Boundary integral equation methods for elliptic problems - part 2

Rikard Ojala

NA, KTH

April 30, 2013

Last lecture...

- We saw that Laplace's equation in Ω could be reformulated as an integral equation on the boundary $\partial\Omega$ of the domain in several ways.
- We derived Green's representation formula and noted that it could be used to solve Laplace's equation.
- Opting for simplicity we instead picked the double layer representation

$$U(x_0) = \frac{1}{2\pi} \int_{\partial\Omega} \mu(x) \frac{\partial}{\partial n} \log(|x_0 - x|) \, \mathrm{d}l(x)$$

  of the solution in Ω
- Using the jump relations we got the following integral equation for the density μ on $\partial\Omega$:

$$\frac{1}{2}\mu(x_0) + \frac{1}{2\pi} \int_{\partial\Omega} \mu(x) \frac{\partial}{\partial n} \log(|x_0 - x|) \, \mathrm{d}l(x) = f(x_0)$$

Last lecture...

- I decided that the equivalent complex representation was easier to handle (which it is). The double layer representation in complex variables is

$$U(z_0) = \frac{1}{2\pi} \int_{\partial\Omega} \mu(z) \mathrm{Im} \left\{ \frac{\mathrm{d}z}{z - z_0} \right\}.$$

and using the jump relations we get the integral equation

$$\frac{1}{2}\mu(z_0) + \frac{1}{2\pi} \int_{\partial\Omega} \mu(z) \mathrm{Im} \left\{ \frac{\mathrm{d}z}{z - z_0} \right\} = f(z_0),$$

for μ on the boundary.

Compact integral operators

- We can write

$$\frac{1}{2}\mu(z_0) + \frac{1}{2\pi}\int_{\partial\Omega}\mu(z)\mathrm{Im}\left\{\frac{\mathrm{d}z}{z-z_0}\right\} = f(z_0),$$

  symbolically as

$$\left(\frac{1}{2}\mathcal{I} + \mathcal{K}\right)\mu = f.$$

- Here, we call $\mathcal{I}$ the identity operator and $\mathcal{K}$ is an integral operator with kernel

$$K(z, z_0) = \frac{1}{2\pi}\mathrm{Im}\left\{\frac{\mathrm{d}z}{z-z_0}\right\}.$$

- Integral operators are in general infinite dimensional, but there is a special kind of integral operator that is "almost finite dimensional".
- We call such operators compact integral operators.

Compact integral operators

- We will need some terminology.
- A linear space is a collection of objects on which the operations addition and multiplication by scalar are defined, obeying the common rules such as commutativity, etc.
- A normed linear space $\mathcal{S}$, also has a norm $\|\cdot\|$ defined on it for which it should hold for $x, y \in \mathcal{S}$ that
  - $\|x\| = 0$ then $x = 0$,
  - $\|ax\| = |a|\|x\|$ for scalars a,
  - $\|x + y\| \leq \|x\| + \|y\|$
- There are many examples of normed linear spaces. For example $\mathbb{R}^n$ equipped with various norms.
- We will discuss spaces of functions here. Popular examples of such are
  - $C([a, b])$, the space of continuous functions over $[a, b]$ with norm $\|f\| = \sup_{x \in [a,b]} f(x)$,
  - $L^2([a, b])$, the space of square integrable functions over $[a, b]$ with norm $\|f\| = \sqrt{\int_a^b |f(x)|^2 \mathrm{d}x}$.

Compact integral operators

- An open ball of radius r centered at x_0 consists of all x such that $\|x - x_0\| < r$.
- A sequence is an ordered list $x_1, x_2, ...$ of elements and a Cauchy sequence is a sequence where elements get closer and closer. That is for any $\varepsilon > 0$, $\|x_n - x_k\| < \varepsilon$ for some N such that $n > N, k > N$.
- A sequence converges to x if $\|x - x_n\| \to 0$ as $n \to \infty$.
- A space for which all Cauchy sequences converge is called a Banach space.
- The closure of a set $\mathcal{X}$ is $\mathcal{X}$ but with the limit points added. For example, the closure of the interval (a, b) on the real line is $[a, b]$.

