CBMS Conference on Fast Direct Solvers

Dartmouth College

June 23 – June 27, 2014

Lecture 1: Introduction

Gunnar Martinsson

The University of Colorado at Boulder

Research support by:

Many thanks to everyone who made these lectures possible!

The organizers:

- **Alexander Barnett (Dartmouth College)**
- **Min Hyung Cho (Dartmouth College)**
- **Adrianna Gillman (Dartmouth College)**
- **Leslie Greengard (NYU and Simons Foundation)**

Financial and logistic support:

Dartmouth College **CBMS** **NSF**

For bibliographic notes, see: amath.colorado.edu/faculty/martinss/2014_CBMS

Contributors to the work presented:

Direct solvers: Tracy Babb, Alex Barnett, James Bremer, Eduardo Corona, Adrianna Gillman, Leslie Greengard, Sijia Hao, Terry Haut, Eric Michielssen, Vladimir Rokhlin, Mark Tygert, Patrick Young, Denis Zorin, etc

Randomized methods in matrix algebra: Nathan Halko, Edo Liberty, Vladimir Rokhlin, Joel Shkolnisky, Arthur Szlam, Joel Tropp, Mark Tygert, Sergey Voronin, etc

High-accuracy discretizations: Alex Barnett, James Bremer, Adrianna Gillman, Leslie Greengard, Sijia Hao, Patrick Young, etc

Coarse graining and model reduction: Ivo Babuška, Alexander Movchan, Greg Rodin, etc

While not formally collaborators, very useful discussions and suggestions from: Greg Beylkin, Laurent Demanet, Ming Gu, Johan Helsing, Mauro Maggioni, François Meyer, Michael O'Neil, Lexing Ying, Jianlin Xia, etc

Legend: green for students and postdocs, red for advisers and mentors.

These 10 lectures will address a central question in computational mathematics: How to efficiently compute approximate solutions to linear boundary value problems (BVPs) of the form

$$\text{(BVP)} \qquad \begin{cases} A\,u(\boldsymbol{x}) = g(\boldsymbol{x}), & \boldsymbol{x} \in \Omega, \\ B\,u(\boldsymbol{x}) = f(\boldsymbol{x}), & \boldsymbol{x} \in \Gamma, \end{cases}$$

where Ω is a domain in $\mathbb{R}^2$ or $\mathbb{R}^3$ with boundary Γ, and where A is an elliptic differential operator (constant coefficient, or not). Examples include:

- The Laplace equation.
- The equations of linear elasticity.
- Stokes' equation.
- Helmholtz' equation (at least at low and intermediate frequencies).
- Time-harmonic Maxwell (at least at low and intermediate frequencies).

Example: Poisson equation with Dirichlet boundary data:

$$\begin{cases} -\Delta\, u(\boldsymbol{x}) = g(\boldsymbol{x}), & \boldsymbol{x} \in \Omega, \\ u(\boldsymbol{x}) = f(\boldsymbol{x}), & \boldsymbol{x} \in \Gamma. \end{cases}$$

Bonus question: How to efficiently construct factorizations of (large!) matrices whose singular values decay.

Discretization of linear Boundary Value Problems

↙

Direct discretization of the differential operator via Finite Elements, Finite Differences, spectral composite methods, ...

↓

$N \times N$ discrete linear system.
Very large, sparse, ill-conditioned.

↓

Fast solvers:
iterative (multigrid), $O(N)$,
direct (nested dissection), $O(N^{3/2})$.

↘

Conversion of the BVP to a Boundary Integral Equation (BIE).

↓

Discretization of (BIE) using Nyström, collocation, BEM, ...

↓

$N \times N$ discrete linear system.
Moderate size, dense,
(often) well-conditioned.

↓

Iterative solver accelerated by fast matrix-vector multiplier, $O(N)$.

Discretization of linear Boundary Value Problems

↙

Direct discretization of the differential operator via Finite Elements, Finite Differences, spectral composite methods, ...

↓

$N \times N$ discrete linear system.
Very large, sparse, ill-conditioned.

↓

Fast solvers:
iterative (multigrid), $O(N)$,
direct (nested dissection), $O(N^{3/2})$.
$O(N)$ direct solvers.

↘

Conversion of the BVP to a Boundary Integral Equation (BIE).

↓

Discretization of (BIE) using Nyström, collocation, BEM, ...

↓

$N \times N$ discrete linear system.
Moderate size, dense,
(often) well-conditioned.

↓

Iterative solver accelerated by fast matrix-vector multiplier, $O(N)$.
$O(N)$ direct solvers.

What does a "direct" solver mean in this context?

Basically, it is a solver that is not "iterative" ...

Given a computational tolerance ε, and a linear system

$$\mathbf{A}\,\mathbf{u} = \mathbf{b}, \tag{2}$$

(where the system matrix **A** is often defined implicitly), a *direct solver* constructs an operator **S** such that

$$||\mathbf{A}^{-1} - \mathbf{S}|| \leq \varepsilon.$$

Then an approximate solution to (2) is obtained by simply evaluating

$$\mathbf{u}_{\text{approx}} = \mathbf{S}\,\mathbf{b}.$$

The matrix **S** is typically constructed in a *data-sparse format* (e.g. $\mathcal{H}$-matrix, HSS, etc) that allows the matrix-vector product **S b** to be evaluated rapidly.

Variation: Find factors **B** and **C** such that $||\mathbf{A} - \mathbf{B}\,\mathbf{C}|| \leq \varepsilon$, and linear solves involving the matrices **B** and **C** are fast. (LU-decomposition, Cholesky, *etc*.)

"Iterative" versus "direct" solvers

Two classes of methods for solving an $N \times N$ linear algebraic system

$$\mathbf{A}\,\mathbf{u} = \mathbf{b}.$$

Iterative methods:

Examples: GMRES, conjugate gradients, Gauss-Seidel, *etc*.

Construct a sequence of vectors $\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3, \ldots$ that (hopefully!) converge to the exact solution.

Many iterative methods access **A** only via its action on vectors.

Often require problem specific preconditioners.

High performance when they work well. $O(N)$ solvers.

Direct methods:

Examples: Gaussian elimination, LU factorizations, matrix inversion, *etc*.

Always give an answer. Deterministic.

Robust. No convergence analysis.

Great for multiple right hand sides.

Have often been considered too slow for high performance computing.

(Directly access elements or blocks of **A**.)

(Exact except for rounding errors.)

Advantages of direct solvers over iterative solvers:

1. Applications that require a very large number of solves for a fixed operator:
   - Molecular dynamics.
   - Scattering problems.
   - Optimal design. (Local updates to the system matrix are cheap.)

   A couple of orders of magnitude speed-up is often possible.

2. Solving problems intractable to iterative methods (singular values do not “cluster”):
   - Scattering problems near resonant frequencies.
   - Ill-conditioning due to geometry (elongated domains, percolation, etc).
   - Ill-conditioning due to lazy handling of corners, cusps, etc.
   - Finite element and finite difference discretizations.

   Scattering problems intractable to existing methods can (sometimes) be solved.

3. Direct solvers can be adapted to construct spectral decompositions:
   - Analysis of vibrating structures. Acoustics.
   - Buckling of mechanical structures.
   - Wave guides, bandgap materials, etc.

   Work in progress . . .

Advantages of direct solvers over iterative solvers, continued:

Perhaps most important: **Engineering considerations.**

Direct methods tend to be more robust than iterative ones.

This makes them more suitable for “black-box” implementations.

Commercial software developers appear to avoid implementing iterative solvers whenever possible. (Sometimes for good reasons.)

The effort to develop direct solvers aims to help in the development of general purpose software packages solving the basic linear boundary value problems of mathematical physics.

Advantages of direct solvers over iterative solvers, subtleties:

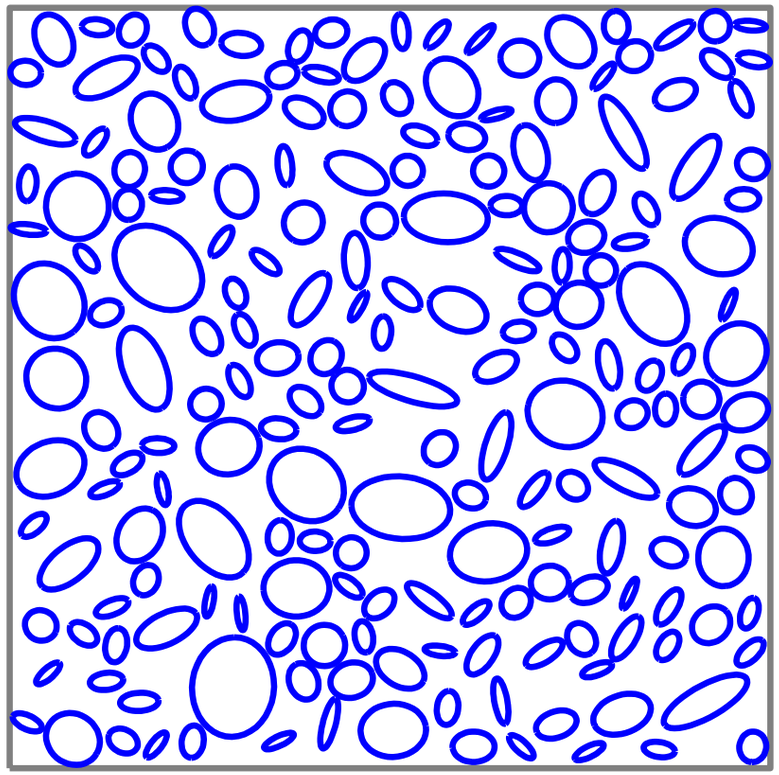

Numerical model reduction.

Advantages of direct solvers over iterative solvers, subtleties:

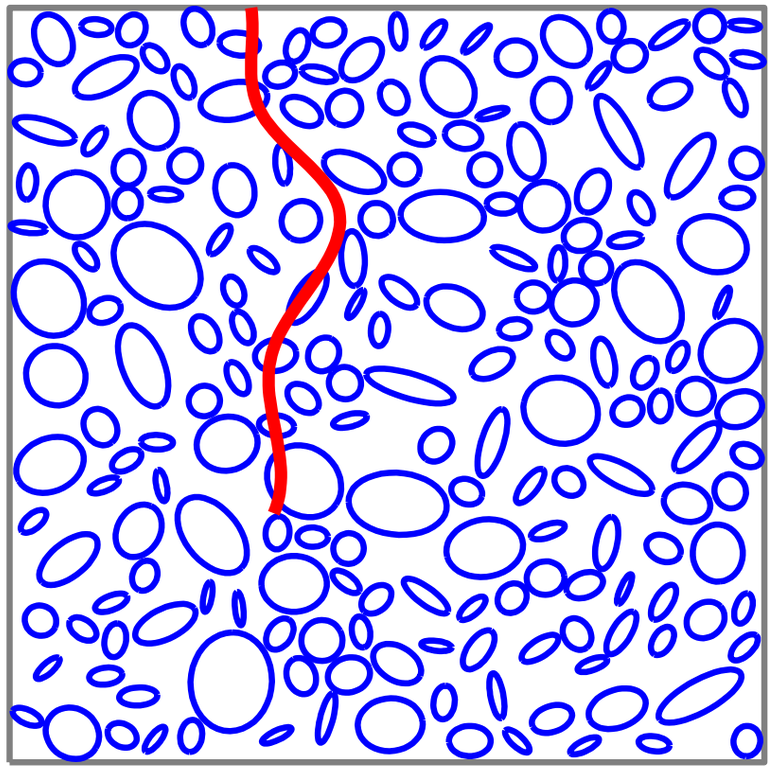

Numerical model reduction.

Advantages of direct solvers over iterative solvers, subtleties:

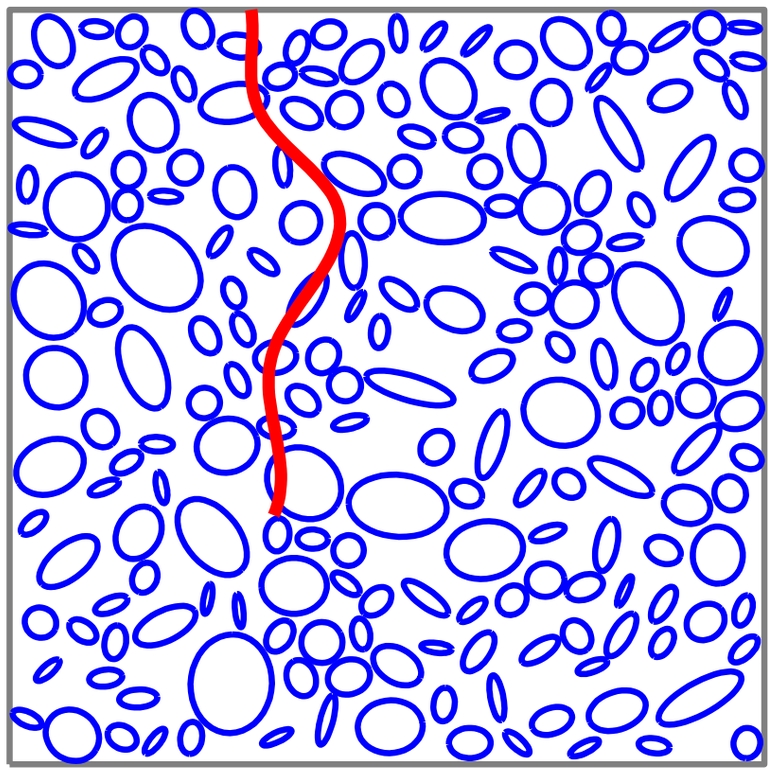

Numerical model reduction.

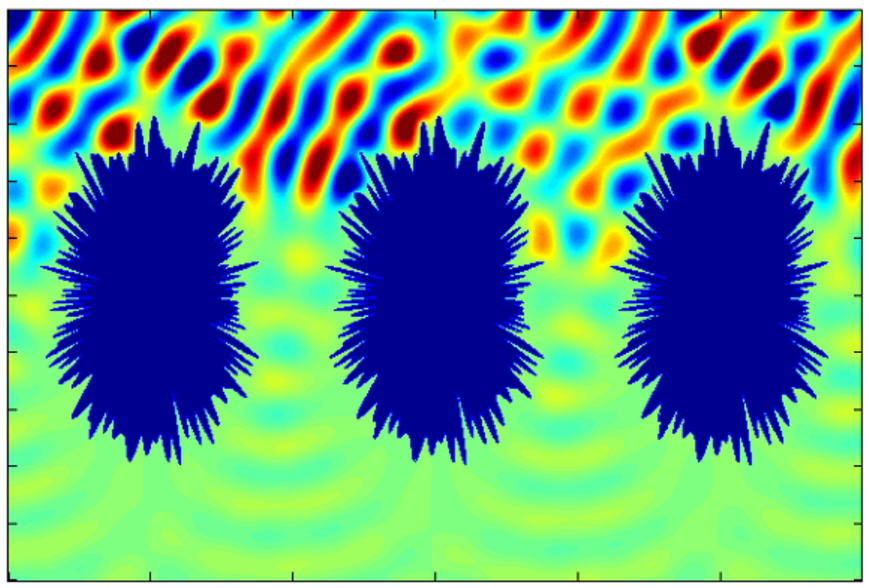

Accurate discretization of corners and edges.

Advantages of direct solvers over iterative solvers, subtleties:

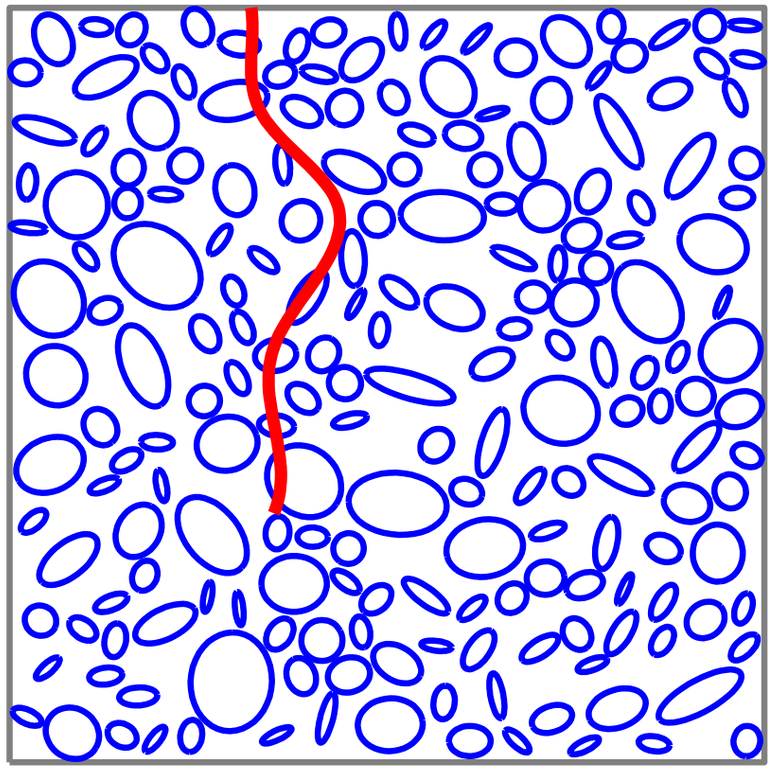

Numerical model reduction.

