
−5.4
−6.0
$\log_{10}$(Error)

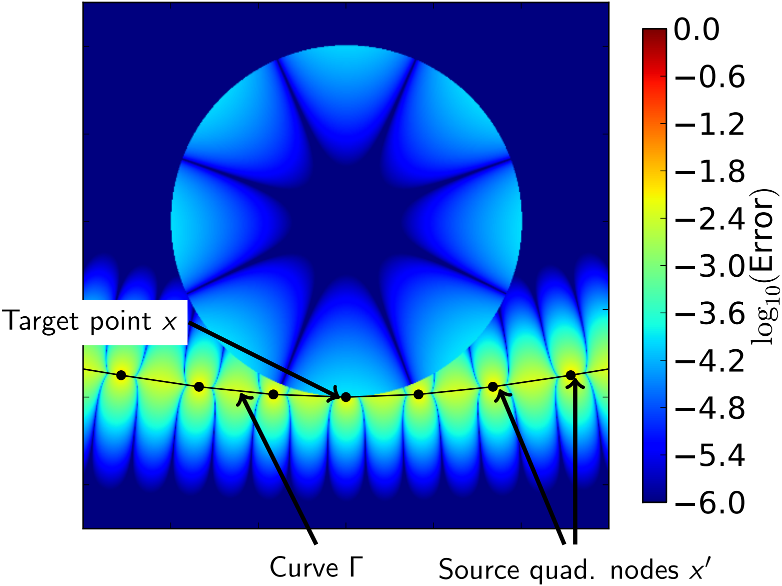
Target point x
Curve Γ
Source quad. nodes x'
0.0
−0.6
−1.2
−1.8
−2.4
−3.0
−3.6
−4.2
−4.8
−5.4
−6.0
$\log_{10}$(Error)

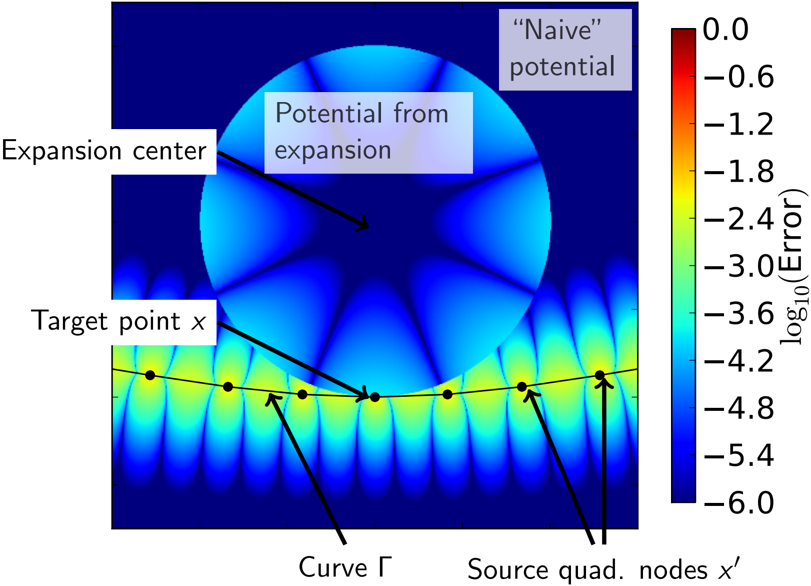
"Naive" potential
Potential from expansion
Expansion center
Target point x
Curve Γ
Source quad. nodes x'
0.0
−0.6
−1.2
−1.8
−2.4
−3.0
−3.6
−4.2
−4.8
−5.4
−6.0
$\log_{10}$(Error)

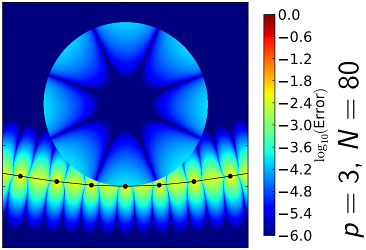
0.0
−0.6
−1.2
−1.8
−2.4
−3.0
−3.6
−4.2
−4.8
−5.4
−6.0
$\log_{10}$(Error)
$p = 3$, $N = 80$