Compact integral operators

- We first need the concept of a compact set.

Definition

A set $\mathcal{X}$ in a Banach space is called compact if and only if every sequence $x_1, x_2, x_3, ...$ of elements in $\mathcal{X}$ contains a converging subsequence.

If a set $\mathcal{X}$ can be covered by a finite number of balls of radius ε for each $\varepsilon > 0$ then the closure of $\mathcal{X}$ is compact.

- These definitions are equivalent and they both hint that compact sets are in some sense small.
- Heine-Borel's theorem says that any closed and bounded subset of $\mathbb{R}^n$ is compact.
- So $[a, b]$ is compact for any finite a and b but the whole real line $\mathbb{R}$ is not.

Compact integral operators

Definition

An integral operator $\mathcal{K}$ taking inputs from $\mathcal{X}$ and producing outputs in $\mathcal{Y}$ is called a compact integral operator if and only if every bounded subset of $\mathcal{X}$ is mapped into a subset of $\mathcal{Y}$ whose closure is compact.

- We could say that a compact integral operator takes inputs from a potentially large set and produces outputs in a smaller set.
- So what are the conditions for an integral operator to be compact?
- Well, it depends on what space we are dealing with. For our purposes we can limit ourselves to the space of square integrable functions on the interval $[a, b]$. Then

Theorem

An integral operator from $L^2([a,b])$ to $L^2([a,b])$ is compact if its kernel $K(x, x_0)$ obeys

$$\int_a^b \int_a^b |K(x, x_0)|^2 \mathrm{d}x\mathrm{d}x_0 < \infty.$$

Compact integral operators

- Let us check how our integral operator

$$\mathcal{K}\mu(z_0) = \int_{\partial\Omega} \mu(z)\mathrm{Im}\left\{\frac{\mathrm{d}z}{z - z_0}\right\}$$

  does.
- Introduce a parameterization $z(t), t \in [0, 2\pi]$ of the boundary. We then get

$$\mathcal{K}\mu(z_0) = \int_0^{2\pi} \mu(z)\mathrm{Im}\left\{\frac{z'(t)}{z(t) - z_0}\right\}\mathrm{d}t,$$

  and using the same parameterization for z_0 we write the kernel as

$$K(t, s) = \mathrm{Im}\left\{\frac{z'(t)}{z(t) - z(s)}\right\}.$$

- Is $K(t, s)$ square integrable? Yes, for smooth boundaries it is, even though the denominator looks problematic for $t = s$. It is in fact continuous, which you will show in the homework.

Compact integral operators

- Great, so our integral operator is compact. What does that buy us?
- First off, we have the following theorem

Theorem

If $\mathcal{K}$ is a compact operator then its spectrum consists of either a finite number of eigenvalues, or a denumerable set of eigenvalues accumulating only at the origin.

- This basically says that either $\mathcal{K}$ has a finite number of eigenvalues, in which case $\mathcal{K}$ is finite dimensional, or
- $\mathcal{K}$ has infinitely many eigenvalues, but they tend toward zero. $\mathcal{K}$ is "almost finite dimensional".
- In either case, making $\mathcal{K}$ finite dimensional by discretizing it should not remove too much information.
- These spectral properties will also be very advantageous later when we discuss iterative solvers for solving the discretized linear systems of equations.

Compact integral operators

- Knowing that $\mathcal{K}$ is compact gives us the opportunity to use another theorem

Theorem

Fredholm's alternative. *If the integral operator $\mathcal{K}$ is compact, then either the homogenous integral equation $(\mathcal{I}+\mathcal{K})\mu = 0$ has a non-trivial solution(μ not identically zero), or $(\mathcal{I}+\mathcal{K})\mu = f$ has a unique solution for every f.*

- Compare this to the finite dimensional case : either the linear system of equations $\mathbf{Ax} = \mathbf{0}$ has a non-trivial solution, or $\mathbf{Ax} = \mathbf{f}$ has a unique solution for every $\mathbf{f}$
- One can show that $(\mathcal{I}+\mathcal{K})\mu = 0$ for our integral operator has only the solution $\mu = 0$, thereby showing by Fredholm's alternative that $(\mathcal{I}+\mathcal{K})\mu = f$ is uniquely solvable for any f. This requires quite a lot of analytic function theory so we will skip that here.