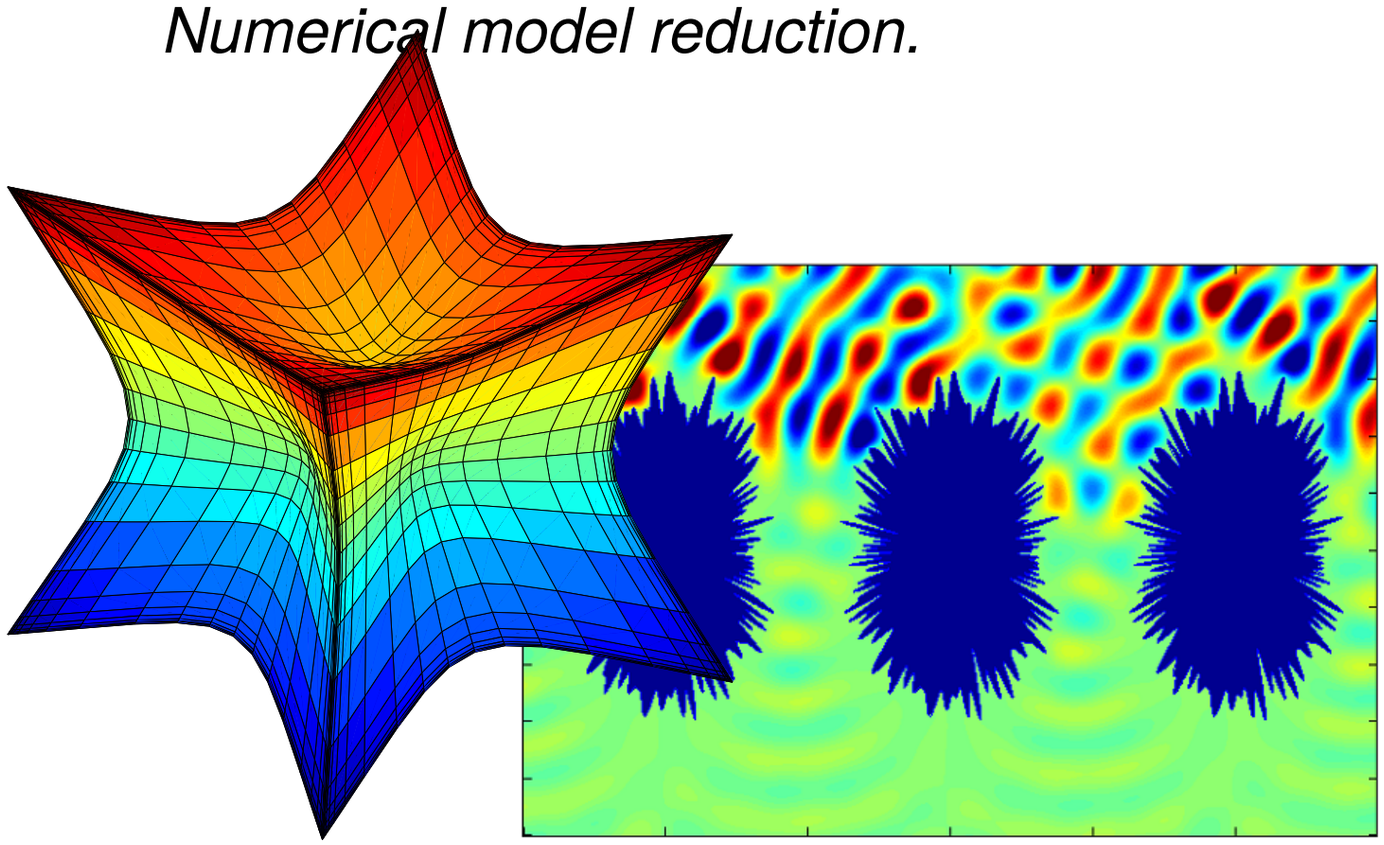

Accurate discretization of corners and edges.

Advantages of direct solvers over iterative solvers, subtleties:

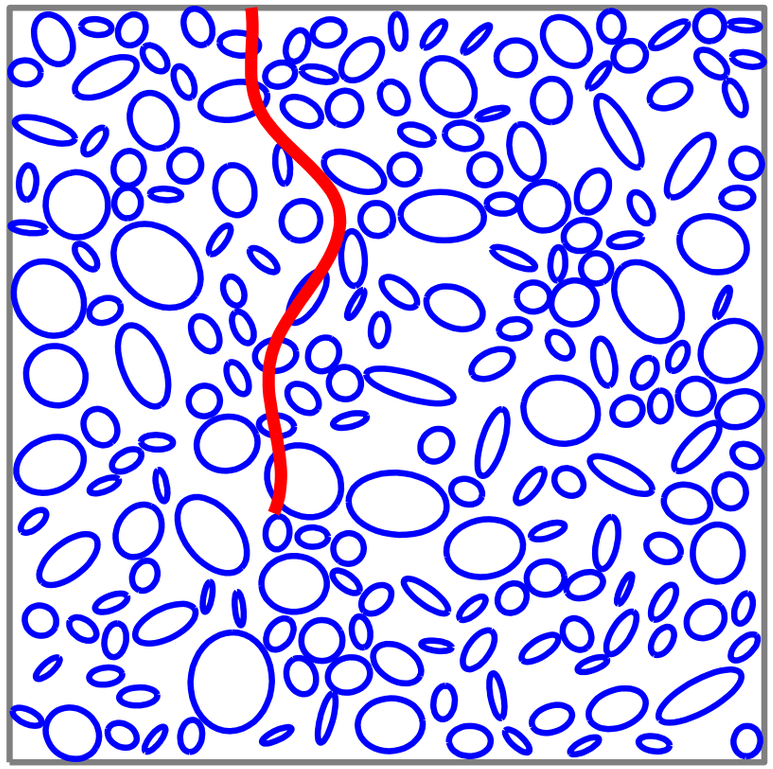

Numerical model reduction.

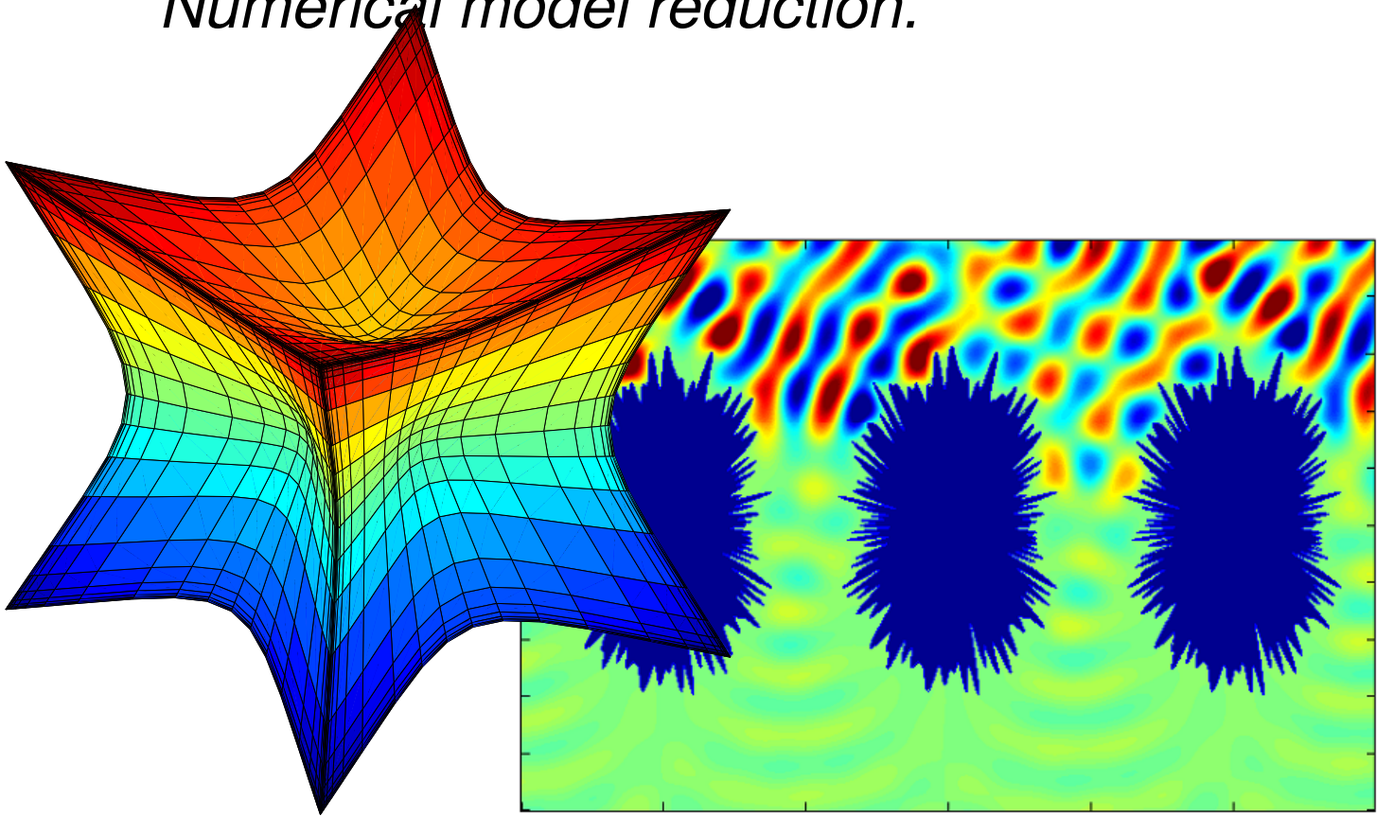

Accurate discretization of corners and edges.

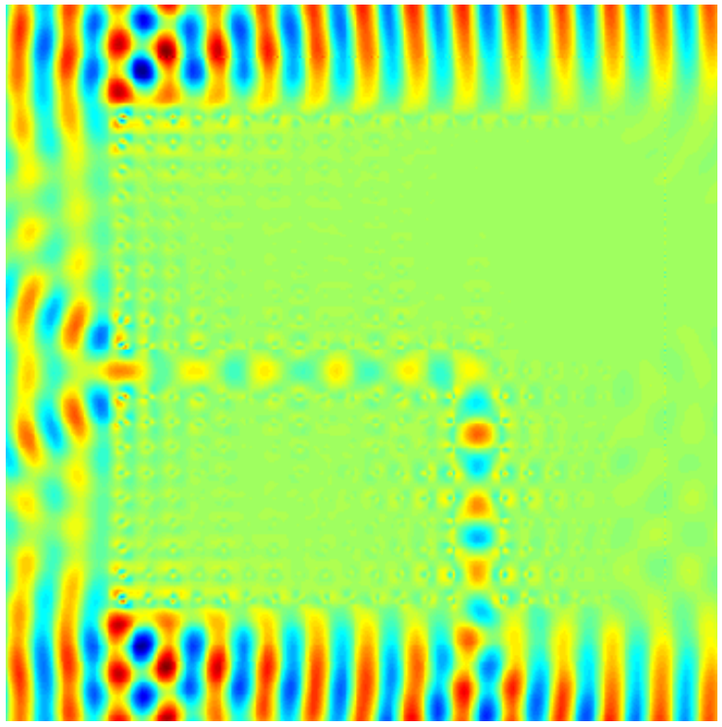

Operator algebra.
“Gluing” operator together.

Advantages of direct solvers over iterative solvers, subtleties:

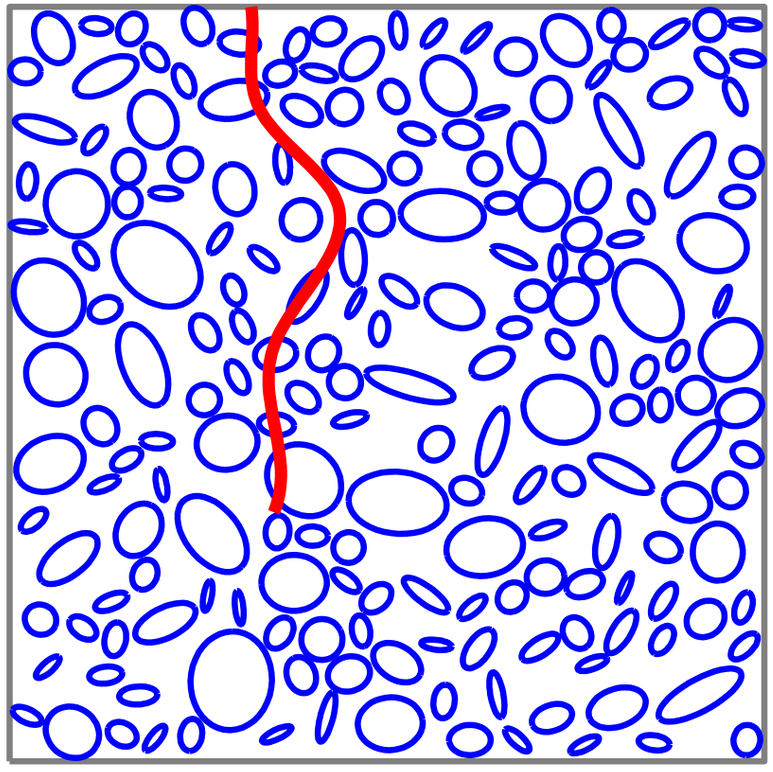

Numerical model reduction.

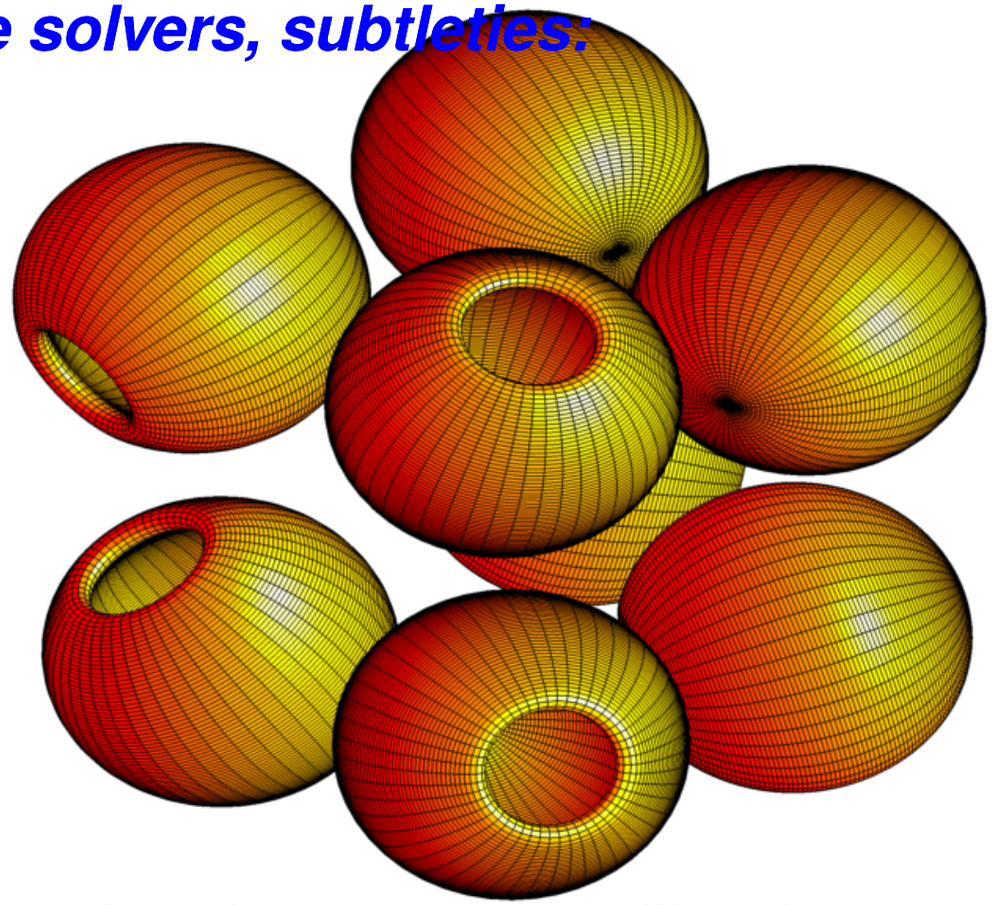

Accelerate iterative (!) solvers.