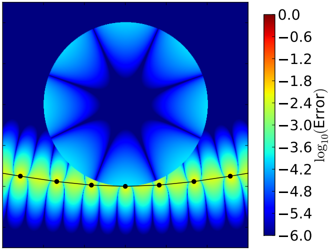
0.0
−0.6
−1.2
−1.8
−2.4
−3.0
−3.6
−4.2
−4.8
−5.4
−6.0
$\log_{10}$(Error)

$p = 3,\ N = 80$

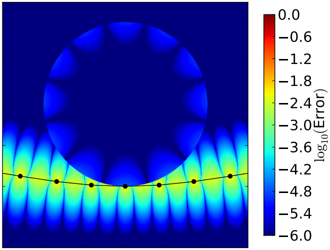
0.0
−0.6
−1.2
−1.8
−2.4
−3.0
−3.6
−4.2
−4.8
−5.4
−6.0
$\log_{10}$(Error)

$p = 6,\ N = 80$

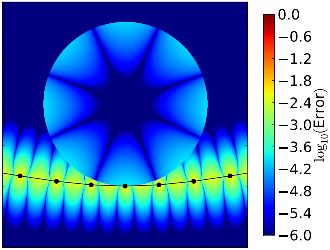

$p = 3$, $N = 80$

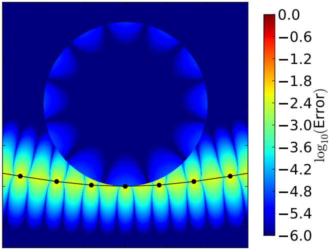

$p = 6$, $N = 80$

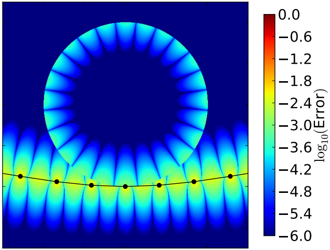

$p = 12$, $N = 80$

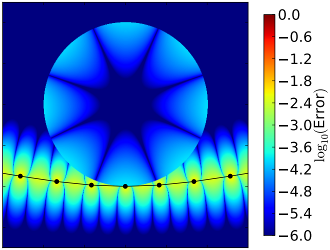

$p = 3$, $N = 80$

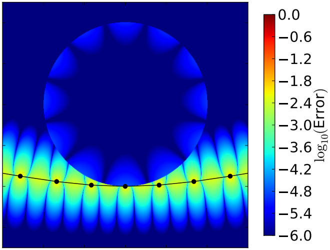

$p = 6$, $N = 80$

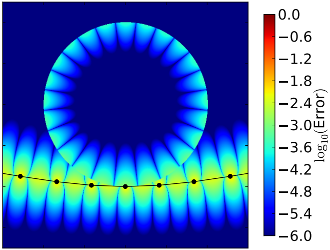

$p = 12$, $N = 80$

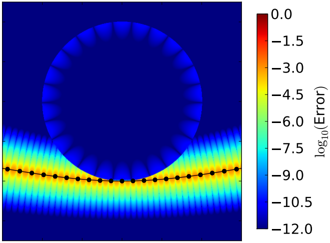

$p = 12$, $N = 240$

QBX in formulas: Notation, Basics

Graf's addition theorem

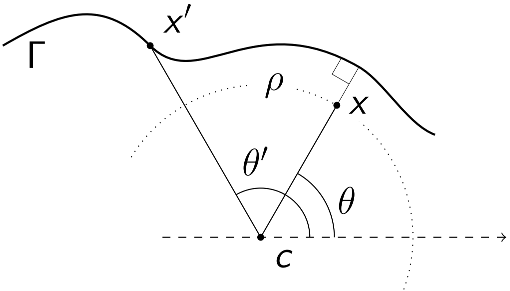

QBX in formulas: Notation, Basics

Graf's addition theorem

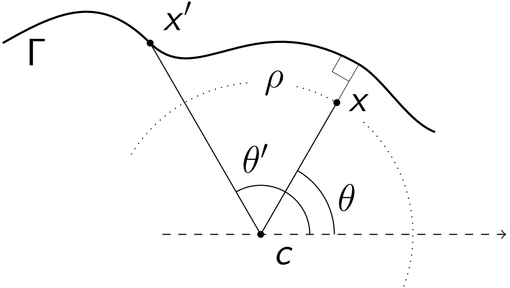

$$H_0^{(1)}(k|x-x'|) = \sum_{l=-\infty}^{\infty} H_l^{(1)}(k|x'-c|)e^{il\theta'} J_l(k|x-c|)e^{-il\theta}$$