Recap and what now?

- To recap, compactness of our integral operator gave us
  - a way to show existence and uniqueness of solutions to the integral equation,
  - that the integral operator is almost finite dimensional, it should discretize nicely, and
  - once discretized, the linear system of equations can be solved rapidly using iterative solvers (but more on that later).
- So, let us discretize the integral equation. We will need to evaluate integrals, and we will do this using the good old trapezoidal rule.
- As it turns out, the trapezoidal rule is astoundingly good for evaluating the integrals we will encounter.

Trapezoidal rule error

- The standard error estimate for the trapezoidal rule says that the error decays as h^2.
- This is a massive overestimation of the error when
  1. the integrand is smooth, many times continuously differentiable,
  2. the integrand is periodic.
- For periodic functions the trapezoidal rule becomes the midpoint rule, but this does not explain the higher order. Why would the error be much lower?
- A first "hand waving" explanation might be that because of periodicity, the over and underestimations cancel over one period.

Trapezoidal rule error

- More precisely we have the following :

Theorem

Let $m \geq 0$, $N \geq 1$, and define $h = 2\pi/N, x_j = jh$ for $j = 0, 1, ..., N-1$. Further assume that $f(x)$ is $2m+2$ times continuously differentiable on $[0, 2\pi]$ for some $m \geq 0$. Then, for the error in the trapezoidal rule

$$\int_0^{2\pi} f(x)\mathrm{d}x - T_N(f) = -\sum_{k=1}^{m} \frac{B_{2k}}{(2k)!} h^{2k} \left(f^{(2k-1)}(2\pi) - f^{(2k-1)}(0) \right)$$
$$-2\pi h^{2m+2} \frac{B_{2m+2}}{(2m+2)!} f^{(2m+2)}(\xi)$$

for some ξ in $[0, 2\pi]$. The B_k are the Bernoulli numbers, and $T_N(f)$ denotes the N-point trapezoidal rule applied to f.

- For smooth (m large), periodic ($f^{(n)}(0) = f^{(n)}(2\pi)$) functions we get very high order, "spectral accuracy".

Nyström discretization

- Given a periodic parameterization $z(t), t \in [0, 2\pi]$ we would like to discretize our integral equation so that we can solve it.
- Introducing the parameterization in our equation we get

$$\frac{1}{2}\mu(z(s)) + \frac{1}{2\pi}\int_{\partial\Omega} \mu(z(t))\mathrm{Im}\left\{\frac{z'(t)}{z(t)-z(s)}\right\}\mathrm{d}t = f(z(s)).$$

  for $s \in [0, 2\pi]$.
- Let us pick N equidistant points $t_j = 2\pi j/N, j = 0, ..., N-1$. The distance between them is $h = 2\pi/N$. Evaluating the integral with the trapezoidal rule we get

$$\frac{1}{2}\mu(z(s)) + \frac{h}{2\pi}\sum_{j=0}^{N-1} \mu(z(t_j))\mathrm{Im}\left\{\frac{z'(t_j)}{z(t_j)-z(s)}\right\} = f(z(s)).$$

- We are halfway there. We can now compute this for every $s \in [0, 2\pi]$, but there are an infinite number of values of s. We need to pick a finite number of them.

Nyström discretization

- The idea in Nyström discretization is to pick the quadrature nodes as the points at which we require the integral equation to hold. We get

$$\frac{1}{2}\mu(z(t_i)) + \frac{h}{2\pi}\sum_{j=0}^{N-1}\mu(z(t_j))\mathrm{Im}\left\{\frac{z'(t_j)}{z(t_j)-z(t_i)}\right\} = f(z(t_i)),$$

  for $i = 0, 1, ..., N-1$.
- If we write $z_j = z(t_j), z'_j = z'(t_j), \mu_j = \mu(z(t_j)), f_j = f(z(t_j))$ we get

$$\frac{1}{2}\mu_i + \frac{h}{2\pi}\sum_{j=0}^{N-1}\mu_j\mathrm{Im}\left\{\frac{z'_j}{z_j-z_i}\right\} = f_i.$$