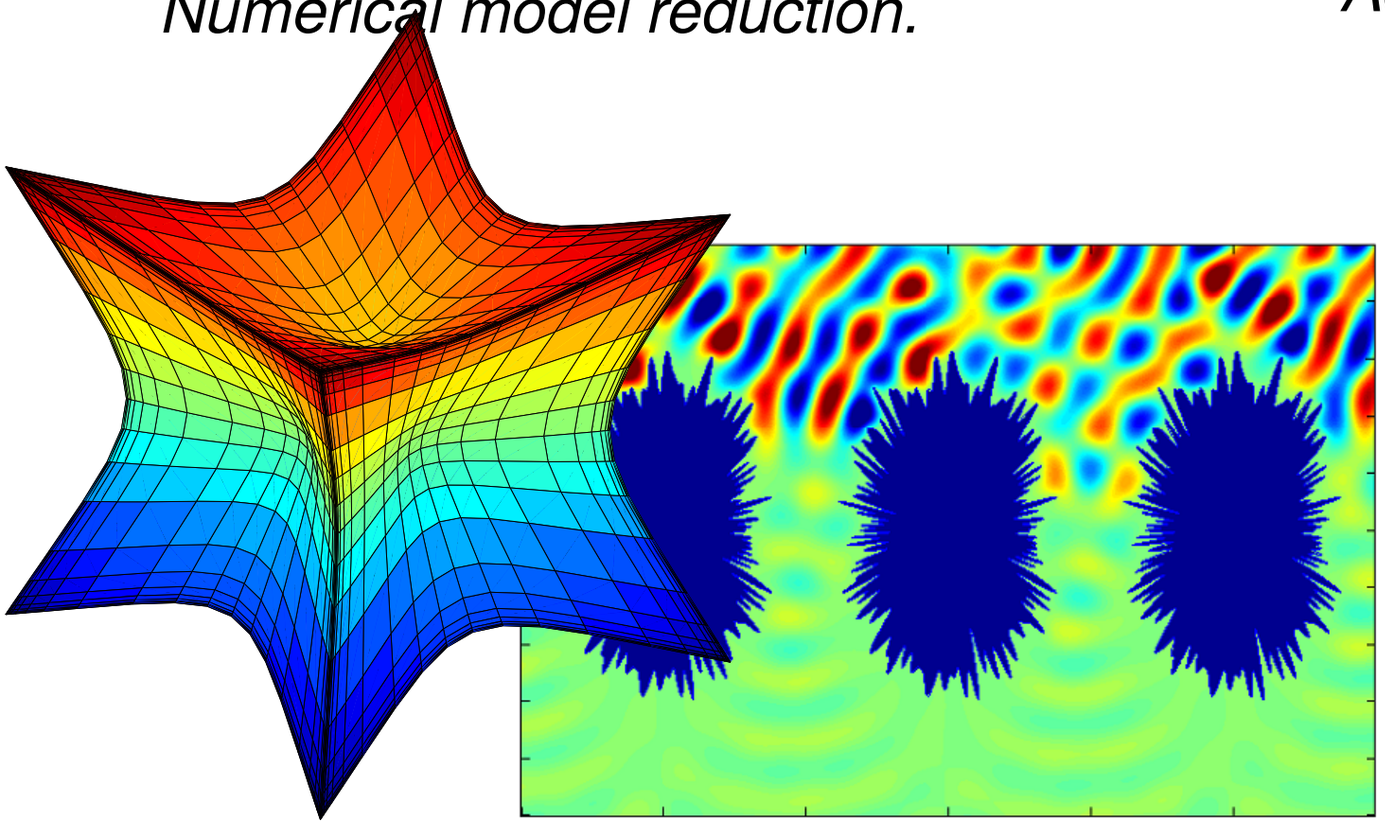

Accurate discretization of corners and edges.

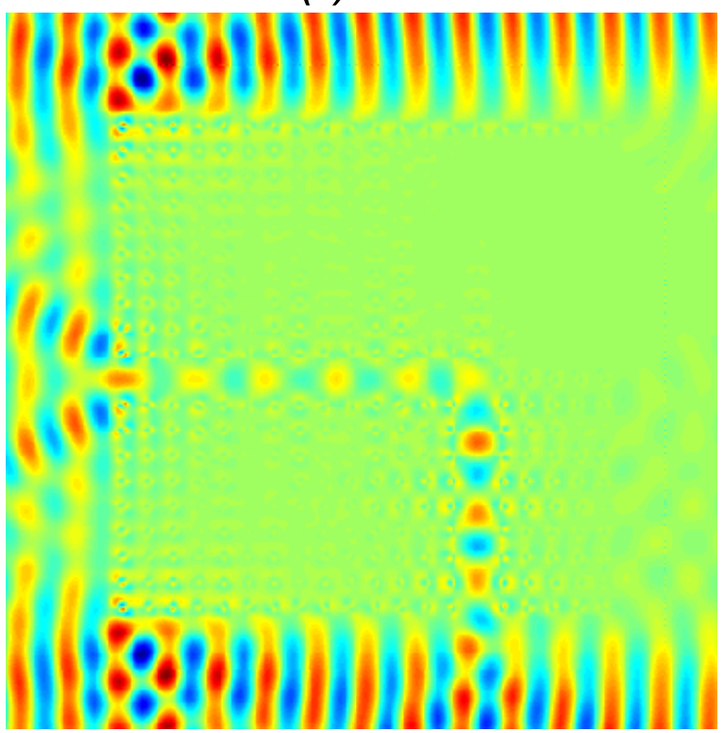

Operator algebra.
“Gluing” operator together.

There are fundamental reasons why we would expect direct solvers to be better

Consider the task of solving a classical boundary value problem such as

$$\begin{cases} -\Delta u(\boldsymbol{x}) = g(\boldsymbol{x}) & \boldsymbol{x} \in \Omega, \\ u(\boldsymbol{x}) = h(\boldsymbol{x}) & \boldsymbol{x} \in \Gamma. \end{cases} \tag{3}$$

The explicit solution formula takes the form

$$u(\boldsymbol{x}) = \int_{\Omega} G(\boldsymbol{x}, \boldsymbol{y})\, g(\boldsymbol{y})\, d\boldsymbol{y} + \int_{\Gamma} H(\boldsymbol{x}, \boldsymbol{y})\, h(\boldsymbol{y})\, dS(\boldsymbol{y}). \tag{4}$$

The functions G and H depend on Ω and are known analytically only for trivial domains.

A direct solver for (3) "numerically" constructs approximations to G and/or H.

Key point: The mathematical operator in (4) is *much* nicer than the operator in (3):

There are fundamental reasons why we would expect direct solvers to be better

Consider the task of solving a classical boundary value problem such as

$$\begin{cases} -\Delta u(\boldsymbol{x}) = g(\boldsymbol{x}) & \boldsymbol{x} \in \Omega, \\ u(\boldsymbol{x}) = h(\boldsymbol{x}) & \boldsymbol{x} \in \Gamma. \end{cases} \tag{3}$$

The explicit solution formula takes the form

$$u(\boldsymbol{x}) = \int_{\Omega} G(\boldsymbol{x}, \boldsymbol{y})\, g(\boldsymbol{y})\, d\boldsymbol{y} + \int_{\Gamma} H(\boldsymbol{x}, \boldsymbol{y})\, h(\boldsymbol{y})\, dS(\boldsymbol{y}). \tag{4}$$

The functions G and H depend on Ω and are known analytically only for trivial domains.

A direct solver for (3) "numerically" constructs approximations to G and/or H.

Key point: The mathematical operator in (4) is *much* nicer than the operator in (3):

- The operator in (4) is *a smoothing operator.*

There are fundamental reasons why we would expect direct solvers to be better

Consider the task of solving a classical boundary value problem such as

$$\begin{cases} -\Delta u(\boldsymbol{x}) = g(\boldsymbol{x}) & \boldsymbol{x} \in \Omega, \\ u(\boldsymbol{x}) = h(\boldsymbol{x}) & \boldsymbol{x} \in \Gamma. \end{cases} \tag{3}$$

The explicit solution formula takes the form

$$u(\boldsymbol{x}) = \int_{\Omega} G(\boldsymbol{x}, \boldsymbol{y})\, g(\boldsymbol{y})\, d\boldsymbol{y} + \int_{\Gamma} H(\boldsymbol{x}, \boldsymbol{y})\, h(\boldsymbol{y})\, dS(\boldsymbol{y}). \tag{4}$$

The functions G and H depend on Ω and are known analytically only for trivial domains.

A direct solver for (3) "numerically" constructs approximations to G and/or H.

Key point: The mathematical operator in (4) is *much* nicer than the operator in (3):

- The operator in (4) is *a smoothing operator.*
- The operator in (4) is *bounded.*

There are fundamental reasons why we would expect direct solvers to be better

Consider the task of solving a classical boundary value problem such as

$$\begin{cases} -\Delta u(\boldsymbol{x}) = g(\boldsymbol{x}) & \boldsymbol{x} \in \Omega, \\ u(\boldsymbol{x}) = h(\boldsymbol{x}) & \boldsymbol{x} \in \Gamma. \end{cases} \tag{3}$$

The explicit solution formula takes the form

$$u(\boldsymbol{x}) = \int_{\Omega} G(\boldsymbol{x}, \boldsymbol{y})\, g(\boldsymbol{y})\, d\boldsymbol{y} + \int_{\Gamma} H(\boldsymbol{x}, \boldsymbol{y})\, h(\boldsymbol{y})\, dS(\boldsymbol{y}). \tag{4}$$

The functions G and H depend on Ω and are known analytically only for trivial domains.

A direct solver for (3) "numerically" constructs approximations to G and/or H.

Key point: The mathematical operator in (4) is *much* nicer than the operator in (3):

- The operator in (4) is *a smoothing operator.*
- The operator in (4) is *bounded.*
- The operator in (4) is in many ways *stable* — small changes in g and h often lead to small changes in u.

Approximating (4) is inherently more tractable.

Example — Poisson equation on $\mathbb{R}^2$

Consider the Poisson equation

$$-\Delta u(\boldsymbol{x}) = g(\boldsymbol{x}) \qquad \boldsymbol{x} \in \mathbb{R}^2, \tag{5}$$

coupled with suitable boundary conditions at infinity to ensure uniqueness.

We can solve (5) analytically. For instance, use a Fourier transform,

$$[\mathcal{F}u](\boldsymbol{t}) = \hat{u}(\boldsymbol{t}) = \frac{1}{2\pi}\int_{\mathbb{R}^2} e^{-i\,\boldsymbol{x}\cdot\boldsymbol{t}}\, u(\boldsymbol{x})\, d\boldsymbol{x}.$$

The Laplace operator in Fourier space becomes a simple multiplication operator,

$$[\mathcal{F}(-\Delta u)](\boldsymbol{t}) = |\boldsymbol{t}|^2\, \hat{u}(\boldsymbol{t}),$$

so equation (5) in Fourier space becomes *diagonal*

$$|\boldsymbol{t}|^2\, \hat{u}(\boldsymbol{t}) = \hat{g}(\boldsymbol{t}), \qquad \boldsymbol{t} \in \mathbb{R}^2. \tag{6}$$

Solving (6) is easy:

$$\hat{u}(\boldsymbol{t}) = \frac{1}{|\boldsymbol{t}|^2}\,\hat{g}(\boldsymbol{t}),$$

and then we map back to physical space to obtain

$$u(\boldsymbol{x}) = [\mathcal{F}^{-1}\hat{u}](\boldsymbol{x}) = \frac{1}{2\pi}\int_{\mathbb{R}^2} e^{i\,\boldsymbol{x}\cdot\boldsymbol{t}}\, \frac{1}{|\boldsymbol{t}|^2}\hat{g}(\boldsymbol{t})\, d\boldsymbol{t}.$$

Example — Poisson equation on $\mathbb{R}^2$

Consider the Poisson equation

$$-\Delta u(\boldsymbol{x}) = g(\boldsymbol{x}) \qquad \boldsymbol{x} \in \mathbb{R}^2, \tag{7}$$

coupled with suitable boundary conditions at infinity to ensure uniqueness. Recall that

$$u(\boldsymbol{x}) = [\mathcal{F}^{-1}\hat{u}](\boldsymbol{x}) = \frac{1}{2\pi}\int_{\mathbb{R}^2} e^{i\,\boldsymbol{x}\cdot\boldsymbol{t}}\,\frac{1}{|\boldsymbol{t}|^2}\hat{g}(\boldsymbol{t})\,d\boldsymbol{t}.$$

We can write the solution operator as an operator in physical space by observing that *multiplication* in Fourier space corresponds to *convolution* in physical space.

Setting $\psi(t) = |\boldsymbol{t}|^{-2}$, we find

$$u = \mathcal{F}^{-1}(\psi\,\hat{g}) = \left(\mathcal{F}^{-1}\psi\right) * \left(\mathcal{F}^{-1}\hat{g}\right) = \left(\mathcal{F}^{-1}\psi\right) * g.$$

Introduce the fundamental solution $\phi = \mathcal{F}^{-1}\psi$. It can be shown that

$$\phi(\boldsymbol{x}) = [\mathcal{F}^{-1}\psi](\boldsymbol{x}) = -\frac{1}{2\pi}\log|\boldsymbol{x}|.$$

Then we obtain

$$u(\boldsymbol{x}) = \int_{\mathbb{R}^2} \phi(\boldsymbol{x}-\boldsymbol{y})\,f(\boldsymbol{x})\,d\boldsymbol{x}, \qquad \boldsymbol{x} \in \mathbb{R}^2. \tag{8}$$

Properties of the solution operator $\mathcal{G}$.

Recall the Poisson equation

$$-\Delta u(\boldsymbol{x}) = g(\boldsymbol{x}) \qquad \boldsymbol{x} \in \mathbb{R}^2. \tag{9}$$

Let $\phi(\boldsymbol{x}) = -\dfrac{1}{2\pi} \log |\boldsymbol{x}|$ denote the fundamental solution, and let $\mathcal{G}$ denote the solution operator,

$$u(\boldsymbol{x}) = [\mathcal{G}g](\boldsymbol{x}) = [\phi * g](\boldsymbol{x}) = \int_{\mathbb{R}^2} \phi(\boldsymbol{x} - \boldsymbol{y})\, g(\boldsymbol{x})\, d\boldsymbol{x}, \qquad \boldsymbol{x} \in \mathbb{R}^2$$

Observation 1 (nice): Viewed as an operator on $\mathbb{R}^2$, $\mathcal{G}$ is nicely smoothing, u has two more derivatives than g.

Properties of the solution operator $\mathcal{G}$.

Recall the Poisson equation

$$-\Delta u(\boldsymbol{x}) = g(\boldsymbol{x}) \qquad \boldsymbol{x} \in \mathbb{R}^2. \tag{9}$$

Let $\phi(\boldsymbol{x}) = -\dfrac{1}{2\pi} \log |\boldsymbol{x}|$ denote the fundamental solution, and let $\mathcal{G}$ denote the solution operator,

$$u(\boldsymbol{x}) = [\mathcal{G}g](\boldsymbol{x}) = [\phi * g](\boldsymbol{x}) = \int_{\mathbb{R}^2} \phi(\boldsymbol{x} - \boldsymbol{y})\, g(\boldsymbol{x})\, d\boldsymbol{x}, \qquad \boldsymbol{x} \in \mathbb{R}^2$$

Observation 1 (nice): Viewed as an operator on $\mathbb{R}^2$, $\mathcal{G}$ is nicely smoothing, u has two more derivatives than g.

Observation 2 (extremely nice!): Let Ω_1 and Ω_2 denote two subsets of $\mathbb{R}^2$ that are "well-separated" (to be defined!). Then the *restriction* of $\mathcal{G}$, to say

$$\mathcal{G}_{1,2} : L^2(\Omega_2) \to L^2(\Omega_1),$$

is *very strongly* smoothing. In fact, $\mathcal{G}_{1,2} g \in C^\infty(\Omega_2)$. The singular values of $\mathcal{G}_{1,2}$ decay exponentially fast.

Properties of the solution operator $\mathcal{G}$.

Recall the Poisson equation

$$-\Delta u(\boldsymbol{x}) = g(\boldsymbol{x}) \qquad \boldsymbol{x} \in \mathbb{R}^2. \tag{9}$$

Let $\phi(\boldsymbol{x}) = -\dfrac{1}{2\pi} \log |\boldsymbol{x}|$ denote the fundamental solution, and let $\mathcal{G}$ denote the solution operator,

$$u(\boldsymbol{x}) = [\mathcal{G}g](\boldsymbol{x}) = [\phi * g](\boldsymbol{x}) = \int_{\mathbb{R}^2} \phi(\boldsymbol{x} - \boldsymbol{y})\, g(\boldsymbol{x})\, d\boldsymbol{x}, \qquad \boldsymbol{x} \in \mathbb{R}^2$$

Observation 1 (nice): Viewed as an operator on $\mathbb{R}^2$, $\mathcal{G}$ is nicely smoothing, u has two more derivatives than g.

