## Chapter-14 Application case study— non-Cartesian magnetic resonance imaging An introduction to statistical estimation methods



**FIGURE 14.1:** Scanner k-space trajectories and their associated reconstruction strategies: (A) Cartesian trajectory with FFT reconstruction, (B) Spiral (or non-Cartesian trajectory in general) followed by gridding to enable FFT reconstruction, (C) spiral (non-Cartesian) trajectory with linear solver based reconstruction. Note: \*Based on Fig 1 of Lustig et al. Fast Fourier Transform for Iterative MR Image Reconstruction, IEEE Int'l Symp. on Biomedical Imaging, 2004.



Courtesy of Keith Thulborn and Ian Atkinson, Center for MR Research, University of Illinois at Chicago

**FIGURE 14.2:** Non-Cartesian k-space sample trajectory and accurate linear-solver-based reconstruction enable new capabilities with exciting medical applications. The improved SNR enables reliable collection of in-vivo concentration data on chemical substance such as sodium in human tissues. The variation or shifting of sodium concentration gives early signs of disease development or tissue death. For example, the sodium map of a human brain shown in this figure can be used to give early indication of brain tumor tissue responsiveness to chemotherapy protocols, enabling individualized medicine.



**FIGURE 14.3:** An iterative linear-solver-based approach to reconstructing non-Cartesian k-space sample data.

```
(A)
                                    (B)
 for (m = 0; m < M; m++) {
                                       for (m = 0; m < M; m++) {
   phiMag[m] = rPhi[m]*rPhi[m] +
                                         rMu[m] = rPhi[m] * rD[m] +
                iPhi[m] * iPhi[m];
                                                  iPhi[m]*iD[m];
                                         iMu[m] = rPhi[m]*iD[m] -
                                                  iPhi[m]*rD[m];
   for (n = 0; n < N; n++) {
     expQ = 2*PI*(kx[m]*x[n] +
                   ky[m]*y[n] +
                                         for (n = 0; n < N; n++) {
                   kz[m]*z[n]);
                                           expFhD = 2*PI*(kx[m]*x[n] +
                                                           ky[m]*y[n] +
                                                           kz[m]*z[n]);
     rQ[n] +=phiMag[m]*cos(expQ);
      iQ[n] +=phiMag[m]*sin(expQ);
    }
                                           cArg = cos(expFhD);
 }
                                           sArg = sin(expFhD);
                                           rFhD[n] += rMu[m]*cArg -
                                                        iMu[m]*sArg;
                                           iFhD[n] += iMu[m]*cArg +
                                                       rMu[m]*sArg;
                                         }
                                       }
```

FIGURE 14.4: Computation of Q and FHD. (A) Q computation, (B) FHD computation.

```
global void cmpFhD(float* rPhi, iPhi, rD, iD,
   kx, ky, kz, x, y, z, rMu, iMu, rFhD, iFhD, int N) {
 int m = blockIdx.x * FHD THREADS PER BLOCK + threadIdx.x;
 rMu[m] = rPhi[m]*rD[m] + iPhi[m]*iD[m];
 iMu[m] = rPhi[m]*iD[m] - iPhi[m]*rD[m];
 for (int n = 0; n < N; n++) {
   floatexpFhD = 2*PI*(kx[m]*x[n] + ky[m]*y[n] + kz[m]*z[n]);
   floatcArg = cos(expFhD); floatsArg = sin(expFhD);
   rFhD[n] += rMu[m]*cArg - iMu[m]*sArg;
   iFhD[n] += iMu[m]*cArg + rMu[m]*sArg;
 }
}
```

**FIGURE 14.5:** First version of the FHD kernel. The kernel will not execute correctly due to conflicts between threads in writing into rFhD and iFhD arrays.

```
(A)
                                     (B)
 for (m = 0; m < M; m++) {
                                      for (m = 0; m < M; m++) {
   rMu[m] = rPhi[m] * rD[m] +
                                        rMu[m] = rPhi[m] * rD[m] +
             iPhi[m]*iD[m];
                                                  iPhi[m] *iD[m];
   iMu[m] = rPhi[m]*iD[m] -
                                         iMu[m] = rPhi[m]*iD[m] -
             iPhi[m]*rD[m];
                                                  iPhi[m] *rD[m];
                                       }
   for (n = 0; n < N; n++) {
                                      for (m = 0; m < M; m++) {
     expFhD = 2*PI*(kx[m]*x[n] +
                                        for (n = 0; n < N; n++) {
                     ky[m]*y[n] +
                                           expFhD = 2*PI*(kx[m]*x[n] +
                     kz[m]*z[n]);
                                                          ky[m]*y[n] +
                                                          kz[m]*z[n]);
     cArg = cos(expFhD);
     sArg = sin(expFhD);
                                          cArg = cos(expFhD);
                                           sArg = sin(expFhD);
     rFhD[n] += rMu[m]*cArg -
                  iMu[m]*sArg;
                                           rFhD[n] += rMu[m]*cArg -
     iFhD[n] += iMu[m]*cArg +
                                                       iMu[m]*sArg;
                  rMu[m]*sArg;
                                           iFhD[n] += iMu[m]*cArg +
   }
                                                       rMu[m]*sArg;
 }
                                        }
                                       }
```

