

Project: Radar Resampling

University of Washington

EE 590 Winter 2017

Eric Wilson and Darrell Ross

February 27, 2017

Contents

1	Introduction	1
1.1	Process Overview	1
1.2	Report Organization	1
2	Concepts	2
2.1	Sample Interval	2
2.2	Least Squares Approximation	2
2.2.1	QR Decomposition	2
2.2.2	Givens Rotation	2
2.2.3	Back Substitution	3
2.3	Resampling	3
3	Algorithms	4
3.1	Interval Sample	4
3.2	Least Squares Approximation	4
3.2.1	Givens Rotations and QR Decomposition	4
3.2.2	Back Substitution	5
3.3	Resampling	5
4	Performance	6
4.1	Interval Sample	6
4.2	Givens Rotations and QR Decomposition	6
4.3	Back Substitution	7
4.4	Resampling	7
5	Analysis	8
5.1	Resampling	8
5.2	QR Decomposition	8
6	Conclusions	9
	Appendix A Algorithm Examples	10
A.1	Least Squares Approximation	10
A.2	Givens Rotation and QR Decomposition	10
A.3	Back Substitution	10
	Appendix B Kernels	10
B.1	Resampling()	10
	Appendix C Program Instructions	11
	Appendix D Attachments	12

1 Introduction

Professor John D. Sahr in the Electrical Engineering department at the University of Washington (UW) runs a passive receiver that samples at 4 giga-samples per second (GSPS). He is interested in down-sampling to 3.5 GSPS in real-time by leveraging CUDA on Nvidia Kepler K10 graphics cards.

This report serves two purposes. First, it meets project report requirements for EE590: Applied High-Performance GPU Computing. Second, it serves as a report for Professor Sahr on how we solved the down-sampling problem using parallelism.

1.1 Process Overview

An overview of our resampling process is shown in Figure 1. Walking through the figure, the signal from the receiver is sampled for t seconds. After we acquire the sample data, we perform a Least Squares Approximation (LSA) to get the curve of best fit with a sixth-order polynomial. Finally, the data is sampled at the lower rate of 3.5 GPSP.




Figure 1: An overview of parts of the resampler.

1.2 Report Organization

- [Concepts](#) provides a conceptual review of each piece of the process.
- [Algorithms](#) covers the sequential and OpenCL algorithms used.
- [Analysis](#) evaluates performance of each piece of the process.

2 Concepts

Each of the entries from Figure 1 are covered in this section.

2.1 Sample Interval

As it turns out, the sample interval is not a simple choice. The interval chosen determines how long a single chunk of data will take to solve. Depending on the size, we may be able to pipeline several solutions at the same time. But pipelining will suffer from diminishing returns due to overhead. So we really need to calculate how much work it will take to run a single process all the way through. A simulation will probably be our best bet to determine optimal sample interval size.

2.2 Least Squares Approximation

The Least Squares Approximation (LSA) is a way of finding the curve of best fit to the sample data. We were given that the polynomial degree would be $k \leq 6$. Performing a Least Squares Approximation (LSA) can be done in many ways. We chose QR Decomposition followed by Back Substitution. Another option considered was Single Value Decomposition (SVD) which is covered in Appendix ??.

2.2.1 QR Decomposition

QR decomposition involves taking a matrix A and decomposing it into an upper triangular matrix R and an orthogonal matrix Q such that $A = QR$. If A is $m \times n$ where $m > n$, then the bottom $(m - n)$ rows of an upper triangular matrix will consist of all zeros. In this situation, R is partitioned. Since $R_2 = 0$, the final result is Q_1R_1 .

$$\begin{aligned} A &= QR \\ &= Q \begin{bmatrix} R_1 \\ R_2 \end{bmatrix} \\ &= [Q_1, Q_2] \begin{bmatrix} R_1 \\ R_2 \end{bmatrix} \\ &= Q_1R_1 \end{aligned}$$

There are multiple techniques available for reaching the upper triangular matrix. We decided to go with the Givens Rotation. Other considered options are covered in Appendix A.2.