Observation 2 (extremely nice!): Let Ω_1 and Ω_2 denote two subsets of $\mathbb{R}^2$ that are "well-separated" (to be defined!). Then the *restriction* of $\mathcal{G}$, to say

$$\mathcal{G}_{1,2} : L^2(\Omega_2) \to L^2(\Omega_1),$$

is *very strongly* smoothing. In fact, $\mathcal{G}_{1,2}g \in C^\infty(\Omega_2)$. The singular values of $\mathcal{G}_{1,2}$ decay exponentially fast.

Observation 3 (not so nice): The operator $\mathcal{G}$ is *global* — every part of the domain "talks to" every other part.

Example of solution of Laplace's equation $-\Delta u = g$

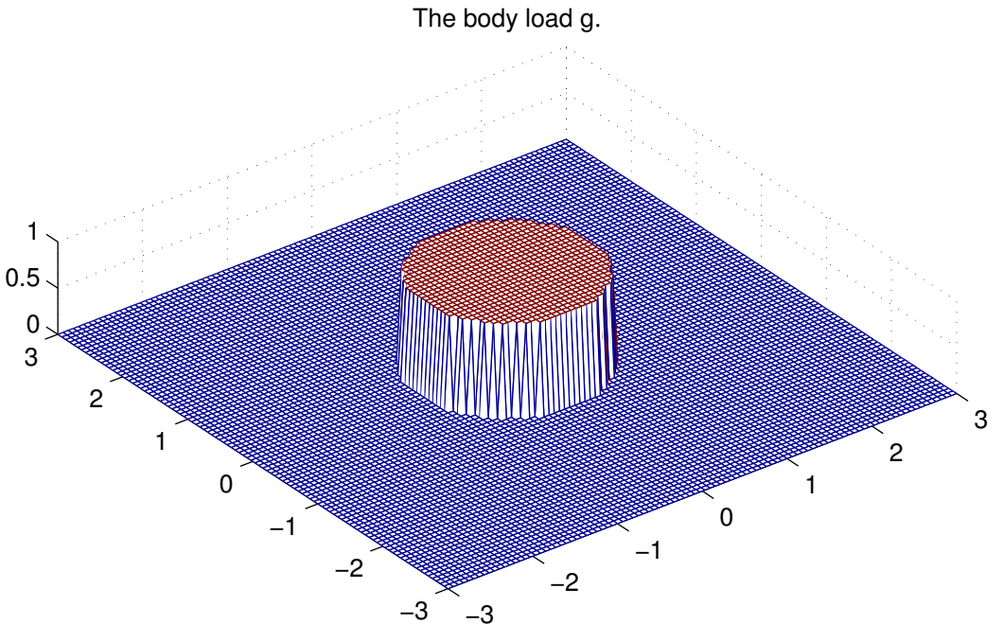

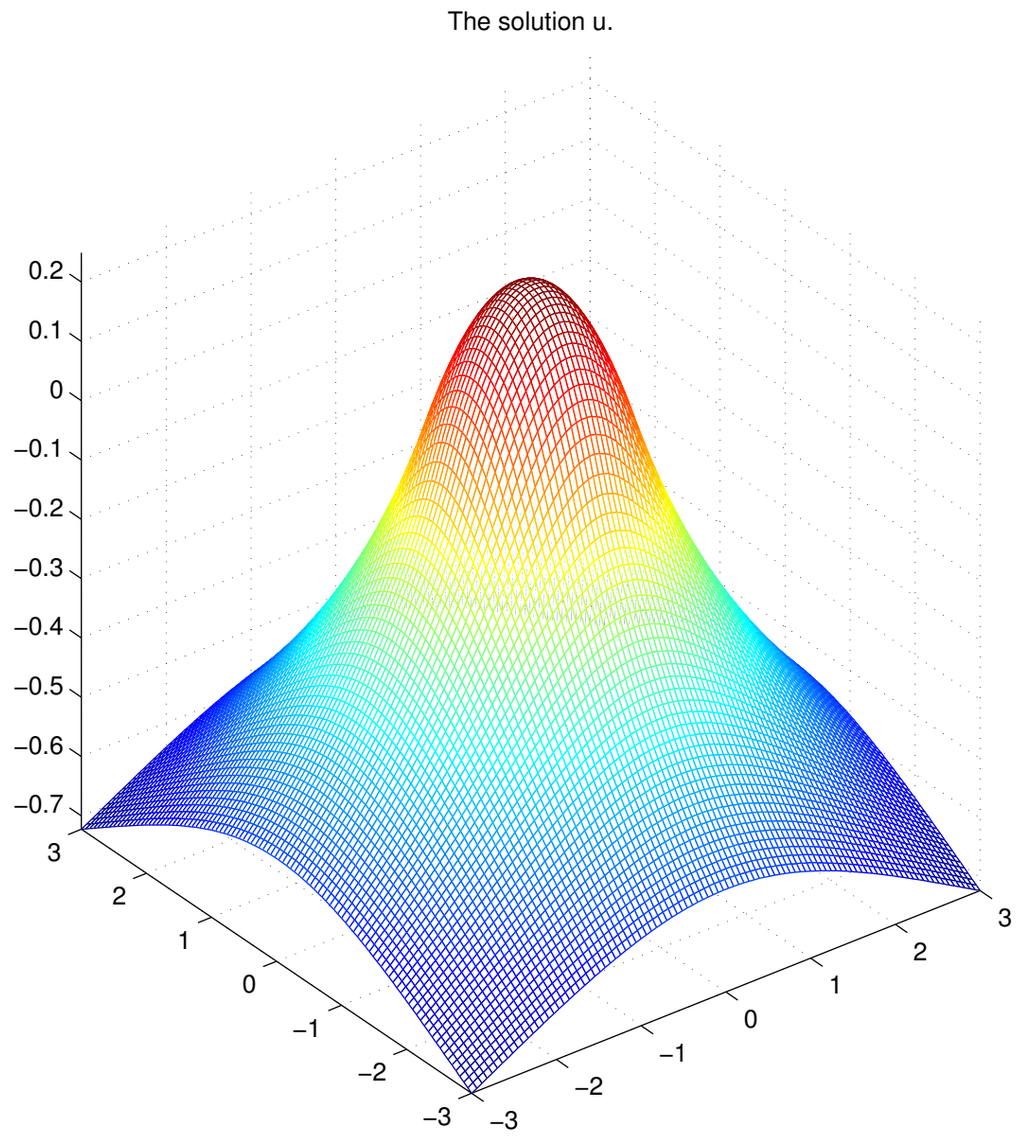

g is piecewise continuous, u is C^1

Example of solution of Laplace's equation $-\Delta u = g$

Now showing derivatives of solution from previous slide.

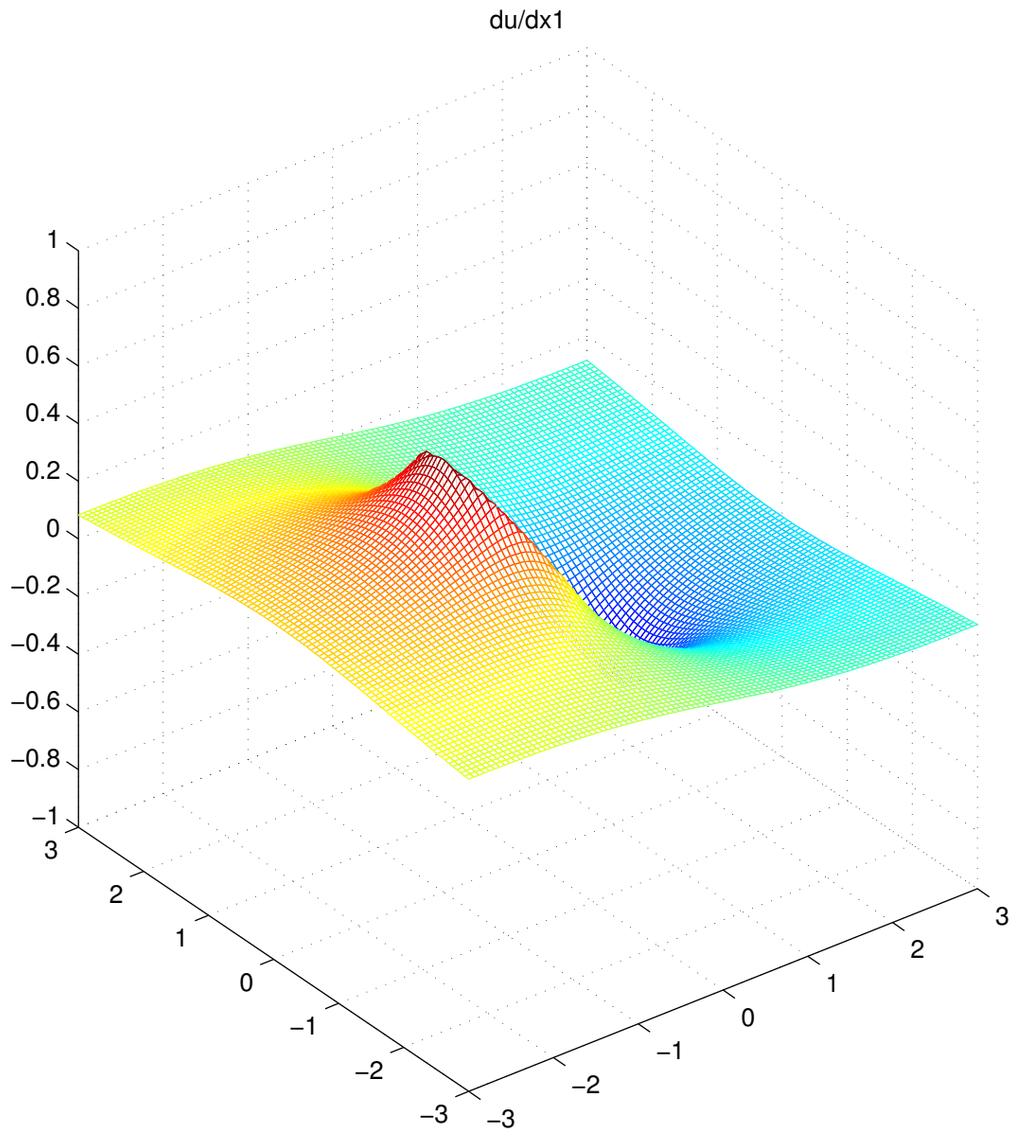

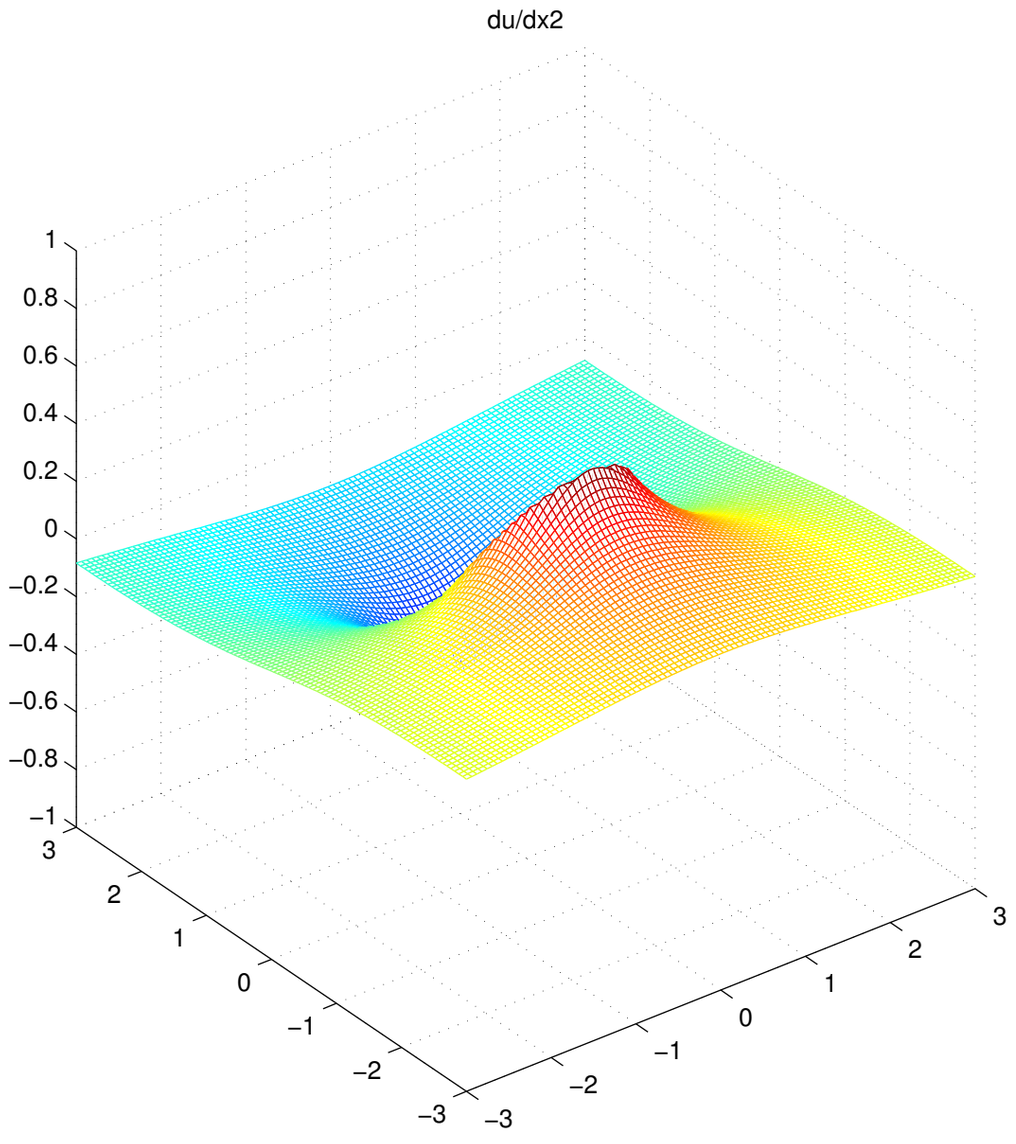

g is piecewise continuous, u is C^1

Example of solution of Laplace's equation $-\Delta u = g$

Recall that in the example we showed, g is piecewise continuous, and u is C^1.

Note however, that away from the lines of discontinuity of g, the function u is C^∞.

We can formalize this statement by splitting the fundamental solution

$$\phi = \phi_{\text{far}} + \phi_{\text{near}}$$

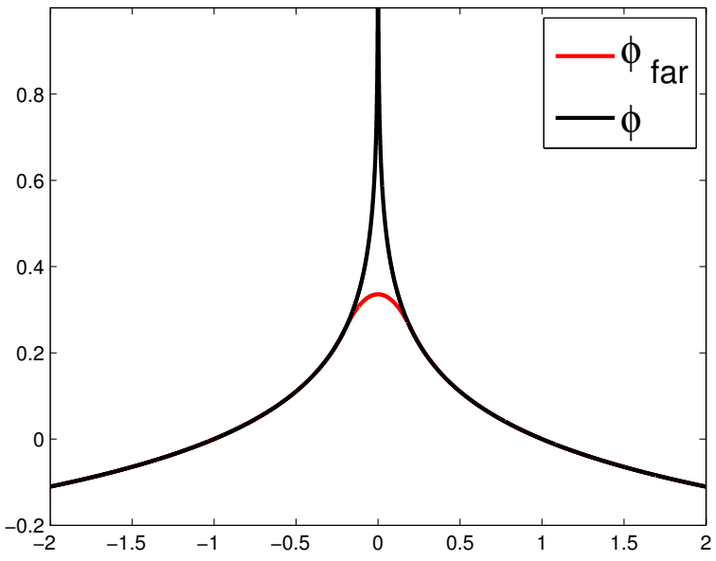

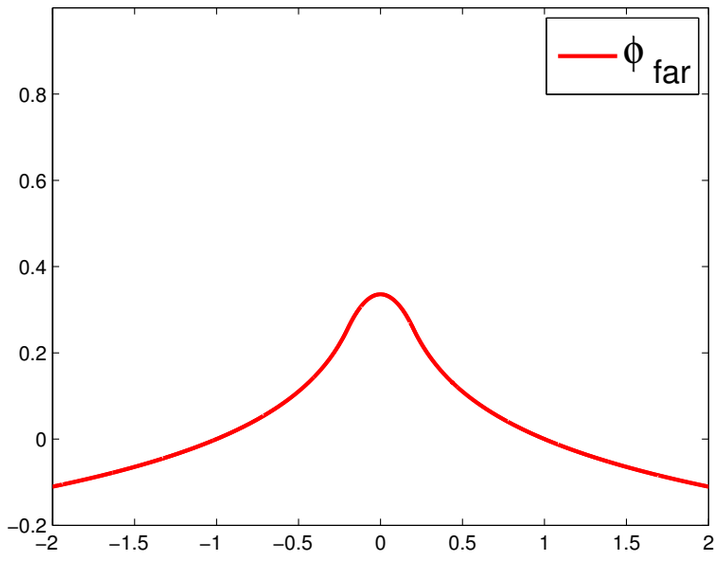

C^∞ function

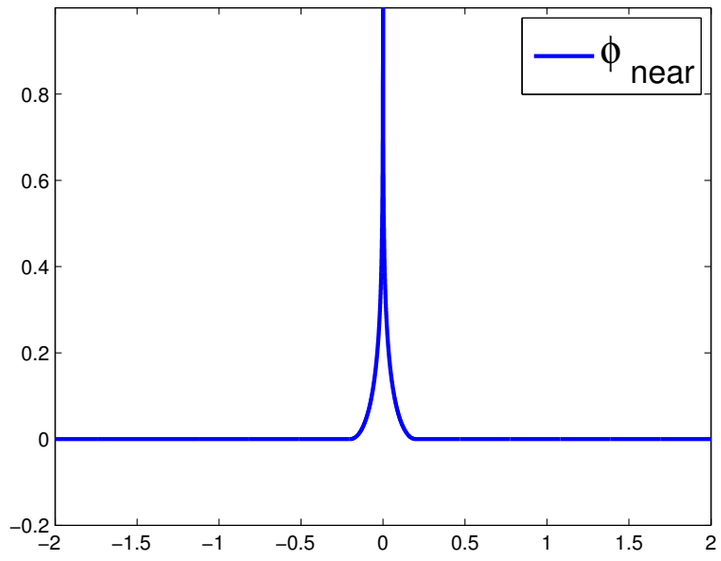

locally supported

C^∞ away from origin

and then writing the solution operator as

$$[\mathcal{G}g](\boldsymbol{x}) = [\mathcal{G}_{\text{near}}g](\boldsymbol{x}) + [\mathcal{G}_{\text{far}}g](\boldsymbol{x}) = [\phi_{\text{near}} * g](\boldsymbol{x}) + [\phi_{\text{far}} * g](\boldsymbol{x}).$$

Example of solution of Laplace's equation $-\Delta u = g$

Suppose we are given point charges $\{q_j\}_{j=1}^{5}$ in a "source domain" Ω_{S}.