QBX in formula

Requires: $|x - c| < |x' - c|$ ("local expansion")

Graf's addition theorem

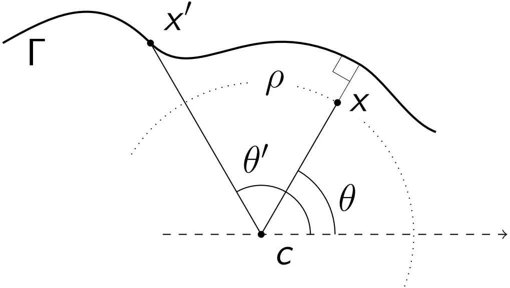

$$H_0^{(1)}(k|x - x'|) = \sum_{l=-\infty}^{\infty} H_l^{(1)}(k|x' - c|)e^{il\theta'} J_l(k|x - c|)e^{-il\theta}$$

QBX in formulas: Formulation, discretization

Compute layer potential on the disk as

$$S_k\sigma(x) = \sum_{l=-\infty}^{\infty} \alpha_l J_l(k\rho) e^{-il\theta}$$

with

$$\alpha_l = \frac{i}{4} \int_\Gamma H_l^{(1)}(k|x' - c|) e^{il\theta'} \sigma(x')\, dx' \quad (l = -\infty, \dots, \infty)$$

$S\sigma$ is a smooth function *up to* Γ.

QBX in formulas: Formulation, discretiza

Now discretize.

Compute layer potential on the disk as

$$S_k\sigma(x) = \sum_{l=-\infty}^{\infty} \alpha_l J_l(k\rho)e^{-il\theta}$$

with

$$\alpha_l = \frac{i}{4}\int_\Gamma H_l^{(1)}(k|x'-c|)e^{il\theta'}\sigma(x')\,dx' \quad (l=-\infty,\ldots,\infty)$$

$S\sigma$ is a smooth function *up to* Γ.

QBX in formulas: Formulation, discretiza

Now discretize.

Compute layer potential on the disk as

$$S_k\sigma(x) = \sum_{l=-p}^{p} \alpha_l J_l(k\rho) e^{-il\theta}$$

with

$$\alpha_l = \frac{i}{4} \int_\Gamma H_l^{(1)}(k|x' - c|) e^{il\theta'} \sigma(x')\, dx' \quad (l = -\infty, \dots, \infty)$$

$S\sigma$ is a smooth function *up to* Γ.

QBX in formulas: Formulation, discretiza

Now discretize.

Compute layer potential on the disk as

$$S_k\sigma(x) = \sum_{l=-p}^{p} \alpha_l J_l(k\rho) e^{-il\theta}$$

with

$$\alpha_l = \frac{i}{4} T_N\left(\int_\Gamma H_l^{(1)}(k|x'-c|) e^{il\theta'} \sigma(x')\, dx' \right) \quad (l = -\infty, \ldots, \infty)$$

$S\sigma$ is a smooth function *up to* Γ.

QBX in formulas: Formulation, discretiza

Now discretize.

Compute layer potential on the disk as

$$S_k\sigma(x) = \sum_{l=-p}^{p} \alpha_l J_l(k\rho) e^{-il\theta}$$

with

$$\alpha_l = \frac{i}{4} T_N\left(\int_\Gamma H_l^{(1)}(k|x' - c|) e^{il\theta'} \sigma(x')\, dx' \right) \quad (l = -\infty, \ldots, \infty)$$

$S\sigma$ is a smooth function *up to* Γ.

Two limits ($p, N \to \infty$)! Experiment showed: *order matters!*

QBX in formulas: Formulation, discretiza

Now discretize.

Compute layer potential on the disk as

$$S_k\sigma(x) = \sum_{l=-p}^{p} \alpha_l J_l(k\rho)e^{-il\theta}$$

with

$$\alpha_l = \frac{i}{4} T_N\left(\int_\Gamma H_l^{(1)}(k|x' - c|)e^{il\theta'}\sigma(x')\,dx' \right) \quad (l = -\infty, \ldots, \infty)$$

Two limits ($p, N \to \infty$)! Experiment showed: *order matters!*

And: failure and repair not actually surprising.