- This is an $N \times N$ equation system with μ_j as unknowns. We can write it as

$$\left(\frac{1}{2}\mathbf{I} + \mathbf{K}\right)\boldsymbol{\mu} = \mathbf{f}$$

Nyström discretization

- Once we have solved

$$\left(\frac{1}{2}\mathbf{I} + \mathbf{K}\right)\boldsymbol{\mu} = \mathbf{f}$$

  for $\boldsymbol{\mu}$ we know the value of the density $\mu(z)$ at the quadrature nodes.

- We may then directly use

$$U(z_0) = \frac{1}{2\pi}\int_{\partial\Omega} \mu(z)\mathrm{Im}\left\{\frac{\mathrm{d}z}{z - z_0}\right\},$$

  discretized in the same way to compute the solution $U(z_0)$ at any point z_0 in the domain.

- This is a natural and simple idea, and there are other upshots as well.

Nyström discretization

- For the error in Nyström discretization we have

Theorem

If $\mathcal{K}_N$ is $\mathcal{K}$ but with its integral evaluated numerically by the N point trapezoidal rule, and x and x_N obeys $(\mathcal{I}+\mathcal{K})x = f$ and $(\mathcal{I}+\mathcal{K}_N)x_N = f$ then

$$\|x - x_N\| \leq C \max_s \left| \int K(t,s)x(t)\mathrm{d}t - h\sum_{j=0}^{N-1} K(t_j,s)x(t_j) \right|$$

- It turns out that Nyström discretization preserves the order of the underlying quadrature rule.
- This is very good for us, since the trapezoidal rule has very high order.

Solving the system

- The matrix $\mathbf{K}$ with elements

$$K_{ij} = \frac{h}{2\pi} \mathrm{Im} \left\{ \frac{z_j'}{z_j - z_i} \right\}$$

  is not sparse.

- This means that solving

$$\left(\frac{1}{2}\mathbf{I} + \mathbf{K} \right) \boldsymbol{\mu} = \mathbf{f}$$

  using Gaussian elimination costs $O(N^3)$ operations for N discretization points. This is too expensive for larger problems.

- We need iterative solvers, but our choice is limited by the fact that the matrix $\frac{1}{2}\mathbf{I} + \mathbf{K}$ is neither symmetric nor positive definite. These are properties often exploited by iterative solvers.

GMRES

- Let's say we wish to solve $\mathbf{Ax} = \mathbf{b}$ (with $\mathbf{A}$ invertible).
- Iterative solvers in general work by computing the matrix-vector products $\mathbf{Ay_n}$ for some $\mathbf{y_n}$ and using the information to construct an approximate solution.
- Each matrix-vector multiplication costs $O(N^2)$ operations for an $N \times N$ matrix, so if we need much less than N such multiplications we gain something.
- The iterative solver that best suits our problem is GMRES (Generalized Minimum RESidual).
- GMRES belongs to the class of Krylov subspace methods, along with for example the conjugate gradient method.
- Unlike the conjugate gradient method it works on general matrices, but as we will see, this is not the only reason for choosing GMRES.

GMRES

- When solving the system $\mathbf{Ax} = \mathbf{b}$, the main idea of GMRES is to look for approximate solutions in the successive Krylov subspaces

$$\mathcal{S}_k = \text{span}\left\{\mathbf{b}, \mathbf{Ab}, \mathbf{A}^2\mathbf{b}, ..., \mathbf{A}^{k-1}\mathbf{b}\right\}.$$

- For an invertible $N \times N$ matrix, the solution $\mathbf{x}$ can be written as a linear combination of the columns in $\mathcal{S}_N$. But we hope to be able to find an accurate solution for k much less than N.
- At step k, we look for the vector x_k in $\mathcal{S}_k$ that minimizes the euclidean norm of the residual

$$\mathbf{r}_k = \mathbf{b} - \mathbf{Ax}_k.$$

- There are many technical details involved here, but we are mainly interested in the convergence properties of GMRES.