We are interested in the potential in a "target domain" Ω_{t}.

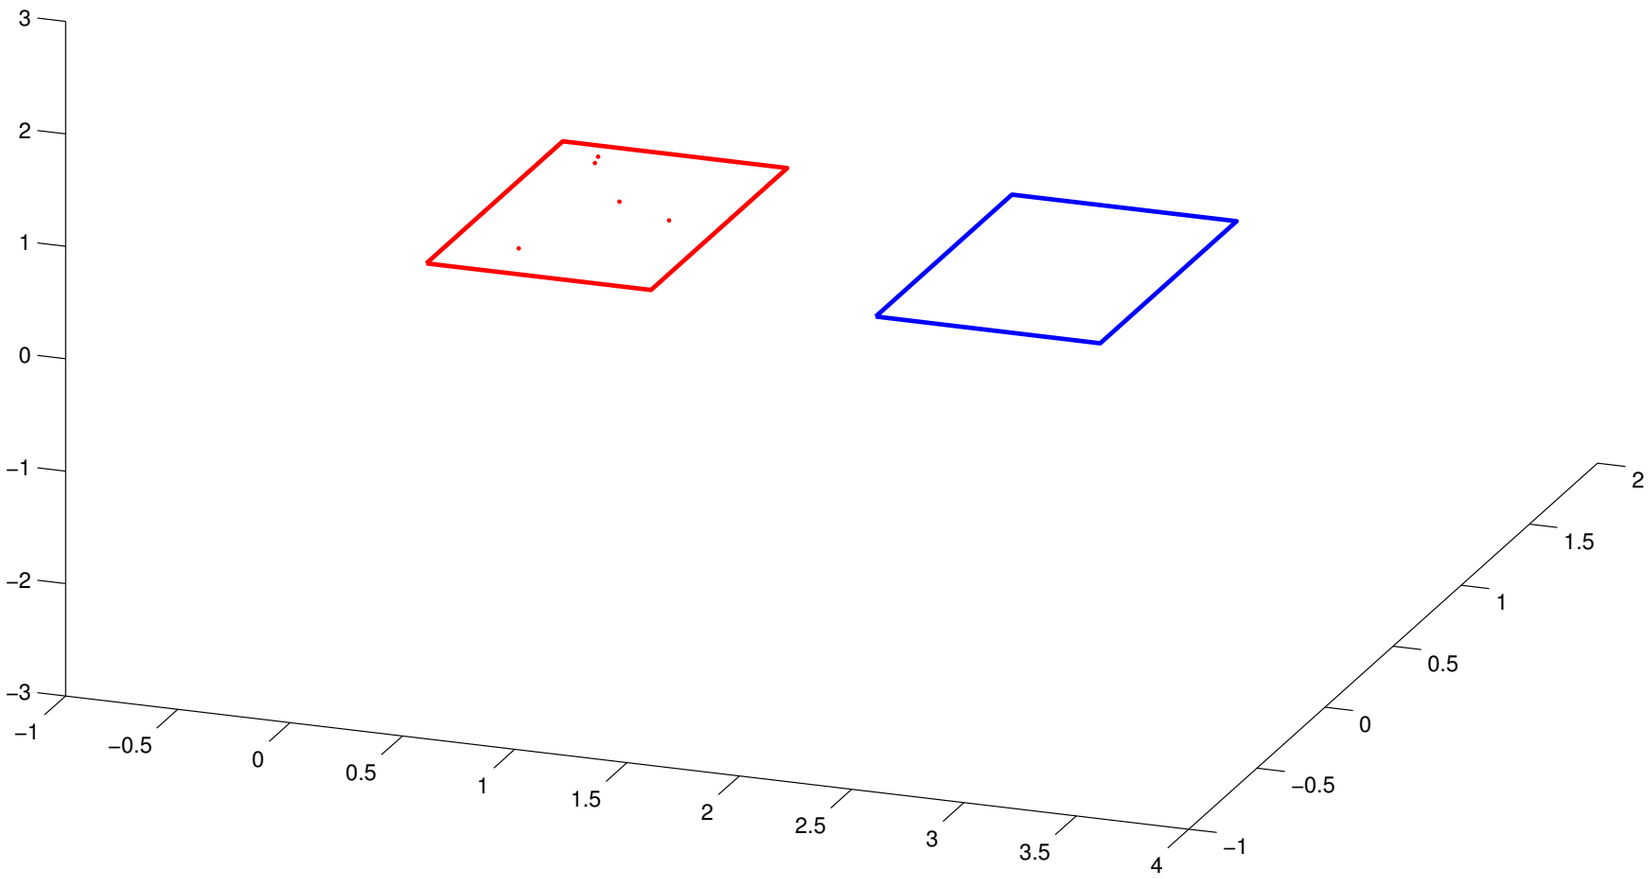

The source domain Ω_{S} (red) and the target domain Ω_{t} (blue).

Example of solution of Laplace's equation $-\Delta u = g$

Suppose we are given point charges $\{q_j\}_{j=1}^{5}$ in a "source domain" Ω_{s}.

We are interested in the potential in a "target domain" Ω_{t}.

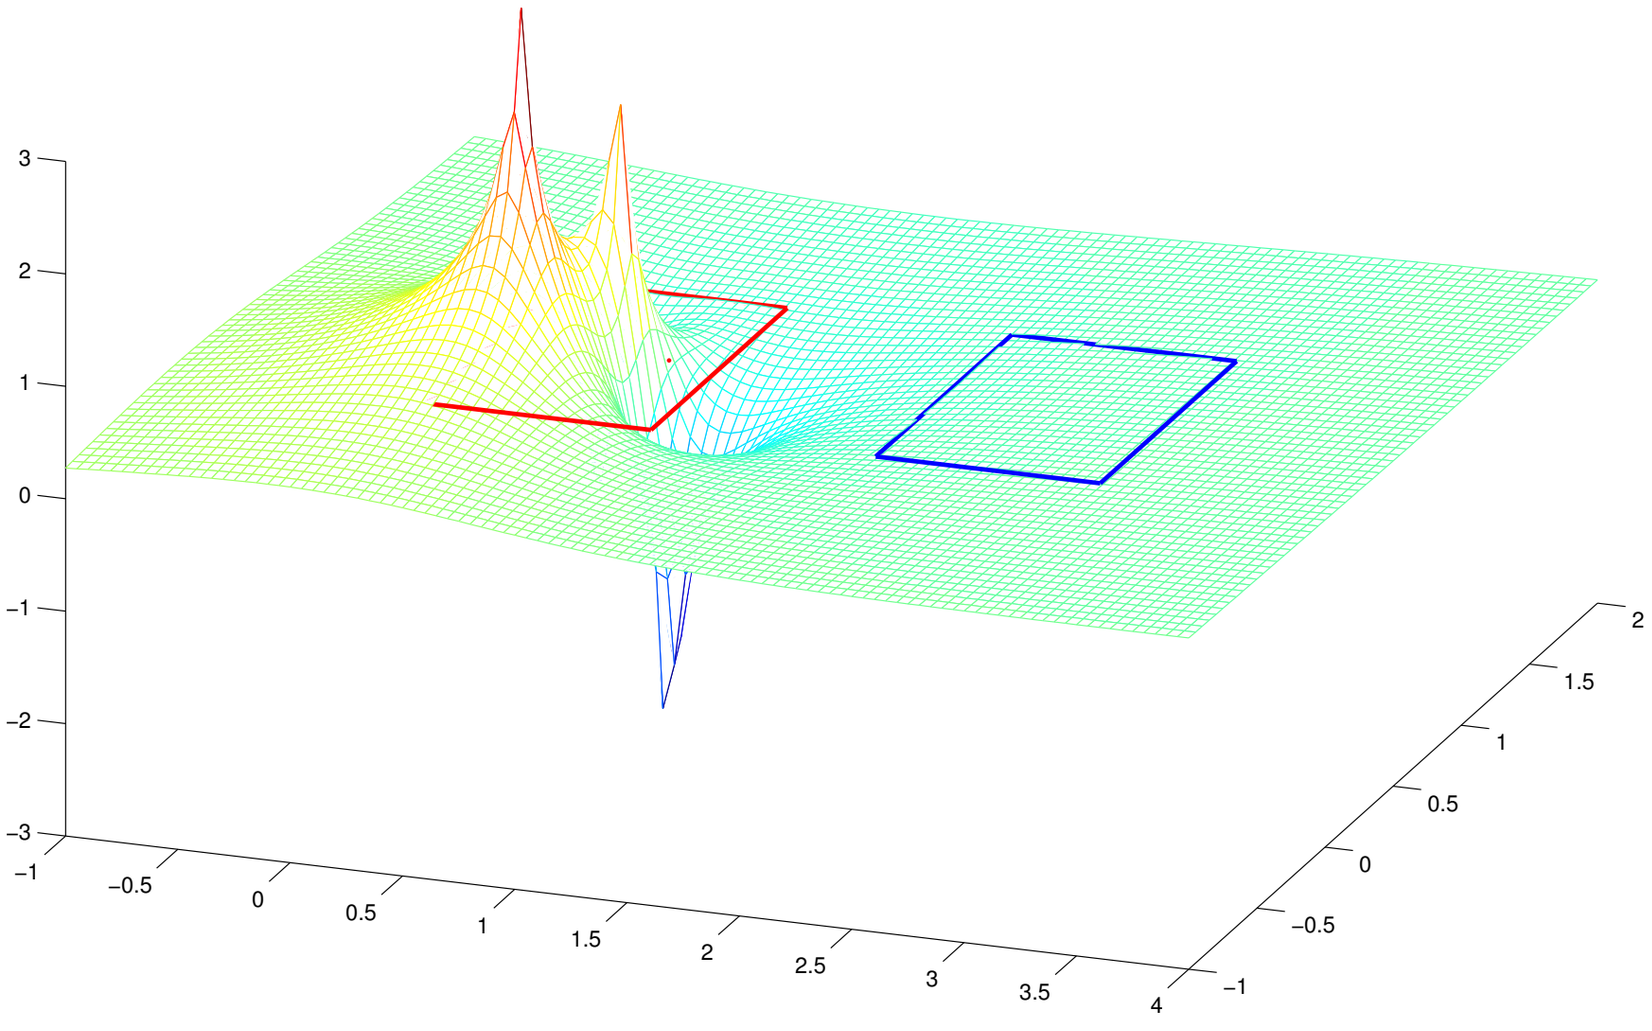

The field generated by the sources (truncated — the peaks go to infinity).

Example of solution of Laplace's equation $-\Delta u = g$

Suppose we are given point charges $\{q_j\}_{j=1}^{5}$ in a "source domain" Ω_{s}.

We are interested in the potential in a "target domain" Ω_{t}.

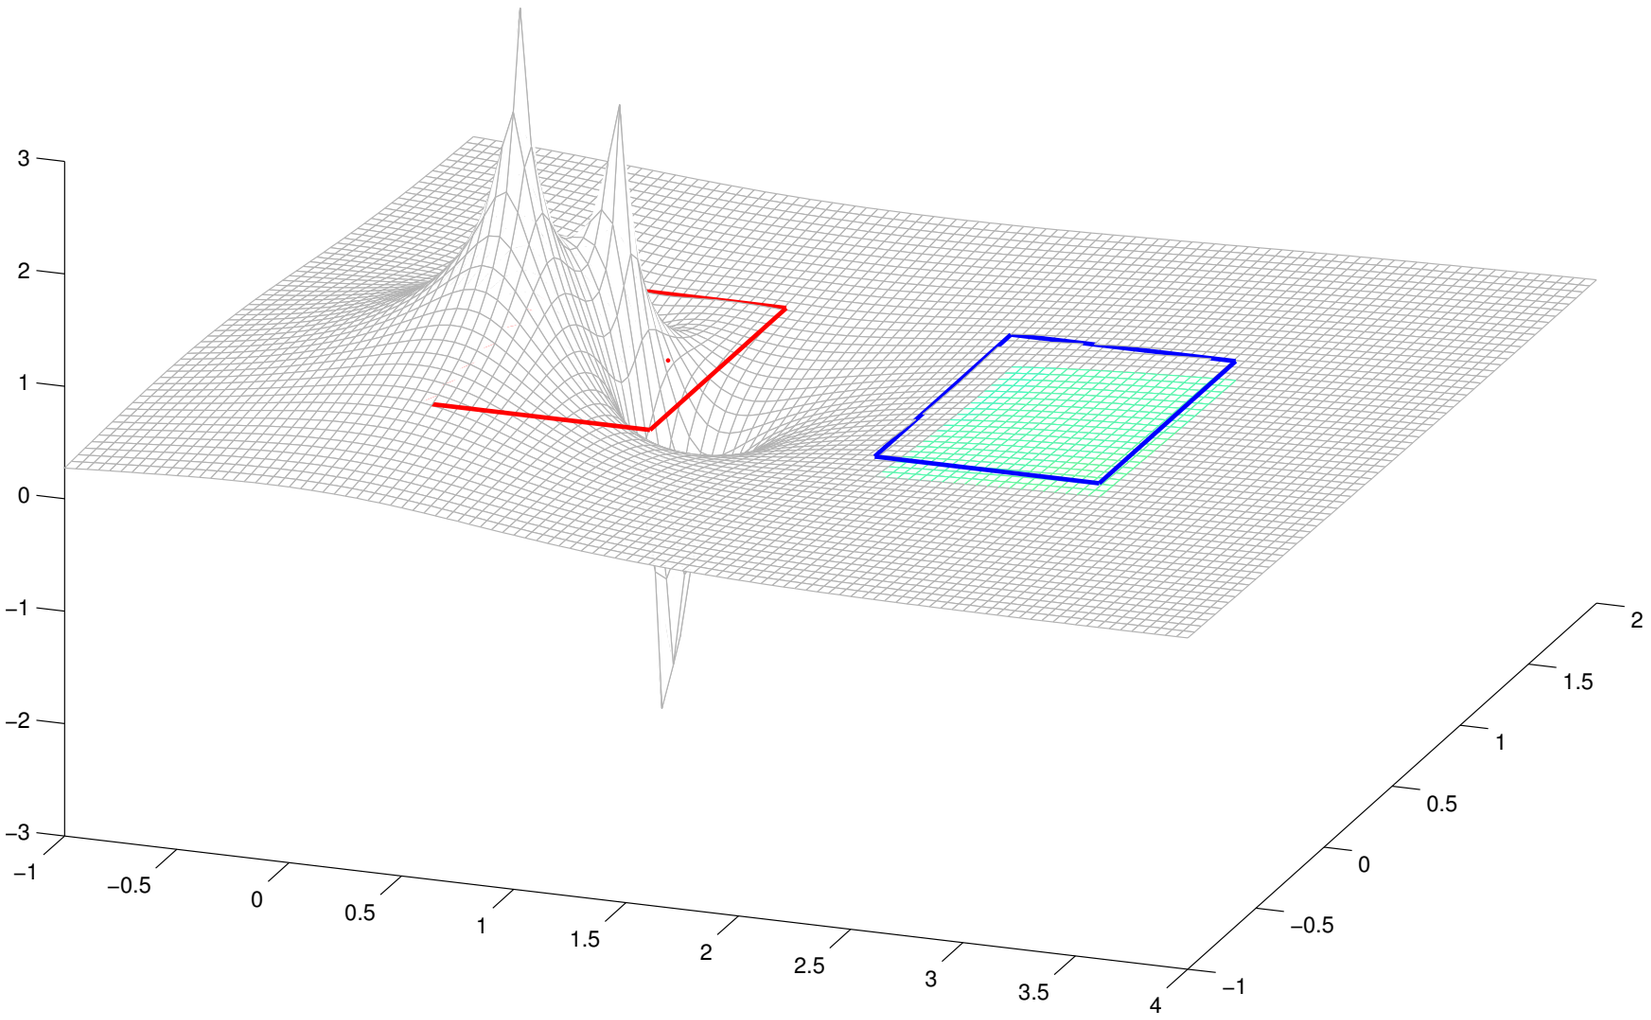

The field generated by the sources (truncated — the peaks go to infinity).