Outline

Error result

$$\left| S\sigma(x) - \sum_{l=-p}^{p} \alpha_l^{\text{QBX}} J_l(k|x-c|) e^{-il\theta_{cx}} \right|$$

$$\leq \left(\underbrace{C_{p,\beta}\ r^{p+1} \|\sigma\|_{\mathcal{C}^{p,\beta}(\Gamma)}}_{\text{Truncation error}} + \underbrace{\tilde{C}_{p,2q,\beta} \left(\frac{h}{4r} \right)^{2q} \|\sigma\|_{\mathcal{C}^{2q,\beta}(\Gamma)}}_{\text{Quadrature error}} \right)$$

Proof sketch:

1. First, assume exact calculation of coefficients
2. Estimate tail of expansion
3. Estimate quadrature error in coefficients (derivatives/...)
4. Sum quadrature errors in truncated expansion

[K, Barnett, Greengard, O'Neil JCP '13]

Outline

Other layer potentials

Can't just do single-layer potentials:

$$\alpha_l^D = \frac{i}{4} \int_\Gamma \frac{\partial}{\partial \hat{n}_{x'}} H_l^{(1)}(k|x' - c|) e^{il\theta'} \mu(x')\, dx'.$$

Even easier for target derivatives (S' et al.):

Take derivative of local expansion.

Analysis says: Will lose an order.

Other layer potentials

Can't just do single-layer potentials:

$$\alpha_l^D = \frac{i}{4} \int_\Gamma \frac{\partial}{\partial \hat{n}_{x'}} H_l^{(1)}(k|x' - c|) e^{il\theta'} \mu(x')\, dx'.$$

Even easier for target derivatives (S' et al.):

Take derivative of local expansion.

Analysis says: Will lose an order.

Slight issue: QBX computes one-sided limits.

Fortunately: Jump relations are known–e.g.

$$(PV) D^* \mu(x)|_\Gamma = \lim_{x^\pm \to x} D\mu(x^\pm) \mp \frac{1}{2}\mu(x).$$

Other layer potentials

Can't just do single-layer potentials:

$$\alpha_l^D = \frac{i}{4} \int_\Gamma \frac{\partial}{\partial \hat{n}_{x'}} H_l^{(1)}(k|x' - c|) e^{il\theta'} \mu(x')\, dx'.$$

Even easier for target derivatives (S' et al.):

Take derivative of local expansion.

Analysis says: Will lose an order.

Slight issue: QBX computes one-sided limits.

Fortunately: Jump relations are known–e.g.

$$(PV)D^*\mu(x)|_\Gamma =$$

Alternative: Two-sided average

Outline

Outline

1 Introduction

2 Developing QBX

3 Method design
- Achieving high order
- Unstructured Geometries
- Acceleration
- Numerical Conditioning

4 Experimental results

5 Conclusions

Achieving high order

$$\text{Error} \leq \left(C \underbrace{r^{p+1}}_{\text{Truncation error}} + C \underbrace{\left(\frac{h}{r}\right)^{q}}_{\text{Quadrature error}} \right) \|\sigma\|$$

Two approaches:

- *Asymptotically convergent*: $r = \sqrt{h}$
  - ⊕ Error $\to 0$ as $h \to 0$
  - ⊖ Low order: $h^{(p+1)/2}$
- *Convergent with controlled precision*: $r = 5h$
  - ⊖ Error $\not\to 0$ as $h \to 0$
  - ⊕ High order: h^{p+1}
    to controlled precision $\varepsilon := (1/5)^q$