**FIGURE 14.6:** Loop fission on the FHD computation. (A) FHD computation, (B) after loop fission.

```
__global___ void cmpMu(float* rPhi, iPhi, rD, iD, rMu, iMu)
{
    int m = blockIdx.x*MU_THREAEDS_PER_BLOCK + threadIdx.x;
    rMu[m] = rPhi[m]*rD[m] + iPhi[m]*iD[m];
    iMu[m] = rPhi[m]*iD[m] - iPhi[m]*rD[m];
}
```

FIGURE 14.7: cmpMu kernel.

```
global void cmpFhD(float* rPhi, iPhi, phimag,
      kx, ky, kz, x, y, z, rMu, imu, int N) {
  int m = blockIdx.x * FHD THREADS PER BLOCK + threadIdx.x;
  for (int n = 0; n < N; n++) {
    float expFhD = 2*PI*(kx[m]*x[n]+ky[m]*y[n]+kz[m]*z[n)];
    float CArg = cos(expFhD);
    float sArg = sin(expFhD);
    atomicAdd(&rFhD[n], rMu[m]*cArg - imu[m]*sArg);
    atomicAdd(&rFhD[n], iMu[m]*cArg + imu[m]*sArg);
  }
}
```

FIGURE 14.8: Second option of the FHD kernel.

```
(A)
                                      (B)
 for (m = 0; m < M; m++) {
                                       for (n = 0; n < N; n++) {
   for (n = 0; n < N; n++) {
                                         for (m = 0; m < M; m++) {
     expFhD = 2*PI*(kx[m]*x[n] +
                                           expFhD = 2*PI*(kx[m]*x[n] +
                     ky[m]*y[n] +
                                                           ky[m]*y[n] +
                     kz[m]*z[n]);
                                                           kz[m]*z[n]);
     cArg = cos(expFhD);
                                           cArg = cos(expFhD);
     sArg = sin(expFhD);
                                           sArg = sin(expFhD);
     rFhD[n] += rMu[m]*cArg -
                                           rFhD[n] += rMu[m]*cArg -
                  iMu[m]*sArg;
                                                        iMu[m]*sArg;
      iFhD[n] +=
                  iMu[m]*cArg +
                                           iFhD[n] += iMu[m]*cArg +
                  rMu[m]*sArg;
                                                        rMu[m]*sArg;
   }
                                         }
 }
                                       }
```

**FIGURE 14.9:** Loop interchange of the FHD computation. (A) Before loop interchange, (B) after loop interchange.

```
global void cmpFHd(float* rPhi, iPhi, phiMag,
     kx, ky, kz, x, y, z, rMu, iMu, int M) {
  int n = blockIdx.x * FHD THREADS PER BLOCK + threadIdx.x;
 for (int m = 0; m < M; m++) {
   float expFhD = 2*PI*(kx[m]*x[n]+ky[m]*y[n]+kz[m]*z[n]);
   float cArg = cos(expFhD);
   float sArg = sin(expFhD);
   rFhD[n] += rMu[m]*cArg - iMu[m]*sArg;
   iFhD[n] += iMu[m]*cArg + rMu[m]*sArg;
  }
}
```

**FIGURE 14.10:** Third option of the FHD kernel.

```
global void cmpFHd(float* rPhi, iPhi, phiMag,
       kx, ky, kz, x, y, z, rMu, iMu, int M) {
  int n = blockIdx.x * FHD_THREADS_PER_BLOCK + threadIdx.x;
  float xn r = x[n]; float yn r = y[n]; float zn r = z[n];
  float rFhDn r = rFhD[n]; float iFhDn r = iFhD[n];
  for (int m = 0; m < M; m++) {
    float expFhD = 2*PI*(kx[m]*xn r+ky[m]*yn r+kz[m]*zn r);
   float cArg = cos(expFhD);
    float sArg = sin(expFhD);
   rFhDn_r += rMu[m]*cArg - iMu[m]*sArg;
    iFhDn r += iMu[m]*cArg + rMu[m]*sArg;
  }
 rFhD[n] = rFhD r; iFhD[n] = iFhD r;
}
```

FIGURE 14.11: Using registers to reduce memory accesses in the FHD kernel.

```
_constant__ float
                   kx_c[CHUNK_SIZE],
                    ky c[CHUNK SIZE],
                    kz_c[CHUNK_SIZE];
void main() {
  for (int i = 0; i < M/CHUNK SIZE; i++);
    cudaMemcpyToSymbol(kx c, &kx[i*CHUNK SIZE], 4*CHUNK SIZE,
                      cudaMemCpyHostToDevice);
    cudaMemcpyToSymbol(ky_c, &ky[i*CHUNK_SIZE], 4*CHUNK_SIZE,
                      cudaMemCpyHostToDevice);
    cudaMemcpyToSymbol(ky_c, &ky[i*CHUNK_SIZE], 4*CHUNK_SIZE,
                      cudaMemCpyHostToDevice);
    ...
    cmpFHD<<<FHD THREADS PER BLOCK, N/FHD THREADS PER BLOCK>>>
       (rPhi, iPhi, phiMag, x, y, z, rMu, iMu, CHUNK SIZE);
  }
  /* Need to call kernel one more time if M is not */
  /* perfect multiple of CHUNK SIZE */
}
```