GMRES

- The convergence theory for Krylov subspace methods is complicated, but for GMRES we can at least say that if the spectrum of the matrix $\mathbf{A}$ is clustered away from the origin, then GMRES converges quickly.
- We can roughly say that the convergence depends on the size of the cluster of eigenvalues.
- Remember that for compact operators the spectrum tends to the origin.
- But we have $\frac{1}{2}\mathcal{I} + \mathcal{K}$ with $\mathcal{K}$ compact so the spectrum tends to $\frac{1}{2}$. The same holds for the discrete version $\frac{1}{2}\mathbf{I} + \mathbf{K}$.
- Because of this, as we increase the size of the system the additional eigenvalues will be very close to $\frac{1}{2}$ and will not increase the size of the cluster of eigenvalues. GMRES will converge as rapidly as before.
- We say that the number of GMRES steps is bounded as the system size grows for these kinds of matrices.

GMRES

- Each matrix-vector multiply costs $O(N^2)$ operations.
- Since the number of iterations is bounded for a given problem, we have decreased the work required by an order of magnitude, from $O(N^3)$ to $O(N^2)$. Not bad.
- However, we need to set up the matrix $\mathbf{K}$ which requires N^2 memory storage. For complicated boundaries, the number of discretization points required can easily exceed 10^5, which translates into 80 Gb of memory.
- We can actually do better, both in terms of operation count and memory. The structure of the matrices arising from boundary integral equation methods is a special one.
- The solution is the fast multipole method.

The fast multipole method

- The elements of our matrix are

$$K_{ij} = \frac{h}{2\pi} \mathrm{Im} \left\{ \frac{z'_j}{z_j - z_i} \right\}.$$

- Ignoring the imaginary part for the moment and setting $q_j = \frac{h}{2\pi} z'_j$ we have

$$K_{ij} = \frac{q_j}{z_j - z_i}.$$

- So multiplying the matrix $\mathbf{K}$ with a real vector $\mathbf{x}$ amounts to evaluating

$$\sum_{j=0}^{N-1} \frac{q_j x_j}{z_j - z_i}$$

  for all i and then taking the imaginary part.
- This is basically an n-body problem.

The fast multipole method

- Basic idea - gravity
- Consider two clusters of $n/2$ particles each.
- Classic evaluation costs $\sim n^2$.
- Instead compute an equivalent center of mass for each cluster.
- Cost drops to $\sim n^2/2$.
- By subdividing the clusters this could drop further.

The fast multipole method

- Some details need to be adressed:
  - What is the error?
  - Higher order "centers of mass"?
  - How to turn this into an efficient algorithm?
- The Fast Multipole Method (FMM) handles this.
  - Rigorous error estimates.
  - Uses multipole expansions for higher order approximations.
  - Uses a tree structure to handle interactions on several scales.