2.2.2 Givens Rotation

A sequence of Givens Rotations are used to construct the upper-right triangular matrix R and the combination of the Givens Rotations results in the matrix Q . Each multiplication of the matrix A by the Givens Rotation matrix G results in a zero being inserted into A , as shown in (1).

$$\begin{bmatrix} c & -s \\ s & c \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} r \\ 0 \end{bmatrix} \tag{1}$$

The Givens rotation matrix G is constructed from c and s where

$$\begin{aligned}c &= \frac{a}{r} \\s &= -\frac{b}{r} \\r &= \sqrt{a^2 + b^2}\end{aligned}$$

2.2.3 Back Substitution

Back Substitution is a quick method of solving for x in (2) without having to invert R . This is quick when R is an upper triangular matrix as it is in our case. An example is provided in Appendix A.3. For our usage, the QR Decomposition with the Givens Rotation will provide us the R and Q matrices so we can quickly solve for x .

$$R \cdot x = Q^T \cdot b \tag{2}$$

2.3 Resampling

Resampling is the process of taking a set of data acquired at one sample rate and converting it to another sample rate. In the case of passive RADAR resampling, the input data rate to the resampler is 4 GSPS and the output rate is 3.5 GSPS. This is a $7/8$ resample rate. To achieve a lower sample rate the discrete time input signal is reconstructed as an analog (continuous-time) signal using a polynomial curve fit. Least squares approximation calculates the coefficients for the polynomial curve fit. The polynomial representing the continuous time signal can now be evaluated at the new sample rate to reconstruct the original discrete time signal at a new sample rate.

3 Algorithms

3.1 Interval Sample

The interval sample will determine the number of data points. The number of data points determines the height of the A matrix which will have a maximum width of $k + 1$ where k is the order of the polynomial – ours is maximum 6th order. There are several constraints to consider:

- A sixth-order polynomial will fail to accurately match the data if too many data points are used. It is important to know the approximate frequencies of the data so we can avoid this issue.
- The GPU has a limit on how much it can process in parallel which will also constrain the sample size.

Thus, choosing the interval requires knowing the rest of our algorithm and the limits of the GPU being used. The timing of the full algorithm t_{solution} can be split into the three segments depicted in Figure 2. In order to maintain real-time output, our pipeline count p must be such that:

$$t_{\text{solution}} = \Delta t \cdot p$$

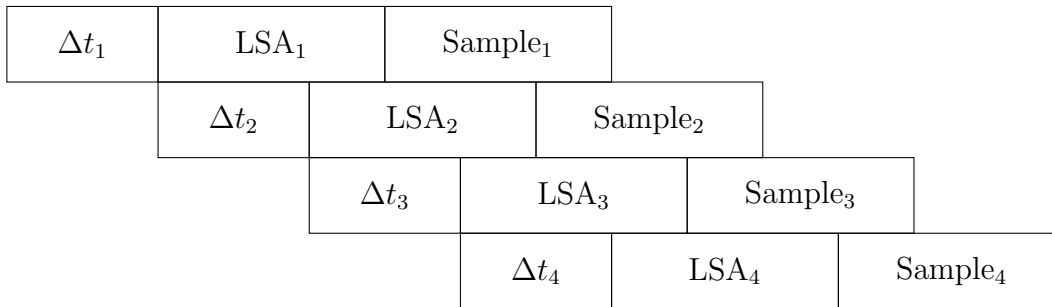


Figure 2: Pipelines with the interval and algorithm.

Once we have a full set of data describing expected execution time and we know what the frequencies are and we know the GPU we are running on, we will finally be able to choose the optimal interval. Until then, we will use sample data.

3.2 Least Squares Approximation

Although the LSA constitutes two algorithms combined together, analysis of them makes more sense when done in one block.

