

Fast Radon Transform via Fast Nonuniform FFTs on GPUs

Stefano Marchesini, Filipe Maia, J. Pien (cuda consultant), Chao Yang, Andre. Schiroztek, Alastair McDowell, Howard Padmore
Lawrence Berkeley National Laboratory, Berkeley, CA

Introduction

Fast Radon and inverse Radon transforms form the computational kernels of many image reconstruction algorithms. We have recently developed a reconstruction algorithm for X-ray phase contrast tomography using a compressive sensing technique. In our approach, we formulate the reconstruction problem as an L1-minimization problem, i.e.

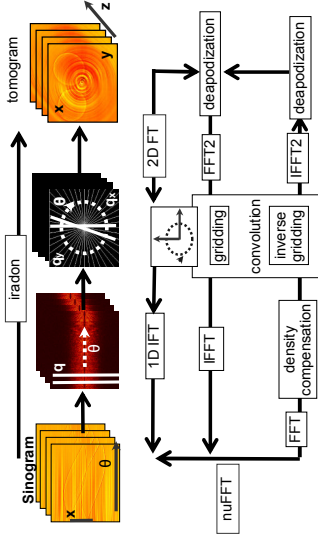
$$\min \|x\|_1$$

$$\text{subject to } \|filter(Radon(x) - data)\|_2 \leq \epsilon$$

where x is the phase contrast image to be reconstructed and $filter$ is a polynomial interpolation operator designed to reduce the ringing artifact in the reconstruction, and ϵ is a regularization parameter we choose in advance. We solve the L1-minimization problem by using the SPGL1 software developed by E. van den Berg and M. P. Friedlander [1]. The software requires us to provide a function to perform $y = Radon(x)$. Due to the large volume of data produced by fast detectors at beam line (BL32) at the advanced light source at Lawrence Berkeley National Lab, we would like perform phase contrast tomographic reconstruction in real time.

Fast Radon Transform

The 2D Radon transform of an image x can be implemented in a number of ways. One of the most efficient ways to perform $Radon(x)$ is to first perform a 2D FFT of x on a regular grid, and then interpolate the transform onto a polar grid before 1D inverse Fourier transforms are applied to interpolated points along the same radial lines.



The interpolation between the Cartesian and polar grid is known as can be carried out using a "gridding" algorithm that maintains desired accuracy and low computational complexity. The gridding algorithm essentially allows us to perform a non-uniform FFT.

The Gridding Algorithm

Inverse gridding: Convolve Kaiser Bessel (KB) kernel from regular samples (x, y) to regular grid points (x_i, y_i)

$$G_1(x_i, y_i) = \sum_{x, y} D(x, y) K(x - x_i, y - y_i)$$

Gridding: Convolve KB kernel from irregular samples (x_i, y_i) to regular grid points (x, y) .

$$G_2(x, y) = \sum_{x_i, y_i} D(x_i, y_i) K(x - x_i, y - y_i)$$


Density compensation:

$$D(x_i, y_i) = \sum_{x, y} \frac{1}{d} K(x - x_i, y - y_i)$$

constant

Deapodization:

GPU Implementation (prototype code in MATLAB and Jacket)



```

% Kaiser-Bessel filter
beta=2.3934*2;
% Kaiser-Bessel window, it could use look up tables (LUT):
KB=8(x,y) KB1(nj+2+1,x,beta)*KB1(nj+2+1,y,beta);
function w=KB1(M,x,beta)
ws=abs(besselj(0,beta));
w=besselj(0,beta*sgt(1-(2*x/M).^2))./ws;
    
```

Gridding:


- Loop through samples
- Sum shifted kernels

```

function G=reg2r(xi,yi,G,KB,nj)
% Input: xi,yi, irregular sample coordinates
% G: regular sample values
% KB: convolution kernel (nj: width)
% output: irregular sampled G
[Ms,Mj]=size(G); % sample values
ni=number(xi); % image size
jfi=nj-nj; % kernel index/ grid location
for ifi=1:ni; % loop over sampling points
nearx=round(xi(ifi)); % nearest grid point x
neary=round(yi(ifi)); % nearest grid point y
% calculate distance between xi and grid points jf+nearx
distx=xi(ifi)-nearx+jfi; %nj long vector
disty=yi(ifi)-neary+jfi; %nj long vector
foledey=mod(nearx+jfi,Ms)*1;%wrap around boundary
foledey=mod(neary+jfi,Mj)*1;%wrap around boundary
%add up G*KB
G(ifi)=sum(sum(G(foledey,foledey).*KB(distx,disty)));
end
    
```

Density compensation:

Loop through irregular samples for comparison



```


function D=irreg2reg(Mi,Yi,KB,as)
% xi,yi, irregular coordinates
% KB, kernel, siz, image boundary
% output: D density
ni=number(Mi);
D=zeros(siz,1);
for i=1:ni; % loop over all sampling points
% distance between i,th sampling point
% and all sampling points
distx=(xi(i)-xi); %1 long vector
disty=(yi(i)-yi); %1 long vector
%wrap around boundary
distx=distx-comnd(distx/size(1));
disty=disty-comnd(disty/size(2));
% Calculate density
D(i)=D(i)+sum(sum(KB(distx,disty)));
end
    
```

Inverse gridding:

- Loop through grid points
- Add sampled kernel
- Slow!

```

function G=irreg2reg(xr,yr,GT,KB,nj,siz)
% xr,yr, irregular grid points
% KB is the kernel, nj kernel width, siz is the size of the
% resulting grid
[Mr=Mr; % boundary of the image
Nr=Nr; % boundary of the image
G=zeros(siz); % initial blank image
jfi=-nj:nj; %kernel index/ grid location
for ifi=1:Mr; % loop over sampling points
nearx=round(xr(ifi)); % nearest grid point x
neary=round(yr(ifi)); % nearest grid point y
% calculate distance between xr and grid points jfi+nearx
distx=xr(ifi)-nearx+jfi; %nj long vector
disty=yr(ifi)-neary+jfi; %nj long vector
foledey=mod(nearx+jfi,Mr)*1;%wrap around boundary
foledey=mod(neary+jfi,Nr)*1;%wrap around boundary
%add up G*KB
G(ifi)=sum(sum(G(foledey,foledey).*KB(distx,disty)));
end
    
```



Improvement:

- Pre-bin (see [2])
- Divide among tiles
- Load balance

Performance:

- FFT (3 msec for 1k by 1k pixels)
- Gridding (5 msec for 180 by 1k pixels)
- Inverse gridding (1 msec)

We achieved 50x speedup on a Tesla GPU
Compared to a fast CPU!

[1] E. van den Berg and M. P. Friedlander, SIAM J. on Sci. Comp. 31, 890-912, 2008.

[2] www.nvidia.com/docs/IO/47905/ECE757_Project_Report_Gregerson.pdf