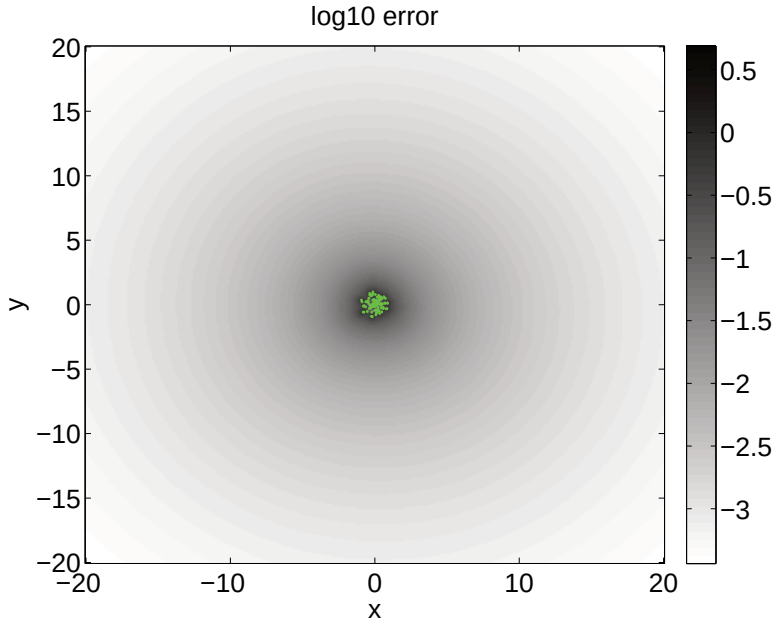

The fast multipole method

- Our goal is to compute

$$\phi(z) = \sum_{j=0}^{N-1} \frac{q_j}{z - z_j},$$

  for many points z.
- We would like to separate the sources z_j from the target z.
- To this end we use the identity

$$\frac{1}{z - z_j} = \frac{1}{z} \frac{1}{1 - \frac{z_j}{z}} = \frac{1}{z} \sum_{k=0}^{\infty} \left(\frac{z_j}{z} \right)^k,$$

  which is valid for $|z| > |z_j|$.
- Putting these two together gives us the *multipole expansion*.

The fast multipole method

The multipole expansion.
Suppose that N sources with strenghts q_k are located at points z_k, with $|z_k| < r$. Then for any z with $|z| > r$ we have

$$\phi(z) = \sum_{k=0}^{\infty} \frac{a_k}{z^{k+1}},$$

where

$$a_k = \sum_{j=0}^{N-1} q_j z_j^k.$$

For a truncated sum with only p terms, we have the error bound

$$\left| \phi(z) - \sum_{k=0}^{p} \frac{a_k}{z^k} \right| \leq C \left| \frac{r}{z} \right|^p$$

The fast multipole method

- So how do we turn this into an algorithm?
- Divide the computational domain into successively finer square grids.
- Do this until the number of particles in each box is small enough for multipole evaluation not to pay off.
- Evaluate multipole expansions whenever the source box is far enough away.

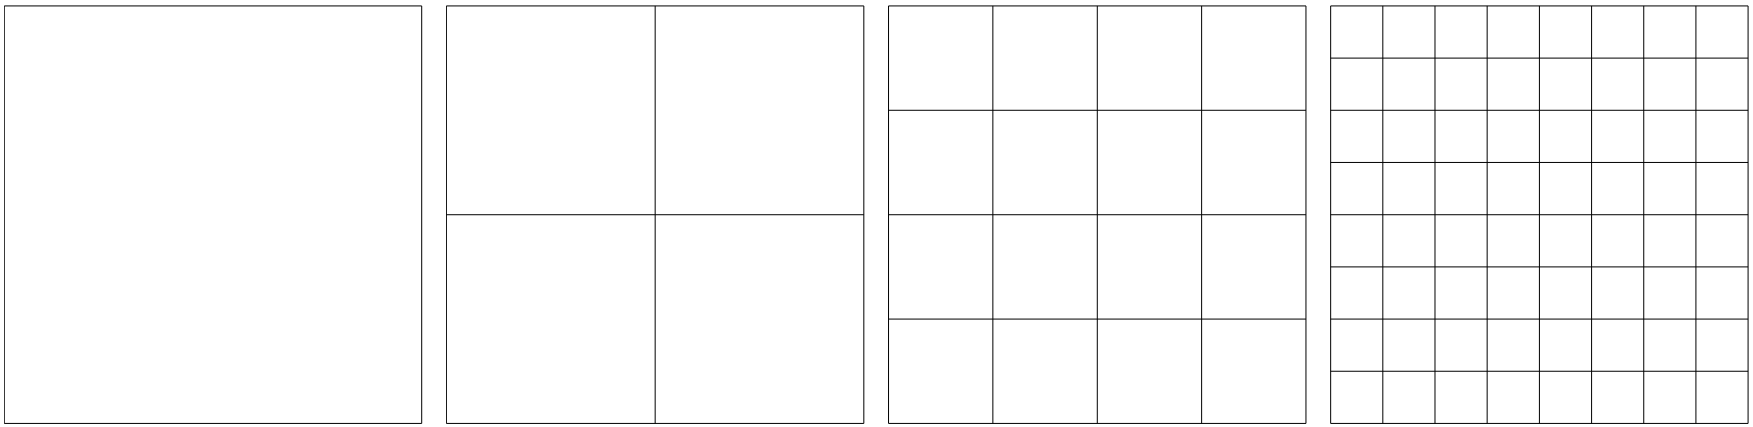

The fast multipole method

- We have to take care not to evaluate interactions more than once.
- Only treat particles in the *interaction list* of each box.
- At the finest level we use direct summation.