3.2.1 Givens Rotations and QR Decomposition

Givens rotations have a limited scope when it comes to speedup via parallelism. This is due to some interdependency between the order in which the Givens rotations are applied. Because the Givens rotations are only applied to two rows during an iteration, the Givens rotations can be parallelized on different columns after a certain number of iterations. If each element in the R matrix is zeroed after iteration n , then the following matrix shows the iteration when each element can be zeroed.

$$\begin{bmatrix} x \\ 6 & x \\ 5 & 7 & x \\ 4 & 6 & 8 & x \\ 3 & 5 & 7 & 9 & x \\ 2 & 4 & 6 & 8 & 10 & x \\ 1 & 3 & 5 & 7 & 9 & 11 & x \end{bmatrix} \quad (3)$$

A Givens rotation kernel can be started on the first column at the start of algorithm execution and works its way up the first column. When the Givens rotation kernel on column 1 gets row 5 on iteration 3, the next Givens rotation kernel can be started and works up column 2. This process can be continued until the last column, which in our case is 7 since the largest number of polynomial coefficients we will be using is $k = 6$ and matrix R has $k + 1$ columns. This means that we can have a maximum of 7 Givens kernels running in parallel at the same time. The following pseudocode implements a method to start each Givens rotation kernel on a column after a certain iteration.

```
function [matrix R, matrix Q] = GivensKernel(matrix A, numRows)
    R = A;
    for (iter = 0..12)
        [R, Q1] = GivensKernel(R[column 1])
        if (iter > 2)
            [R, Q2] = GivensKernel(R[column 2])
        if (iter > 4)
            [R, Q3] = GivensKernel(R[column 3])
        if (iter > 6)
            [R, Q4] = GivensKernel(R[column 4])
        if (iter > 8)
            [R, Q5] = GivensKernel(R[column 5])
        if (iter > 10)
            [R, Q6] = GivensKernel(R[column 6])
        if (iter > 11)
            [R, Q7] = GivensKernel(R[column 7])
    endfor
    Q = Q1*Q2*Q3*Q4*Q5*Q6*Q7;
endfunction
```

3.2.2 Back Substitution

Back Substitution is really a linear operation. For our problem the final matrix we solve will be small enough (7x7) to make parallelization unnecessary. We do not currently intend to use a kernel for this part of the solution.

3.3 Resampling

Resampling is simply passing our new time interval values into the continuous LSA solution. Since we are guaranteed an order ≤ 6 , we can do this in OpenCL using a *cl_float8* for the coefficients. This produces the Resample() kernel shown in Appendix B.1.

4 Performance

This sections covers performance estimates.

4.1 Interval Sample

[Section needs populated once we have run some full simulations and drawn conclusions about what the right sample interval is.]

4.2 Givens Rotations and QR Decomposition

A single Givens Rotation requires work in the following manner. First the calculation of the Givens values:

$$\begin{aligned} c &= \frac{a}{r} & W &= 1 \\ s &= \frac{b}{r} & W &= 1 \\ r &= \sqrt{a^2 + b^2} & W &\approx 4 \\ \therefore W_{cs} &= 6 \end{aligned}$$

Next, the Givens matrix is multiplied by the A matrix which has size $M \times K$. The Givens matrix will only change two rows of the A matrix. This reduces the work to:

$$\begin{aligned} W_{gr} &= 3 \cdot 2 \cdot K = 6K && 2 \text{ products, 1 sum, 2 rows, } K \text{ columns} \\ \therefore W_{GR} &= W_{cs} + W_{gr} = 6 + 6K \end{aligned}$$

In addition, we must maintain the $K \times K$ Q matrix by multiplying it by the new $K \times K$ G matrix each time. This follows the same pattern as the product with the A matrix in terms of only affecting two rows:

$$W_{QG} = 3 \cdot 2 \cdot M = 6M \quad 2 \text{ products, 1 sum, 2 rows, } M \text{ samples}$$

So a single Givens Rotation calculation runs us:

$$W_T = 6 + 6K + 6M$$

The number of Givens Rotations necessary will be equivalent to the original $M \times K$ matrix minus the upper triangular portion of $2(k+1)$ entries.