Example of solution of Laplace's equation $-\Delta u = g$

Suppose we are given point charges $\{q_j\}_{j=1}^{5}$ in a "source domain" Ω_{s}.

We are interested in the potential in a "target domain" Ω_{t}.

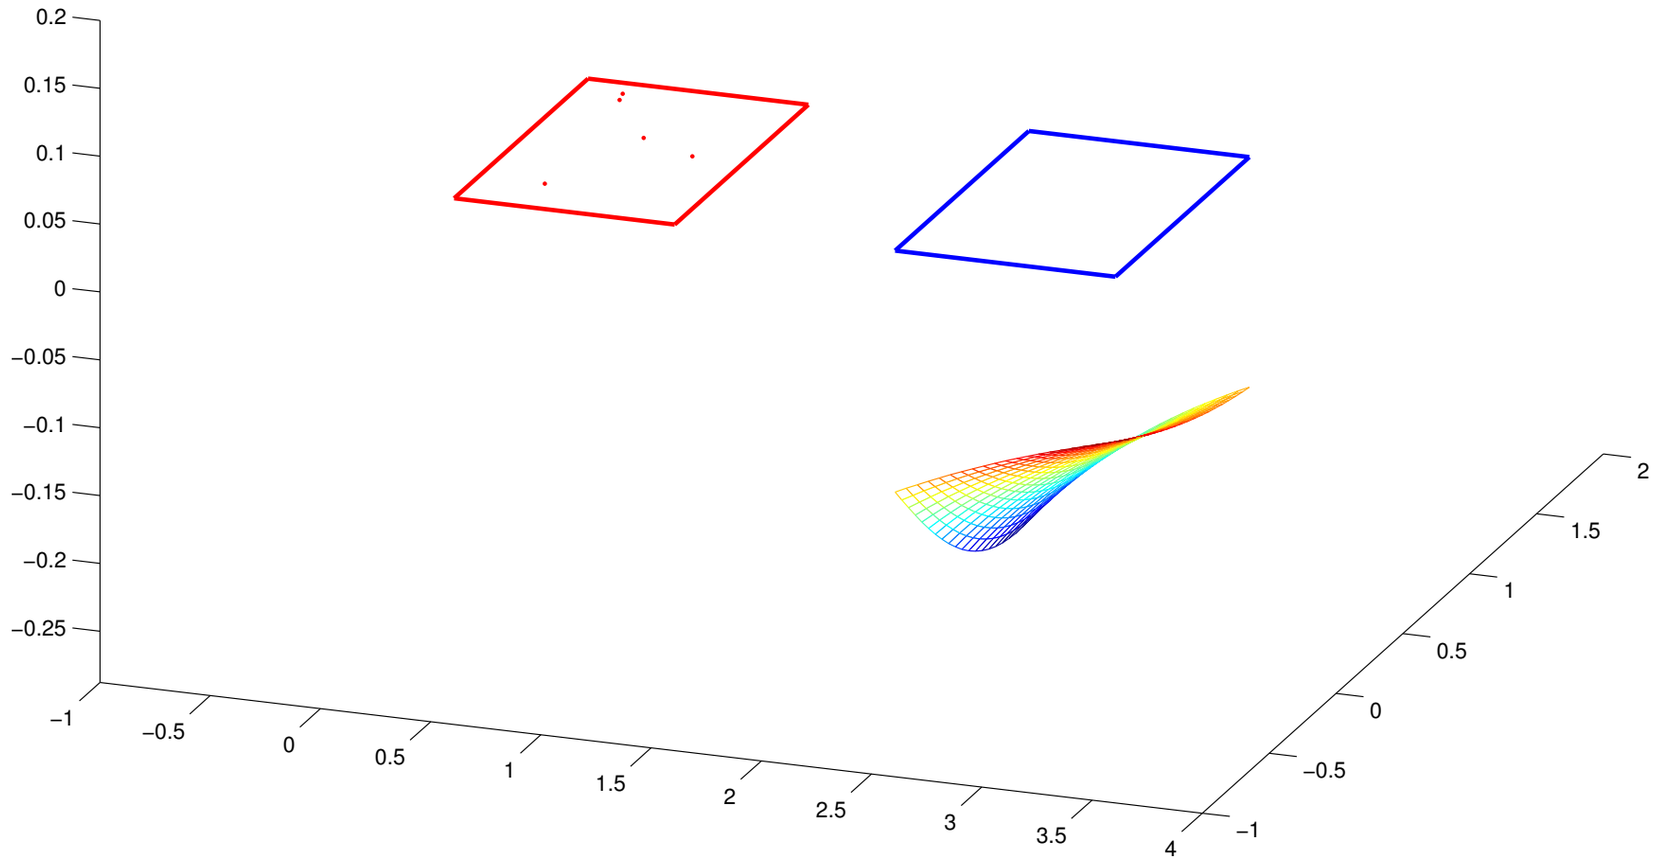

The field generated by the sources (truncated — the peaks go to infinity).

Example of solution of Laplace's equation $-\Delta u = g$

Suppose we are given point charges $\{q_j\}_{j=1}^{5}$ in a "source domain" Ω_{s}.

We are interested in the potential in a "target domain" Ω_{t}.

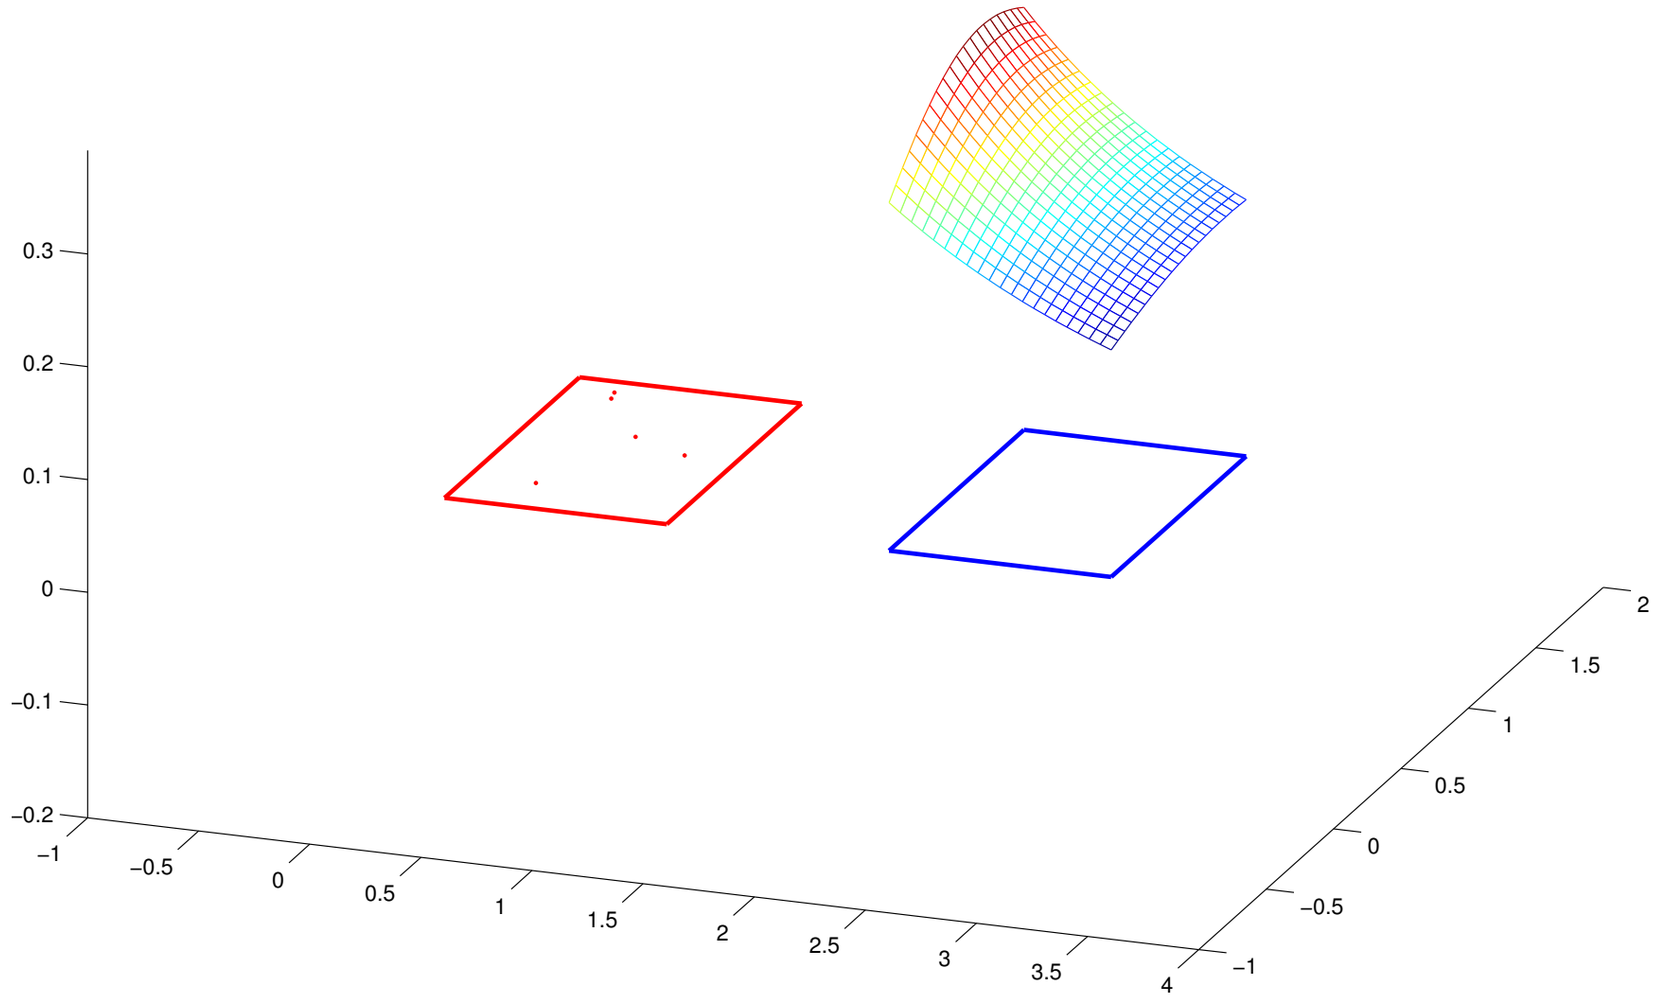

Field from a different (random) distribution of charges.

Example of solution of Laplace's equation $-\Delta u = g$

Suppose we are given point charges $\{q_j\}_{j=1}^{50}$ in a "source domain" $\Omega_{\rm S}$.

We are interested in the potential in a "target domain" $\Omega_{\rm t}$.

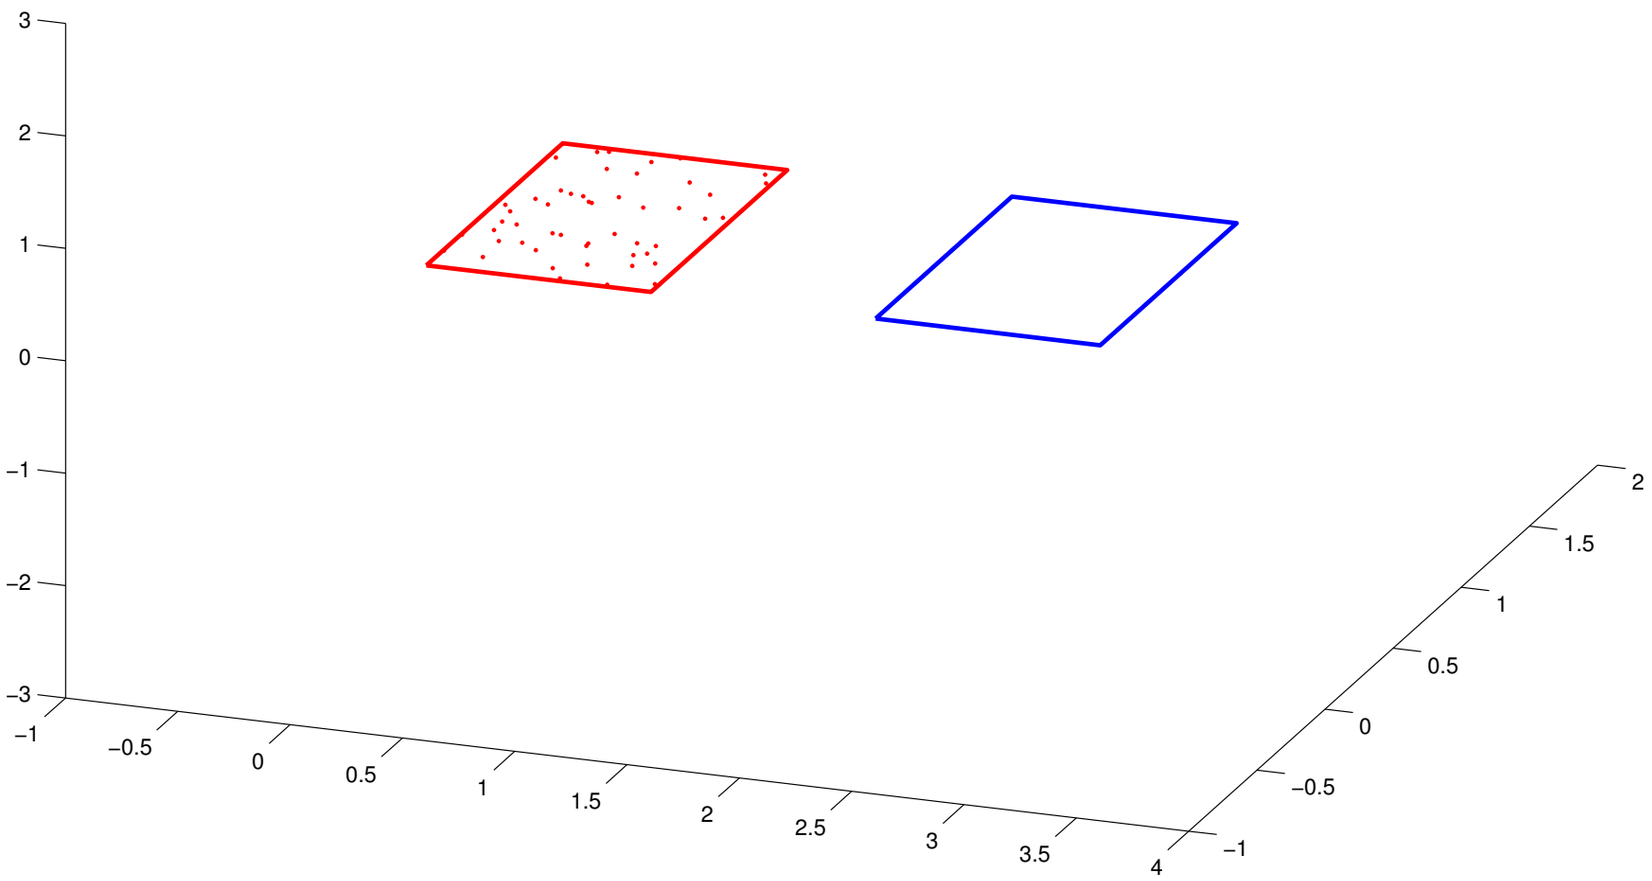

Now consider lots of charges in $\Omega_{\rm S}$.

Example of solution of Laplace's equation $-\Delta u = g$

Suppose we are given point charges $\{q_j\}_{j=1}^{50}$ in a "source domain" Ω_{s}.

We are interested in the potential in a "target domain" Ω_{t}.