Outline

"Global" QBX: Dealing with geometry

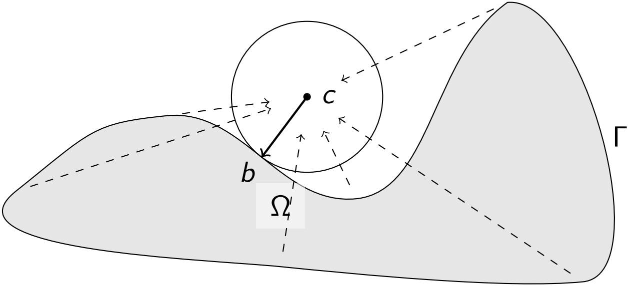

"Global" QBX, part II

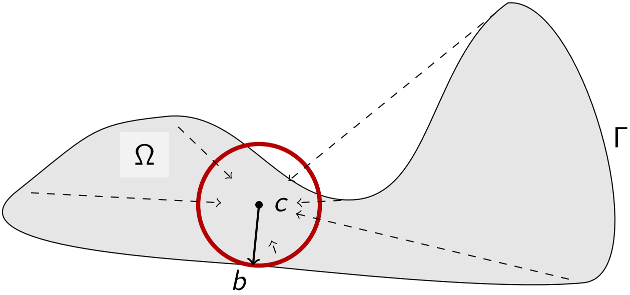

“Local” QBX

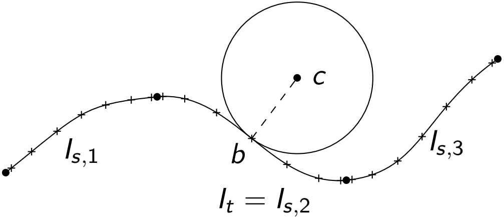

"Local" QBX

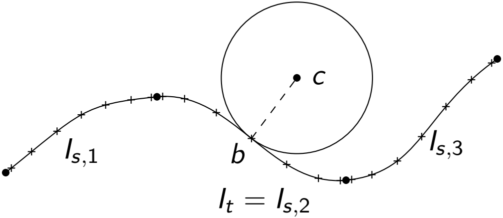

Makes geometry processing much simpler

“Local” QBX

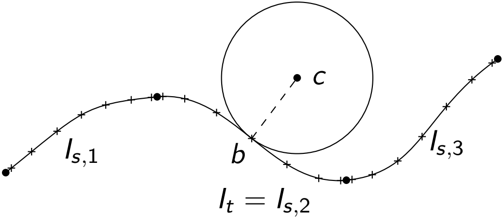

kes geometry process-
much simpler

Problem: Expanded field becomes non-smooth (because of end singularities)

"Local" QBX

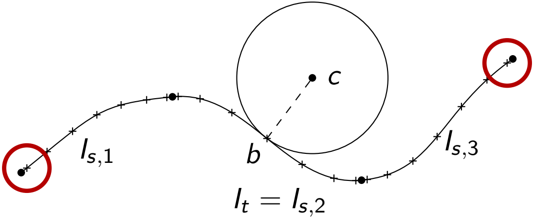

Problem: Expanded field becomes non-smooth (because of end singularities)

kes geometry process-
much simpler

"Local" QBX

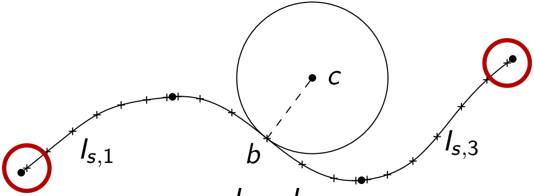

Problem: Expanded field becomes non-smooth (because of end singularities)

Idea: Manage as additional, finite error contribution (using p, $h \propto r$)