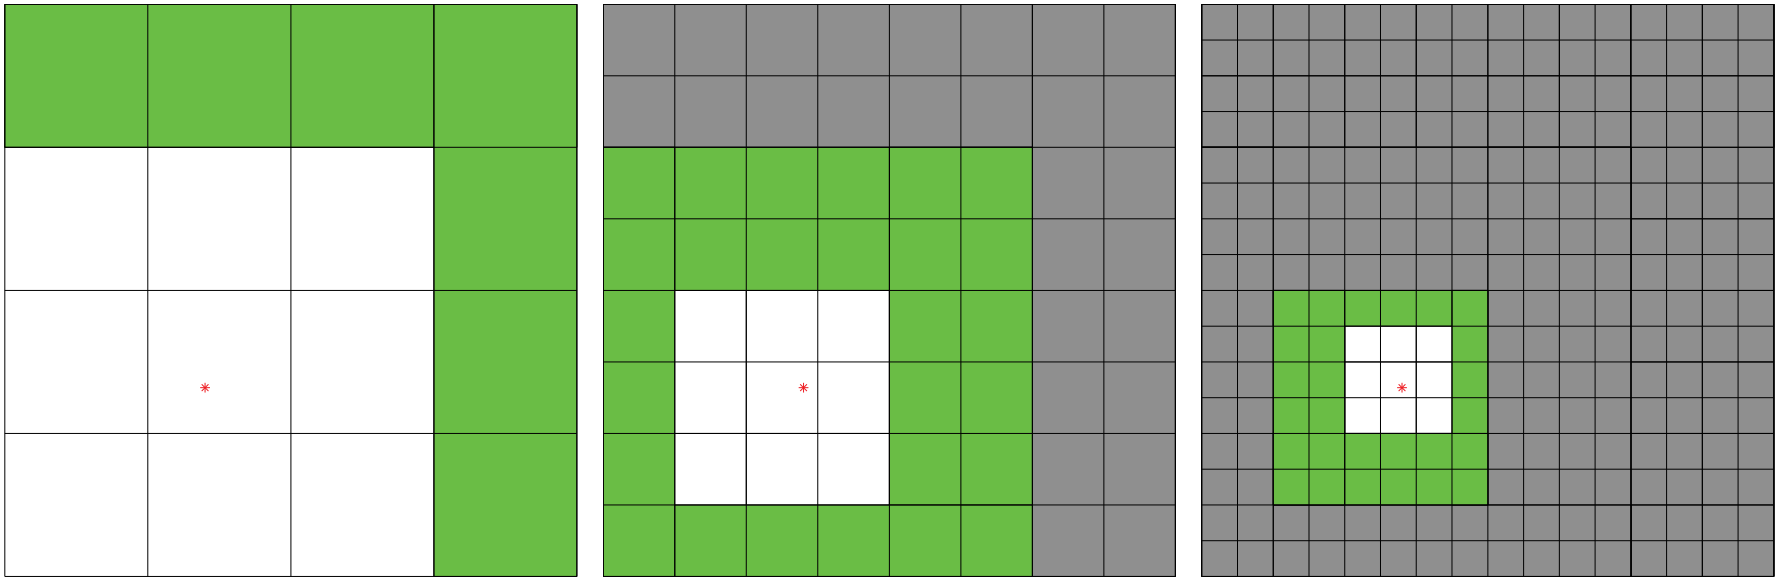

The fast multipole method

- The algorithm described here is $O(N \log N)$, and technically not the FMM per se.
- For the $O(N)$ algorithm, the FMM, further machinery is needed.
- See *A short course on fast multipole methods* by Beatson and Greengard if you are curious