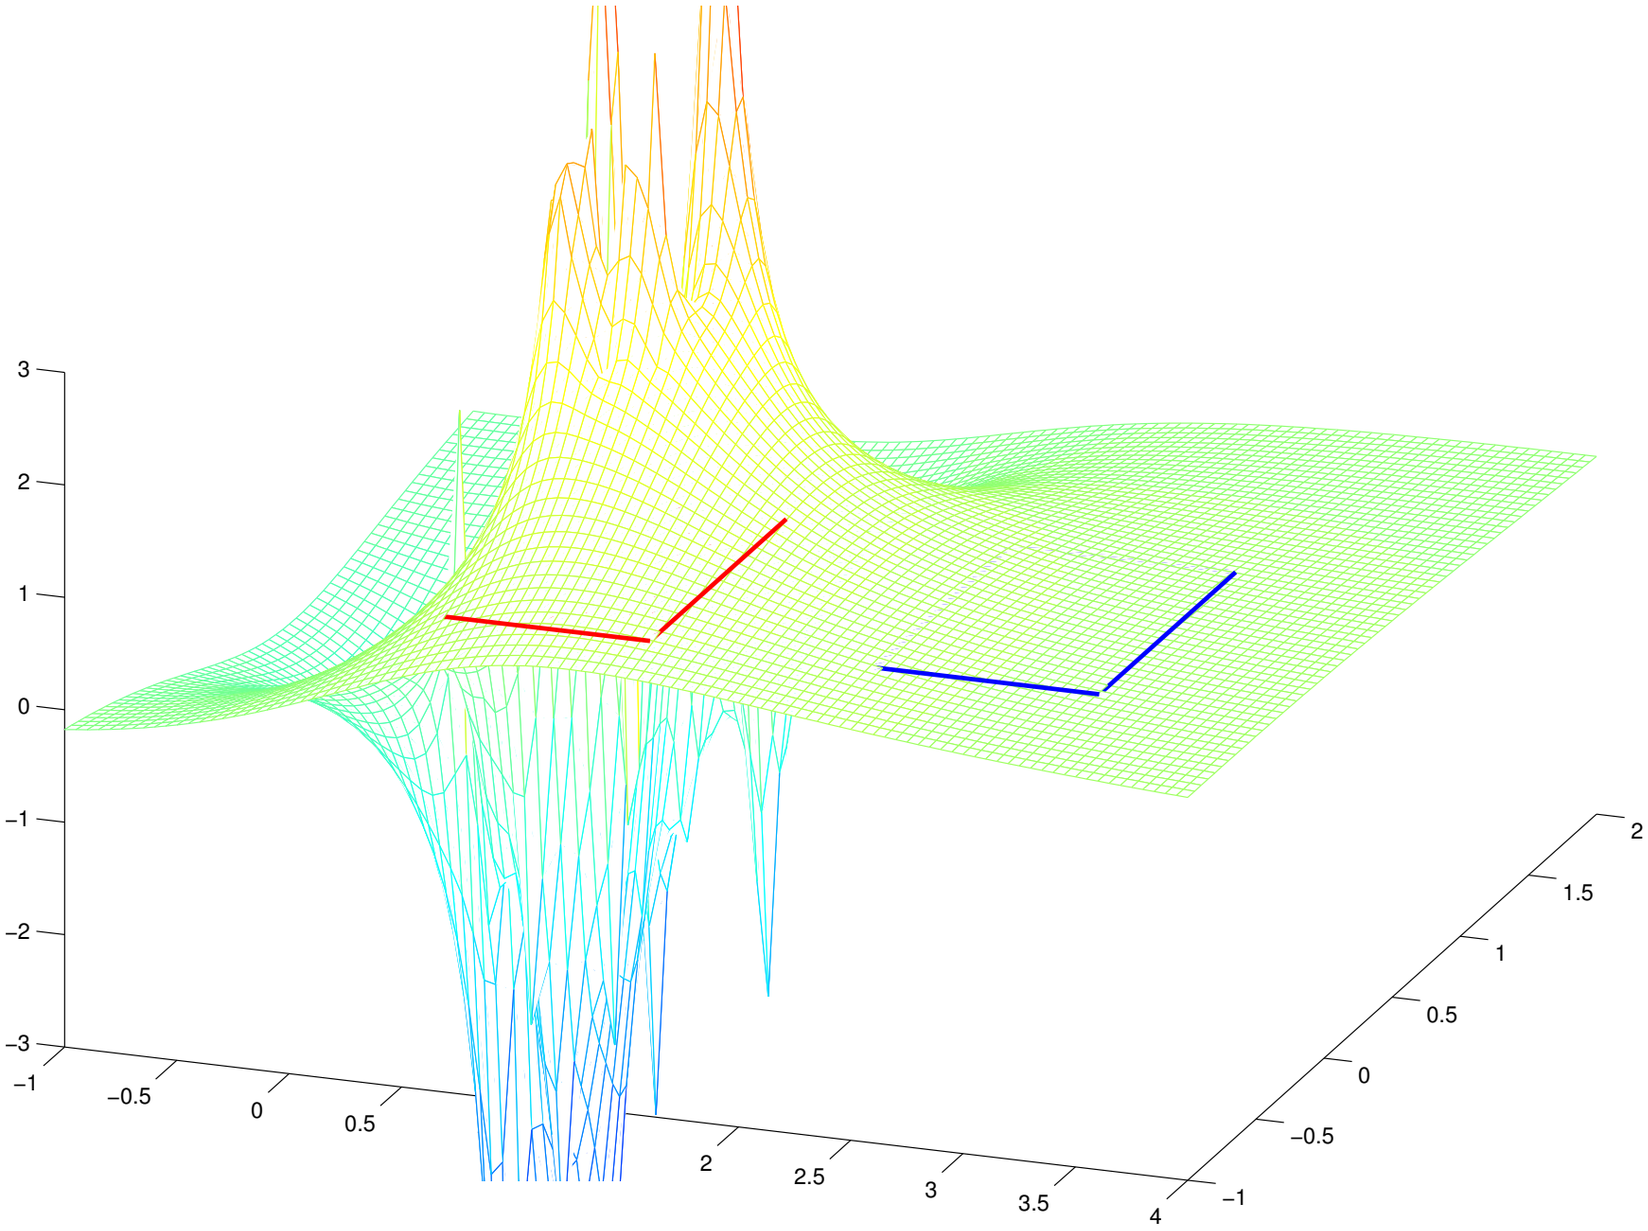

The field generated by the sources (truncated — the peaks go to infinity).

Example of solution of Laplace's equation $-\Delta u = g$

Suppose we are given point charges $\{q_j\}_{j=1}^{50}$ in a "source domain" Ω_{s}.

We are interested in the potential in a "target domain" Ω_{t}.

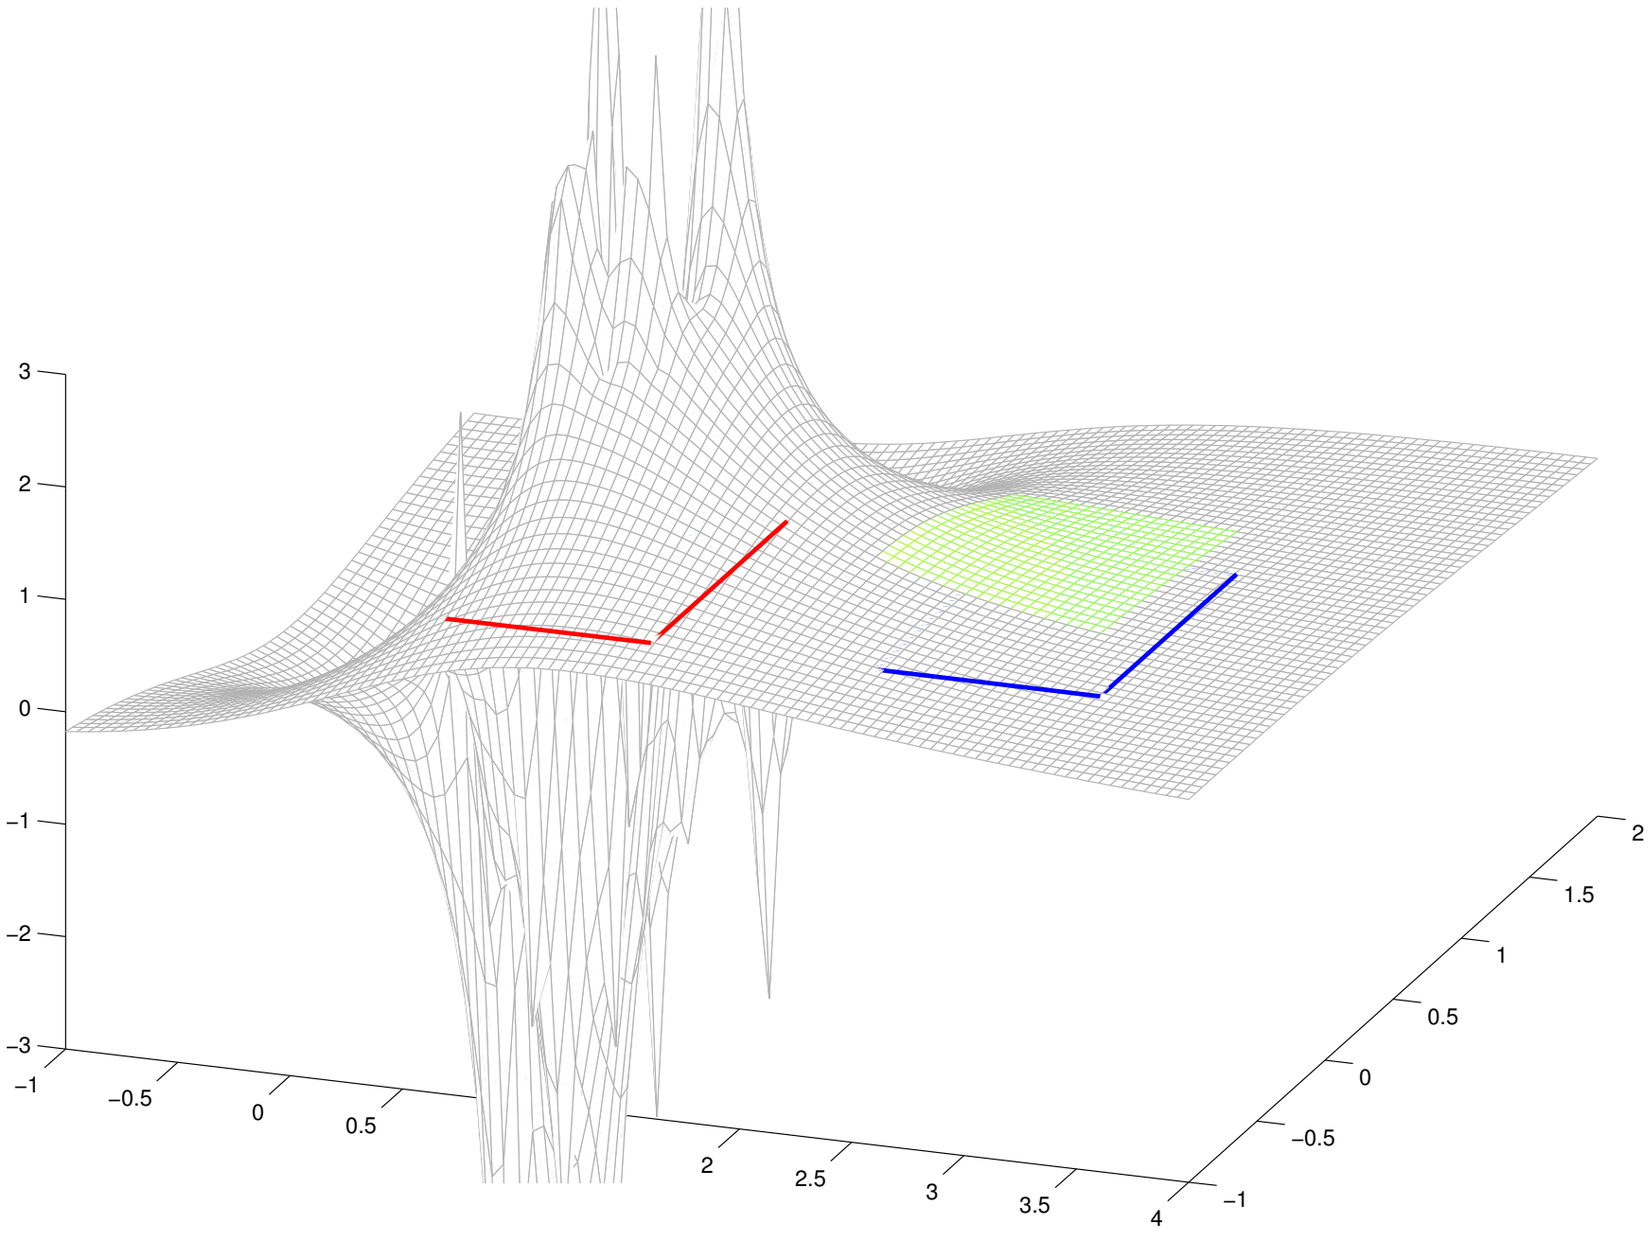

The field generated by the sources (truncated — the peaks go to infinity).

Example of solution of Laplace's equation $-\Delta u = g$

Suppose we are given point charges $\{q_j\}_{j=1}^{50}$ in a "source domain" $\Omega_{\rm s}$.

We are interested in the potential in a "target domain" $\Omega_{\rm t}$.

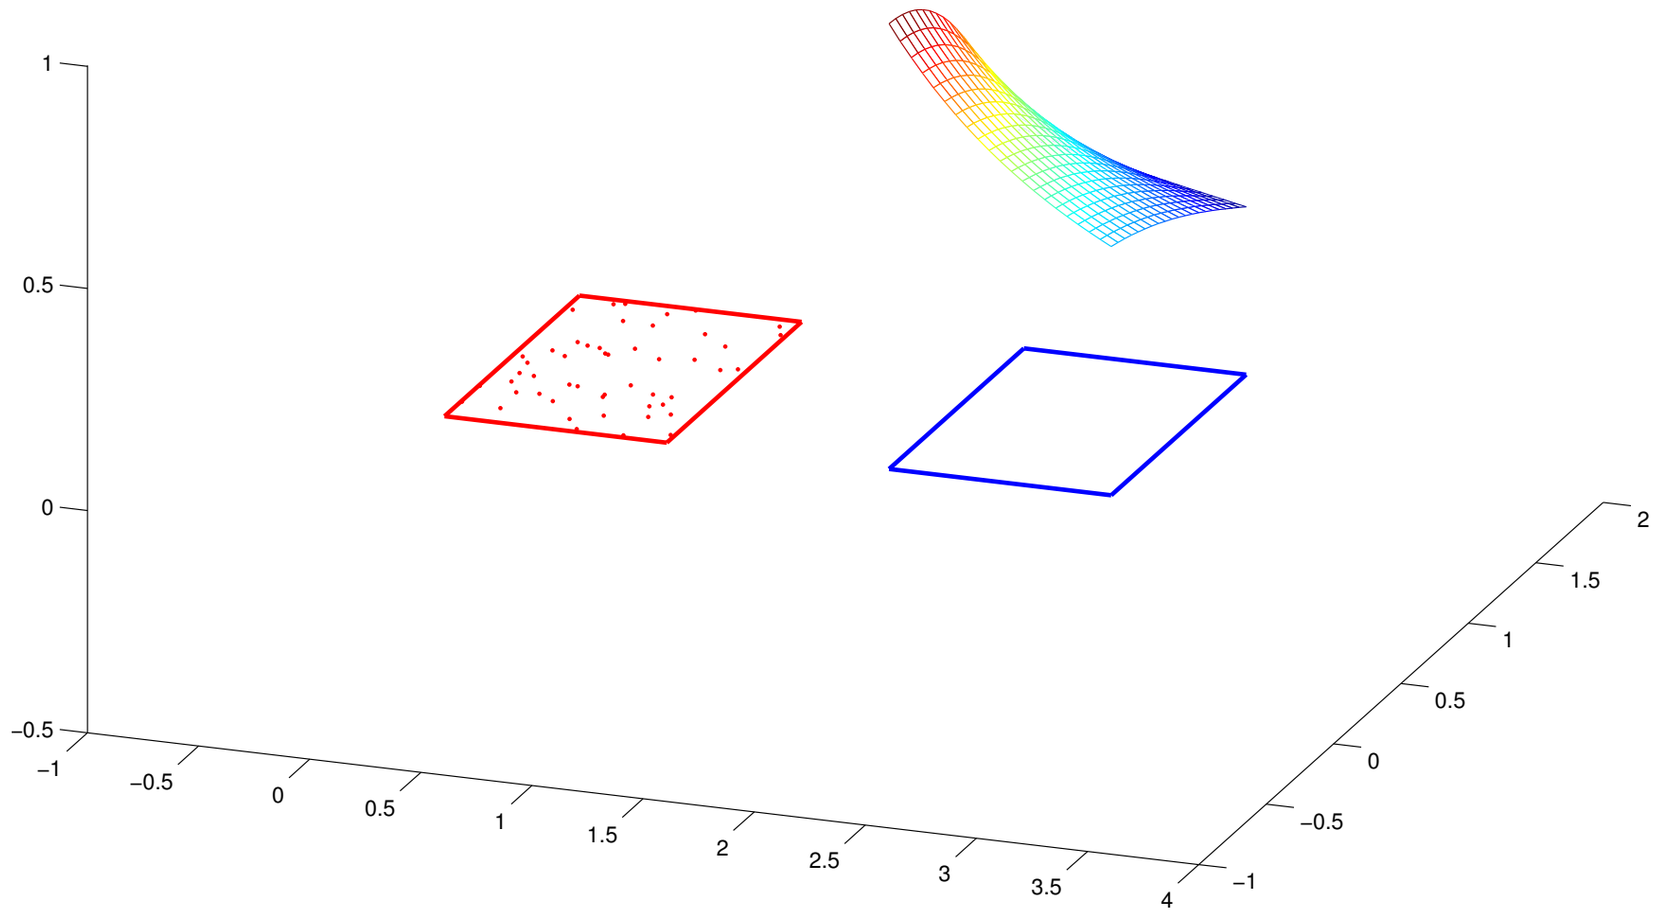

And a close-up.