kes geometry process-
much simpler

Outline

1 Introduction

2 Developing QBX

3 Method design
- Achieving high order
- Unstructured Geometries
- Acceleration
- Numerical Conditioning

4 Experimental results

5 Conclusions

Fast Multipole Methods

Fast Multipole Methods

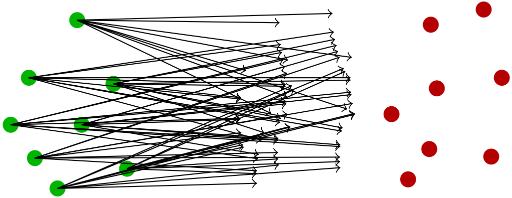

Fast Multipole Methods

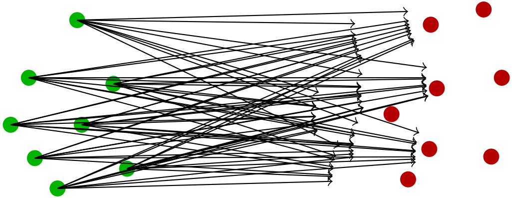

Fast Multipole Methods

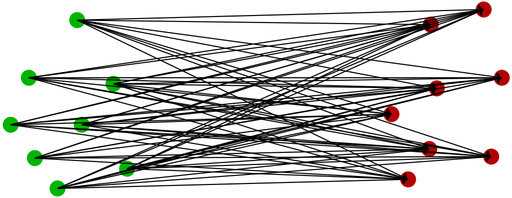

Fast Multipole Methods

Fast Multipole Methods

Fast Multipole Methods

Fast Multipole Methods

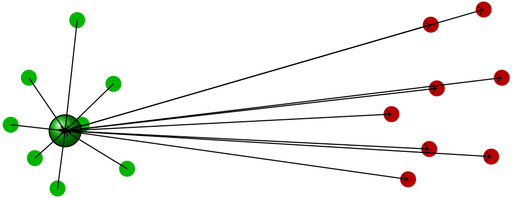

Fast Multipole Methods

Fast Multipole Methods

Fast Multipole Methods

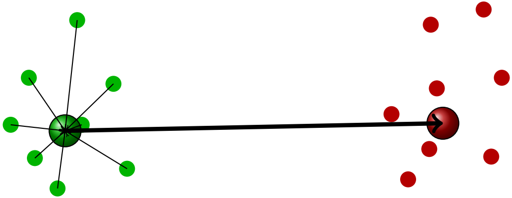

Fast Multipole Methods

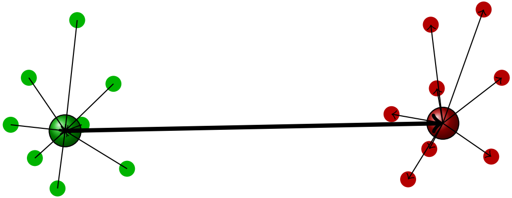

Fast Multipole Methods

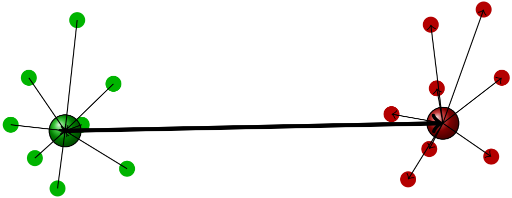

Only works if sources are ‘far enough’ away from targets.

Good: true for most particle pairs

Fast Multipole Methods

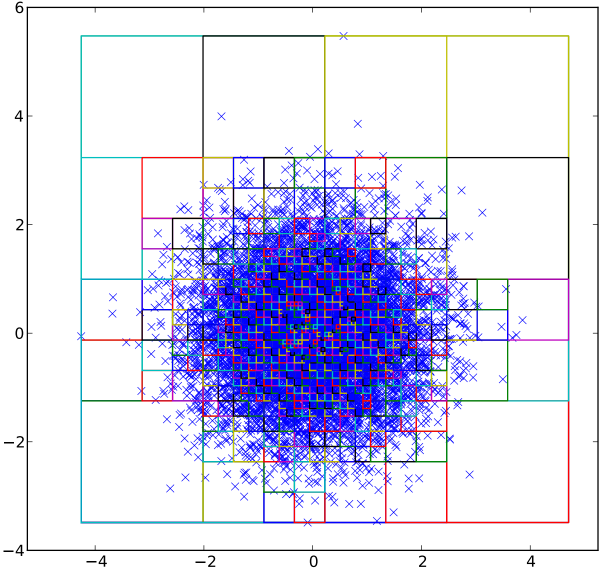

QBX + FMM: Two possibilities (ongoing)

“Global” QBX

- Requires modified FMM
- ✓ Cheap
- ✓ Accurate
- ✗ Not robust wrt complicated geometry
  - requires complicated “fallback”

“Local” QBX

- Requires modified FMM
- ✗ (Relatively) expensive
- ✗ (Controlled) error contribution from end singularities
- ✓ Robust wrt geometry

QBX + FMM: Two possibilities (ongoing)

“Global” QBX

- Requires modified FMM
- ✓ Cheap
- ✓ Accurate
- ✘ Not robust wrt complicated geometry
  - requires complicated “fallback”

“Local” QBX

- Requires modified FMM
- ✘ (Relatively) expensive
- ✘ (Controlled) error contribution from end singularities
- ✓ Robust wrt geometry