FIGURE 14.12: Chunking k-space data to fit into constant memory.

```
_global___ void cmpFHd(float* rPhi, iPhi, phiMag,
     x, y, z, rMu, iMu, int M) {
  int n = blockIdx.x * FHD THREADS PER BLOCK + threadIdx.x;
 float xn r = x[n]; float yn r = y[n]; float zn r = z[n];
  float rFhDn r = rFhD[n]; float iFhDn r = iFhD[n];
 for (int m = 0; m < M; m++) {
   float expFhD =
       2*PI*(kx c[m]*xn_r+ky_c[m]*yn_r+kz_c[m]*zn_r);
   float cArg = cos(expFhD);
    float sArg = sin(expFhD);
   rFhDn r += rMu[m]*cArg - iMu[m]*sArg;
    iFhDn r += iMu[m]*cArg + rMu[m]*sArg;
  }
 rFhD[n] = rFhD r; iFhD[n] = iFhD r;
}
```

FIGURE 14.13: Revised FHD kernel to use constant memory.



(B)

	kx[i]	ky[i] kz[i] kx ky k	٢Z
--	-------	---------------------	----

**FIGURE 14.14:** Effect of k-space data layout on constant cache efficiency. (A) k-space data stored in separate arrays, (B) k-space data stored in an array whose elements are structs.

```
struct kdata {
   float x, float y, float z;
};
 _constant__ struct kdata k_c[CHUNK_SIZE];
...
void main() {
 for (int i = 0; i < M/CHUNK\_SIZE; i++) {
   cudaMemcpyToSymbol(k c,k,12*CHUNK SIZE, cudaMemCpyHostToDevice);
   cmpFhD<<<FHD THREADS PER BLOCK, N/FHD THREADS PER BLOCK>>>(...);
  }
FIGURE 14.15: Adjusting k-space data layout to improve cache efficiency.
```

```
global void cmpFhD(float* rPhi, iPhi, phiMag,
    x, y, z, rMu, iMu, int M) {
 int n = blockIdx.x * FHD THREADS PER BLOCK + threadIdx.x;
 float xn_r = x[n]; float yn_r = y[n]; float zn_r = z[n];
 float rFhDn r = rFhD[n]; float iFhDn r = iFhD[n];
 for (int m = 0; m < M; m++) {
   float expFhD = 2*PI*(k[m].x*xn r+k[m].y*yn r+k[m].z*zn r);
   float cArg = cos(expFhD);
   float sArg = sin(expFhD);
   rFhDn_r += rMu[m]*cArg - iMu[m]*sArg;
   iFhDn r += iMu[m]*cArg + rMu[m]*sArg;
  }
 rFhD[n] = rFhD_r; iFhD[n] = iFhD_r;
}
```

**FIGURE 14.16:** Adjusting for the k-space data memory layout in the FHD kernel.

```
_global__ void cmpFHd(float* rPhi, iPhi, phiMag,
     x, y, z, rMu, iMu, int M) {
 int n = blockIdx.x * FHD THREADS PER BLOCK + threadIdx.x;
 float xn_r = x[n]; float yn_r = y[n]; float zn_r = z[n];
 float rFhDn_r = rFhD[n]; float iFhDn_r = iFhD[n];
 for (int m = 0; m < M; m++) {
   float expFhD = 2*PI*(k[m].x*xn r+k[m].y*yn r+k[m].z*zn r);
   float cArg = cos(expFhD);
   float sArg = sin(expFhD);
   rFhDn_r += rMu[m]*cArg - iMu[m]*sArg;
   iFhDn_r += iMu[m]*cArg + rMu[m]*sArg;
 }
 rFhD[n] = rFhD r; iFhD[n] = iFhD r;
}
```

**FIGURE 14.17:** Using hardware \_\_\_\_\_sin() and \_\_\_\_cos() functions.

$$MSE = \frac{1}{mn} \sum_{i} \sum_{j} (l(i, j) - l_0(i, j))^2$$
$$PSNR = 20 \log_{10} \left( \frac{\max(l_0(i, j))}{\sqrt{MSE}} \right)$$

**FIGURE 14.18:** Metrics used to validate the accuracy of hardware functions. I<sup>0</sup> is perfect image. I is reconstructed image. PSNR is peak signal-to-noise ratio.







PSNR = 27.6 dB

41.7% error PSNR = 16.8 dB



(4) CPU.SP 12.0% error PSNR = 27.6 dB

(6) GPU.RegAlloc 12.1% error PSNR = 27.6 dB





12.1% error

PSNR = 27.6 dB

(5) GPU.Base

12.1% error

PSNR = 27.6 dB

(9) GPU.FastTrig

(7) GPU.Coalesce 12.1% error PSNR = 27.6 dB

12.1% error PSNR = 27.5 dB

FIGURE 14.19: Validation of floating-point precision and accuracy of the different FHD implementations.