$$W_{\text{total}} = (MK - 2(K+1))(6 + 6K + 6M)$$

Calculating the memory operations is a bit more straight forward assuming we write a kernel that can do all givens rotations without having to return each time:

$$\begin{aligned} Q_r &= DMK && M \times K \text{ matrix with } D \text{ bytes per value} \\ Q_w &= DMM && M \times M \text{ matrix with } D \text{ bytes per value} \\ D &= 4 && 4 \text{ bytes per float} \\ \therefore Q &= Q_r + Q_w = 4M^2 + 4MK \end{aligned}$$

$$\therefore AI = \frac{W}{Q} = \frac{(MK - 2(K + 1))(6 + 6K + 6M)}{4M^2 + 4MK}$$

This is all easier to see with an example. For our example, we will use an input matrix of size 5x3. This would map to our problem as 5 samples of a second order polynomial. Running the numbers:

$$\begin{aligned} W_{\text{total}} &= 9(6 + 18 + 30) = 486 \\ Q &= 160 \\ AI &= \frac{486}{160} = 3.03 \end{aligned}$$

4.3 Back Substitution

Assuming all coefficients are non-zero and not one:

$W_7 = 1$	1 product
$W_6 = 3$	2 products, 1 sum
$W_5 = 5$	3 products, 2 sums
$W_4 = 7$	4 products, 3 sums
...	
$W_1 = 13$	7 products, 6 sums
$\therefore W_T = 49$	

4.4 Resampling

Reviewing the Resample() kernel, we can calculate the Arithmetic Intensity AI . In this instance, we assume that the work done by the power function W_{pow} is equivalent to $P - 1$ where P is the exponent value. Applying this to an evaluation for a single time x_i produces the output y_i :

$$\begin{aligned} y_i &= a_0x_i^0 + a_1x_i^1 + a_2x_i^2 + a_3x_i^3 + a_4x_i^4 + a_5x_i^5 + a_6x_i^6 \\ \therefore W &= 1N + 2N + 3N + 4N + 5N + 6N + 5 = 26N && \text{21 products and 5 sums for each input N} \\ Q_r &= DN + 8D && \text{read N entries and one float8} \\ Q_w &= DN && \text{write N entries} \\ \therefore Q &= Q_r + Q_w = 2DN + 8D \\ D &= 4 && \text{4 bytes per float} \\ \therefore Q &= 8N + 32 \\ AI &= \frac{W}{Q} = \frac{26N}{8N + 32} \approx \frac{26}{8} \\ \therefore AI &= 3.25 \end{aligned}$$

5 Analysis

5.1 Resampling

Performance was evaluated using an Intel HD 5500 GPU which has a single-precision floating point maximum of 38.4 GFLOPS and a bandwidth of 12.8 Gbps. With these specs, an AI of at least $38.4/12.8=3$ FLOPS/byte is desired. Since this kernel achieves 3.25 FLOPS/byte, it should be an effective application of parallelism.

When analyzed using OpenCL Kernel Development feature in Intel CodeBuilder in Visual Studio 2015, the GPU was faster as shown in Table 1. Presumably this difference will be more stark on

Processor	Sample Size	Time(ms)
Intel® Core™ i3-5010U CPU @ 2.1GHz	1024	0.078
Intel® HD Graphics 5500	1024	0.041

Table 1: Simulation results for the *Resample* kernel.

higher powered GPUs with larger sample sizes. For reference, a sample size of 1024 would be a Δt of only $0.256 \mu s$.

5.2 QR Decomposition

The sequential QR decomposition algorithm was executed on an Intel I7-6500U CPU. The input data size was a 32x32 matrix. Execution time was averaged over 10,000 iterations.

Processor	Sample Size	Time(ms)
Intel® Core™ i7-6500U CPU @ 2.5GHz	1024	0.388

Table 2: Simulation results for the *QR* kernel.

6 Conclusions

Appendix A Algorithm Examples

This section provides working examples for each algorithm used to help demonstrate how they work.