Complicated because of smoothness assumption

But: Smoothness assumption key strength

Outline

Spectral behavior

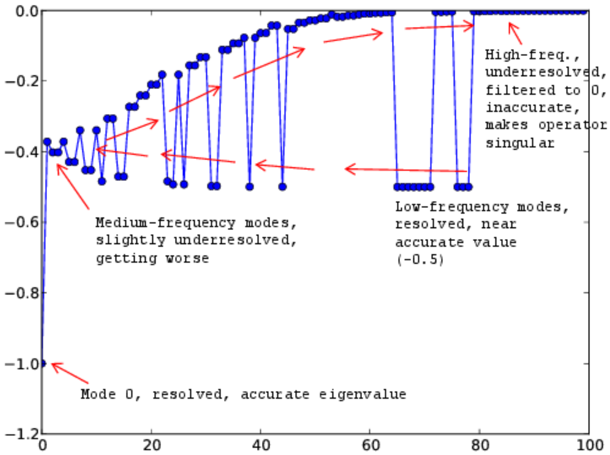

Spectral behavior

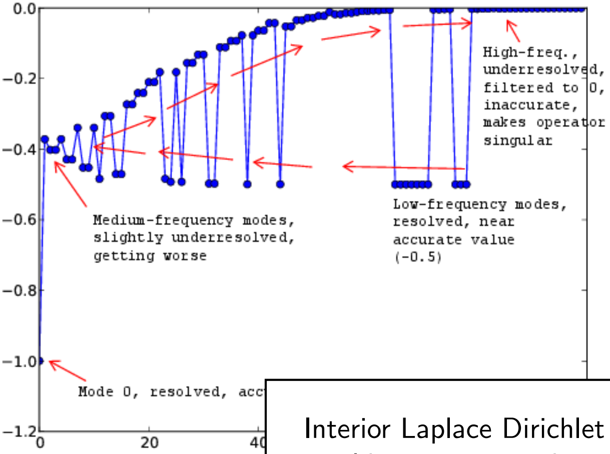

High-freq., underresolved, filtered to 0, inaccurate, makes operator singular

Medium-frequency modes, slightly underresolved, getting worse

Low-frequency modes, resolved, near accurate value (-0.5)

Mode 0, resolved, acc

Interior Laplace Dirichlet problem would try to invert this operator.

Spectral behavior, part II

- QBX *wants* to approximate a compact operator–let it:

$$D\mu(x) = \frac{1}{2}\left(\lim_{x^+\to x} D\mu(x^+) + \lim_{x^-\to x} D\mu(x^-)\right).$$

Simply use two QBX applications.

- *Predictably benign spectral behavior* at high frequencies.

Important for iterative solvers (e.g. GMRES)

Not many competing schemes have that!

Outline

BVP on a smooth domain

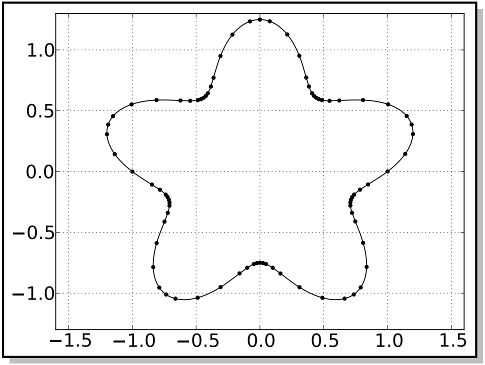

BC	Side	k	p	$M = 70$	$M = 105$	$M = 130$	EOC
Dir.	int	1	1	1.7e-04 (16)	7.8e-05 (16)	8.0e-05 (16)	**1.4**
			3	1.8e-06 (16)	4.4e-07 (16)	1.7e-07 (16)	**3.8**
			5	5.7e-08 (16)	5.6e-09 (16)	3.7e-09 (16)	**4.6**
		6	1	2.9e-02 (25)	1.3e-02 (25)	7.6e-03 (24)	**2.2**
			3	8.1e-05 (25)	1.8e-05 (25)	6.1e-06 (25)	**4.1**
			5	9.0e-07 (25)	9.7e-08 (25)	2.0e-08 (25)	**6.1**
Neu.	int	1	1	7.5e-02 (20)	5.4e-02 (20)	4.1e-02 (21)	**0.9**
			3	9.7e-04 (19)	3.2e-04 (19)	1.5e-04 (19)	**3.0**
			5	1.2e-05 (18)	2.0e-06 (18)	6.6e-07 (18)	**4.7**
		6	1	3.7e-01 (61)	3.2e-01 (61)	2.4e-01 (61)	**0.7**
			3	2.8e-03 (60)	9.6e-04 (60)	4.4e-04 (60)	**3.0**
			5	4.1e-05 (50)	6.5e-06 (50)	1.8e-06 (50)	**4.9**