Kernel expansions — proof that the "rank of interaction" is low

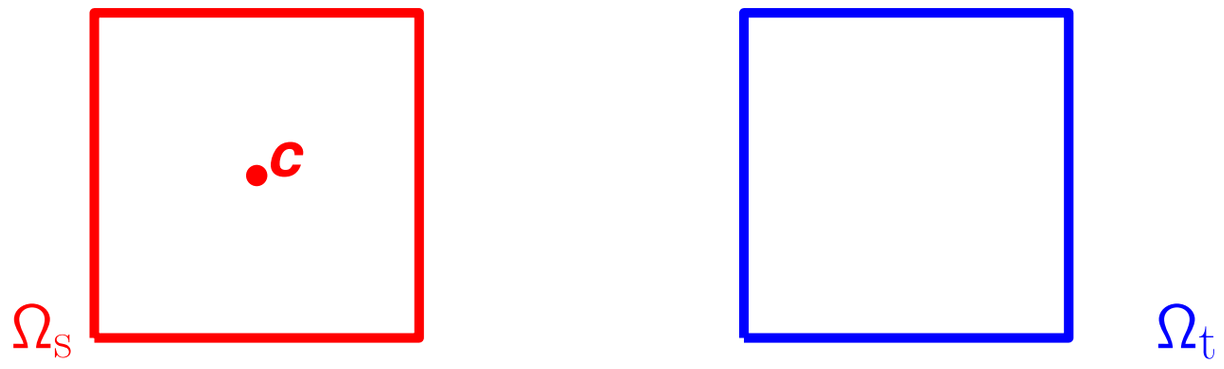

Let u denote the potential in Ω_{t} induced by charges g in Ω_{S} (without factor $-1/2\pi$):

$$u(\boldsymbol{x}) = \int_{\Omega_{\mathrm{S}}} \log|\boldsymbol{x} - \boldsymbol{y}|\, g(\boldsymbol{y})\, d\boldsymbol{y}, \qquad \boldsymbol{x} \in \Omega_{\mathrm{t}}.$$

For convenience of notation, we will switch to complex variables. Observe that for $\boldsymbol{x} = (x_1, x_2) = x_1 + i\, x_2 = r\, e^{it}$, we have

$$\log(\boldsymbol{x}) = \log|\boldsymbol{x}| + i \arg(\boldsymbol{x}) = \log r + i\, t.$$

We can therefore consider the problem of evaluating the complex valued integral

$$u(\boldsymbol{x}) = \int_{\Omega_{\mathrm{S}}} \log(\boldsymbol{x} - \boldsymbol{y})\, g(\boldsymbol{y})\, d\boldsymbol{y}, \qquad \boldsymbol{x} \in \Omega_{\mathrm{t}}.$$

When g is real, the "real-valued" u is simply the real part of the integral.

Let u denote the potential in Ω_{t} induced by sources g in Ω_{s}:

$$u(\boldsymbol{x}) = \int_{\Omega_{\mathrm{s}}} \log(\boldsymbol{x}-\boldsymbol{y})\, g(\boldsymbol{y})\, d\boldsymbol{y}, \qquad \boldsymbol{x} \in \Omega_{\mathrm{t}}. \tag{10}$$

We will use the McLaurin expansion

$$\log(1-z) = -\sum_{p=1}^{\infty} \frac{1}{p} z^p, \qquad \text{valid for } |z| < 1.$$

Let $\boldsymbol{c}$ denote the center of Ω_{s}. Then

$$\begin{aligned} \log(\boldsymbol{x}-\boldsymbol{y}) &= \log\big((\boldsymbol{x}-\boldsymbol{c}) - (\boldsymbol{y}-\boldsymbol{c})\big) \\ &= \log(\boldsymbol{x}-\boldsymbol{c}) + \log\left(1 - \frac{\boldsymbol{y}-\boldsymbol{c}}{\boldsymbol{x}-\boldsymbol{c}}\right) = \log(\boldsymbol{x}-\boldsymbol{c}) - \sum_{p=1}^{\infty} \frac{1}{p} \frac{(\boldsymbol{y}-\boldsymbol{c})^p}{(\boldsymbol{x}-\boldsymbol{c})^p}. \end{aligned} \tag{11}$$

Let r_{s} denote the radius of Ω_{s}. Combining (10) and (11), we get

$$u(\boldsymbol{x}) = \log(\boldsymbol{x}-\boldsymbol{c}) \int_{\Omega_{\mathrm{s}}} g(\boldsymbol{y})\, d\boldsymbol{y} + \sum_{p=1}^{\infty} \frac{r_{\mathrm{s}}^p}{-p(\boldsymbol{x}-\boldsymbol{c})^p} \int_{\Omega_{\mathrm{s}}} \left(\frac{\boldsymbol{y}-\boldsymbol{c}}{r_{\mathrm{s}}}\right)^p g(\boldsymbol{y})\, d\boldsymbol{y}.$$

Set $\hat{g}_0 = \int_{\Omega_{\mathrm{s}}} g(\boldsymbol{y})\, d\boldsymbol{y}$ and $\hat{g}_p = \int_{\Omega_{\mathrm{s}}} \left(\frac{\boldsymbol{y}-\boldsymbol{c}}{r_{\mathrm{s}}}\right)^p g(\boldsymbol{y})\, d\boldsymbol{y}$ and observe that $|\hat{g}_p| \le ||g||_{L^1(\Omega_{\mathrm{s}})}$.
It follows that we can write the potential u as a sum

$$u(\boldsymbol{x}) = \log(\boldsymbol{x}-\boldsymbol{c})\, \hat{g}_0 + \sum_{p=1}^{\infty} \frac{r_{\mathrm{s}}^p}{-p(\boldsymbol{x}-\boldsymbol{c})^p}\, \hat{g}_p.$$

The sum converges exponentially fast whenever $|\boldsymbol{x}-\boldsymbol{c}| > r_{\mathrm{s}}$.

We can write the series expansion without complex variables by taking the real part, and find that, with $\boldsymbol{x} - \boldsymbol{c} = r\, e^{it}$ and $\boldsymbol{y} - \boldsymbol{c} = r'\, e^{it'}$,

$$u(\boldsymbol{x}) = \log r \int_{\Omega_S} g(r', t')\, dA(r', t') +$$
$$\sum_{p=1}^{\infty} \left[\frac{-1}{p} \left(\frac{r_S}{r} \right)^p \cos(pt) \int_{\Omega_S} \left(\frac{(r')}{r_S} \right)^p \cos(pt')\, g(r', t')\, dA(r', t') + \right.$$
$$\left. \frac{-1}{p} \left(\frac{r_S}{r} \right)^p \sin(pt) \int_{\Omega_S} \left(\frac{(r')}{r_S} \right)^p \sin(pt')\, g(r', t')\, dA(r', t') \right].$$

Note that we expand the field in a rapidly converging sum of the simple functions

$$C_0(r, t) = \log(r)$$
$$C_p(r, t) = \frac{1}{r^p} \cos(pt)$$
$$S_p(r, t) = \frac{1}{r^p} \sin(pt).$$

Example — the Laplace equation with Dirichlet data on a circle.

Set

$$\Omega = \{\boldsymbol{x} \in \mathbb{R}^2 : \ |\boldsymbol{x}| \leq R\},$$

and consider the equation

$$\begin{cases} -\Delta u(\boldsymbol{x}) = 0, & \boldsymbol{x} \in \Omega, \\ u(\boldsymbol{x}) = h(\boldsymbol{x}), & \boldsymbol{x} \in \Gamma. \end{cases}$$

It is convenient to work with polar coordinates (r, t) so that $x_1 = r\cos t$ and $x_2 = r\sin t$. Then g is a periodic function of t. We write its Fourier expansion as

$$h(t) = \sum_{n=-\infty}^{\infty} h_n e^{int}$$

where

$$h_n = \frac{1}{2\pi}\int_{-\pi}^{\pi} e^{-ins}\, h(s)\, ds.$$

Now consider the function

$$u(r, t) = \sum_{n=-\infty}^{\infty} h_n \left(\frac{r}{R}\right)^{|n|} e^{int}.$$

Then u obviously satisfies the boundary condition, and it is easily verified that $-\Delta u = 0$. Due to uniqueness, we must have found the correct solution.

Recall that the solution to the Dirichlet problem can be written

$$u(r,t) = \sum_{n=-\infty}^{\infty} h_n \left(\frac{r}{R}\right)^{|n|} e^{int}.$$

where $h_n = \dfrac{1}{2\pi}\displaystyle\int_{-\pi}^{\pi} e^{-ins}\, h(s)\, ds$. By combining the two expressions, we find that

$$u(r,t) = \int_{-\pi}^{\pi} G(r, t-s)\, h(s)\, ds$$

where G is the *Green's function* of the problem

$$G(r,t-s) = \frac{1}{2\pi} \sum_{n=-\infty}^{\infty} \left(\frac{r}{R}\right)^{|n|} e^{in(t-s)}.$$

Observe smoothing: away from Γ, the solution is analytic.

Green's function for a circle

Let $\Omega = \{\boldsymbol{x} \in \mathbb{R}^2 : |\boldsymbol{x}| \leq R\}$ denote the unit circle, and consider the equation

$$\begin{cases} -\Delta u(\boldsymbol{x}) = 0, & \boldsymbol{x} \in \Omega, \\ u(\boldsymbol{x}) = h(\boldsymbol{x}), & \boldsymbol{x} \in \Gamma. \end{cases}$$

Recall that the solution can be written

$$u(r,t) = \int_{-\pi}^{\pi} G(r, t-s)\, h(s)\, ds$$

where G is the *Green's function* of the problem

$$G(r, t-s) = \frac{1}{2\pi} \sum_{n=-\infty}^{\infty} \left(\frac{r}{R}\right)^{|n|} e^{in(t-s)}.$$

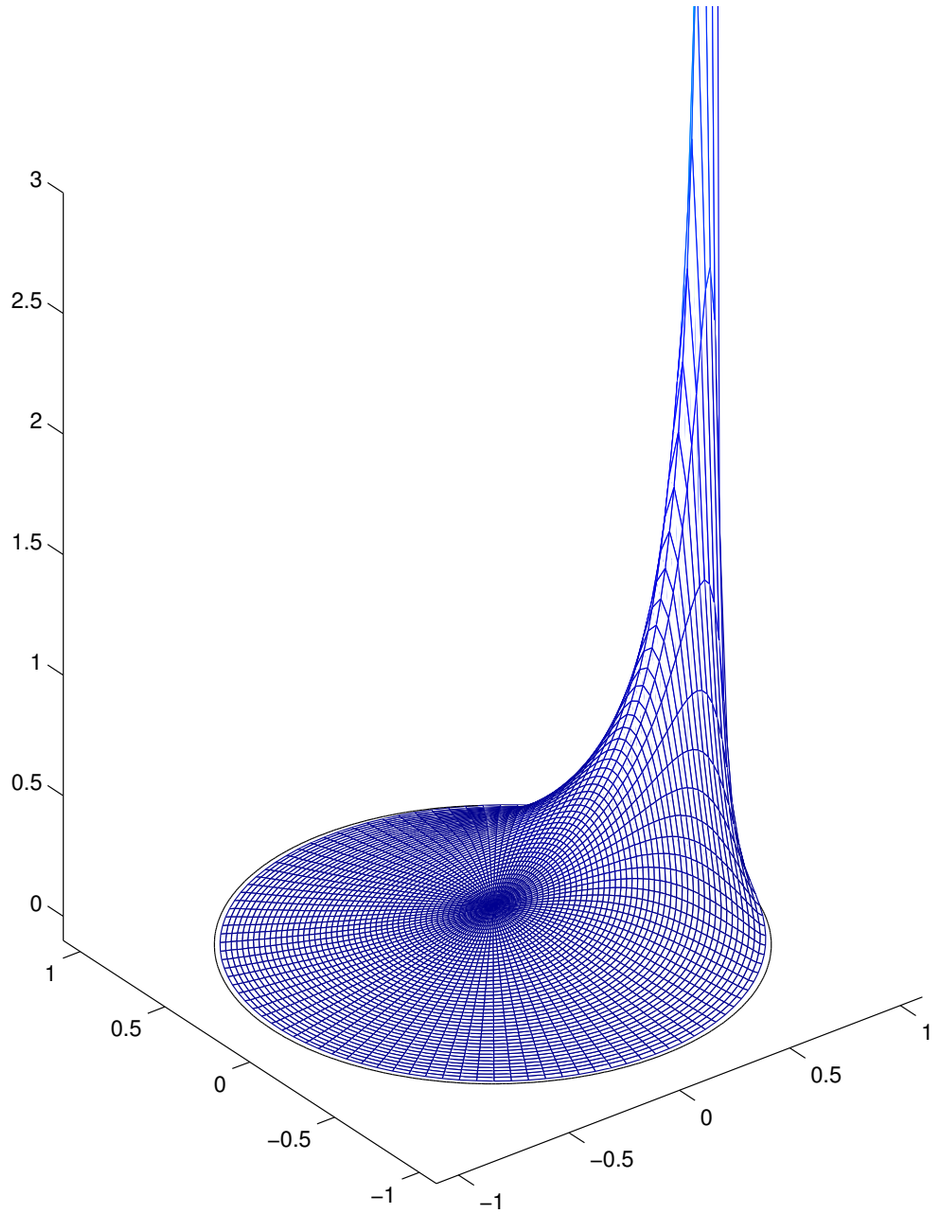

The Green's function.

Now consider a Dirichlet problem on a general domain:

$$\begin{cases} -\Delta u(\boldsymbol{x}) = 0, & \boldsymbol{x} \in \Omega, \\ u(\boldsymbol{x}) = h(\boldsymbol{x}), & \boldsymbol{x} \in \Gamma. \end{cases} \tag{12}$$

We know, *in principle*, that there exists a "Green's function" H such that

$$u(\boldsymbol{x}) = \int_\Gamma H(\boldsymbol{x}, \boldsymbol{y})\, h(\boldsymbol{y})\, ds(\boldsymbol{y}). \tag{13}$$

H is known analytically only for trivial domains (rectangles, spheres, half-planes, etc).

However, under very mild conditions, there exists a function q on Γ such that

$$u(\boldsymbol{x}) = \int_\Gamma \phi(\boldsymbol{x} - \boldsymbol{y})\, q(\boldsymbol{y})\, ds(\boldsymbol{y}). \tag{14}$$

We now use a known kernel function — the free space fundamental solution $\phi(\boldsymbol{x} - \boldsymbol{y}) = -\frac{1}{2\pi} \log |\boldsymbol{x} - \boldsymbol{y}|$. The techniques already described now apply, and show that the solution is smooth on any domain that is sufficiently removed from the boundary.

Note: We can obtain an equation for q by inserting the boundary condition $u|_\Gamma = h$ from (12) into (13),

$$g(\boldsymbol{x}) = \int_\Gamma \phi(\boldsymbol{x} - \boldsymbol{y})\, q(\boldsymbol{y})\, ds(\boldsymbol{y}), \qquad \boldsymbol{x} \in \Gamma.$$

This equation can be solved to obtain q. More on this later ...

Summary:

- Solution operators of elliptic PDEs are exceptionally benign mathematical objects:
  - Behave like *integration operators.*
  - Typically bounded.
  - Typically very stable in many ways.
  - Extreme information loss over distance.

  One could say that the objective of the "direct solver" research effort is to numerically build approximations to the analytic solution operators.

- Advantages of building solution operators:
  - Very fast when the same equation is to be solved multiple times.
  - Capable of solving problems for which iterative solvers may not converge.
  - They *parallelize very well.*
  - They enable operator algebra.

- The principal difficulty we will need to wrestle with is that the solution operators are *global* (as opposed to differential operators which are local).

  We can attain algorithms with very good scaling (often $O(N)$) despite the fact that the matrices are dense by exploiting that the dense matrices are *rank structured.*

  However, direct solvers typically do require more memory than iterative solvers.