A.1 Least Squares Approximation

A.2 Givens Rotation and QR Decomposition

A.3 Back Substitution

Recall that as we complete our QR Decomposition, we have our Q matrix and our R matrix and we know b so only to solve for x in the following:

$$R \cdot x = Q^T \cdot b$$

Solving by example:

$$\left[\begin{array}{ccc|c} 1 & -2 & 1 & 4 \\ 0 & 1 & 6 & -1 \\ 0 & 0 & 1 & 2 \end{array} \right]$$

$$\begin{aligned} x - 2y + z &= 4 \\ y + 6z &= -1 \\ z &= 2 \end{aligned}$$

Back substitution means that we plug in the last entry, $z = 2$ to the second equation to solve for y and then plug the z and the results for y into the first equation to solve for x .

$$\begin{aligned} & \Rightarrow z = 2 \\ \therefore y + 12 &= -1 & \Rightarrow y = 11 \\ \therefore x - 22 + 2 &= 4 & \Rightarrow x = 28 \end{aligned}$$

Appendix B Kernels

B.1 Resampling()

```
__kernel void Resample(float8 coeffs, __global float* t, __global float* result)
{
    const unsigned int id = get_global_id(0);
    float val = t[id];
    float out = 0;
    for(unsigned int i=0; i<6; ++i)
        out += coeffs[i]*pow(val,i);

    result[id] = out;
}
```

Appendix C Program Instructions

When we complete our program, this section will contain instructions on how to use it.

Appendix D Attachments

- *control.cpp* - C++ code used to execute sequential and kernel functions.

```

1  #include "CL/cl.h"
2  #include "ocl.h"
3  #include "tools.h"
4  #include "utils.h"
5  #include "data.h"
6  #include "control.h"
7  #include "profiler.h"
8  #include "enums.h"
9  #include <iostream>
10 #include <vector>
11 #include <algorithm>
12
13 namespace
14 {
15     const char* FILENAME = "resample.cl";
16 }
17
18 ControlClass::ControlClass()
19     : GroupManager("Control")
20 {
21     groups_ = GroupFactory();
22 }
23
24
25 std::map<int, ProblemGroup*> ControlClass::GroupFactory()
26 {
27     std::map<int, ProblemGroup*> pgs;
28
29     ProblemGroup* InputControl = GroupManagerInputControlFactory();
30     pgs[InputControl->GroupNum()] = InputControl;
31
32     ProblemGroup* projectFuncs = new ProblemGroup(1, "Control");
33     projectFuncs->problems_[projectFuncs->problems_.size() + 1] = new Problem(&exCL_Resample, "OpenCL: Apply
34     sixth-order polynomial");
35     projectFuncs->problems_[projectFuncs->problems_.size() + 1] = new Problem(&exSeq_Resample, "Sequential: Apply
36     sixth-order polynomial");
37     pgs[projectFuncs->GroupNum()] = projectFuncs;
38     return pgs;
39 }
40
41 ////////////////////////////////////////////////// RESAMPLE USING POLYNOMIAL APPROXIMATION //////////////////////////////////////
42 int exCL_Resample(ResultsStruct* results)
43 {
44     cl_int err;
45     ocl_args ocl(CL_DEVICE_TYPE_GPU);
46
47     // Create Local Variables and Allocate Memory
48     // The buffer should be aligned with 4K page and size should fit 64-byte cached line
49     cl_uint sampleSize = 1024;
50     cl_uint optimizedSizeFloat = ((sizeof(cl_float) * sampleSize - 1) / 64 + 1) * 64;