Dirichlet problem on corner domain

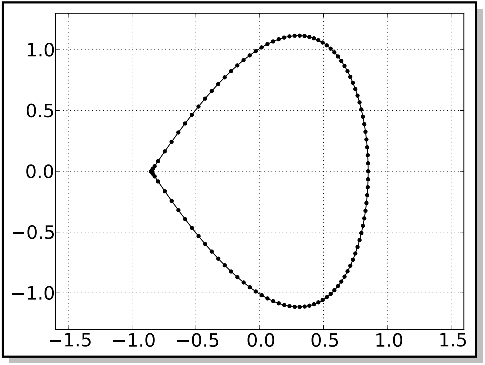

Dirich

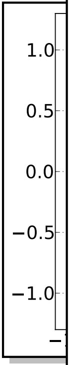

BC type	Side	k	p	$M = 80$	$M = 138$	EOC
Dirichlet	int	1	1	1.6e-03 (20)	2.1e-04 (21)	**3.7**
			3	5.9e-06 (20)	2.0e-07 (21)	**6.2**
			5	3.3e-08 (20)	3.3e-09 (21)	**4.2**
		6	1	6.9e-02 (38)	8.7e-03 (38)	**3.8**
			3	1.1e-04 (38)	3.8e-06 (38)	**6.1**
			5	1.0e-06 (38)	2.4e-08 (38)	**6.9**
	ext	1	1	3.4e-04 (19)	5.2e-05 (19)	**3.5**
			3	2.2e-06 (19)	8.2e-08 (19)	**6.0**
			5	1.3e-08 (19)	1.6e-09 (19)	**3.9**
		6	1	1.6e-02 (33)	1.9e-03 (33)	**3.9**
			3	4.5e-05 (33)	1.3e-06 (33)	**6.5**
			5	1.4e-07 (33)	1.3e-08 (33)	**4.4**

QBX in 3D

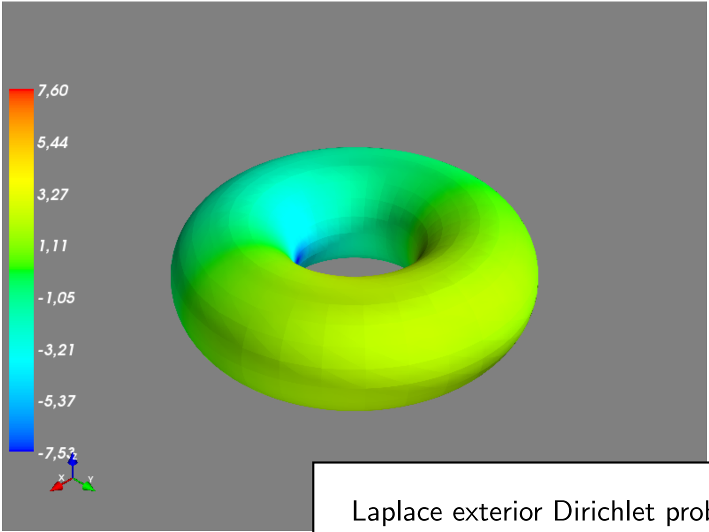

Laplace exterior Dirichlet problem

Outline

QBX: Why exciting?

Mathematically:

- General: Dimension, PDE, BC, kernel, surface discretization
- Fast enough for on-the-fly (non-stored) quadrature
- Benign conditioning
  (iterative methods $\rightarrow 10^{-15} \ll$ discr. error)
- Works easily for hypersingular kernels

✓

Computationally: (ongoing)

- Little data, many flops: high arithmetic intensity
- Good match for hierarchical parallelism
- Locally homogeneous, batched work

Questions?

?

Thank you for your attention!

`http://www.cs.illinois.edu/~andreask/`