```

```

51  cl_float* inputA = (cl_float*)_aligned_malloc(optimizedSizeFloat, 4096);
52  cl_float* outputC = (cl_float*)_aligned_malloc(optimizedSizeFloat, 4096);
53  if (NULL == inputA || NULL == outputC)
54  {
55      LogError("Error: _aligned_malloc failed to allocate buffers.\n");
56      return -1;
57  }
58  // Generate Random Input
59  data::generateInputCLSeq(inputA, sampleSize, 1);
60
61  // Create OpenCL buffers from host memory for use by Kernel
62  cl_float8      coeffs = {1,2,3,4,5,6,7,0};
63  cl_mem         srcA;           // hold first source buffer
64  cl_mem         dstMem;        // hold destination buffer
65  if (CL_SUCCESS != CreateReadBufferArg(&ocl.context, &srcA, inputA, sampleSize, 1))
66      return -1;
67  if (CL_SUCCESS != CreateWriteBufferArg(&ocl.context, &dstMem, outputC, sampleSize, 1))
68      return -1;
69
70  // Create and build the OpenCL program - imports named cl file.
71  if (CL_SUCCESS != ocl.CreateAndBuildProgram(FILENAME))
72      return -1;
73
74  // Create Kernel - kernel name must match kernel name in cl file
75  ocl.kernel = clCreateKernel(ocl.program, "Resample", &err);
76  if (CL_SUCCESS != err)
77  {
78      LogError("Error: clCreateKernel returned %s\n", TranslateOpenCLError(err));
79      return -1;
80  }
81
82  // Set OpenCL Kernel Arguments - Order Indicated by Final Argument
83  if (CL_SUCCESS != SetKernelArgument(&ocl.kernel, &coeffs, 0))
84      return -1;
85  if (CL_SUCCESS != SetKernelArgument(&ocl.kernel, &srcA, 1))
86      return -1;
87  if (CL_SUCCESS != SetKernelArgument(&ocl.kernel, &dstMem, 2))
88      return -1;
89
90  // Enqueue Kernel (wrapped in profiler timing)
91  ProfilerStruct profiler;
92  profiler.Start();
93  size_t globalWorkSize[1] = { sampleSize };
94  // hard code work group size after finding optimal solution with KDF Sessions
95  size_t localWorkSize[1] = { 16 };
96  if (CL_SUCCESS != ocl.ExecuteKernel(globalWorkSize, 1, localWorkSize))
97      return -1;
98  profiler.Stop();
99  float runTime = profiler.Log();
100
101  if (!SKIP_VERIFICATION)
102  {

```



```

103     // Map Host Buffer to Local Data
104     cl_float* resultPtr = NULL;
105     if (CL_SUCCESS != MapHostBufferToLocal(&ocl.commandQueue, &dstMem, sampleSize, 1, &resultPtr))
106     {
107         LogError("Error: clEnqueueMapBuffer failed.\n");
108         return -1;
109     }
110
111     // VERIFY DATA
112     // We mapped dstMem to resultPtr, so resultPtr is ready and includes the kernel output !!!
113     // Verify the results
114     bool failed = false;
115     /// @TODO WRITE SEQUENTIAL VERIFICATION CODE
116     /*
117     float cumSum = 0.0;
118     for (size_t i = 0; i < sampleSize; ++i)
119     {
120         cumSum += inputA[i];
121         if (resultPtr[i] != cumSum)
122         {
123             LogError("Verification failed at %d: Expected: %f. Actual: %f.\n", i, cumSum, resultPtr[i]);
124             failed = true;
125         }
126     }
127     */
128     if (!failed)
129         LogInfo("Verification passed.\n");
130
131     // Unmap Host Buffer from Local Data
132     if (CL_SUCCESS != UnmapHostBufferFromLocal(&ocl.commandQueue, &dstMem, resultPtr))
133         LogInfo("UnmapHostBufferFromLocal Failed.\n");
134 }
135
136 _aligned_free(inputA);
137 _aligned_free(outputC);
138
139 if (CL_SUCCESS != clReleaseMemObject(srcA))
140     LogError("Error: clReleaseMemObject returned '%s'.\n", TranslateOpenCLError(err));
141 if (CL_SUCCESS != clReleaseMemObject(dstMem))
142     LogError("Error: clReleaseMemObject returned '%s'.\n", TranslateOpenCLError(err));
143
144 results->WindowsRunTime = runTime;
145 results->HasWindowsRunTime = true;
146 results->OpenCLRunTime = ocl.RunTimeMS();
147 results->HasOpenCLRunTime = true;
148 return 0;
149 }
150
151 int exSeq_Resample(ResultsStruct* results)
152 {
153     return 0;
154 }

```

```

155
156
157 //////////////// QR DECOMPOSITION ////////////////
158 int exCL_QRD(ResultsStruct* results)
159 {
160     return 0;
161 }
162
163 /*
164  * Sequential QR decomposition function
165  */
166 void QR(cl_float* R, cl_float* Q, cl_uint arrayWidth, cl_uint arrayHeight)
167 {
168     cl_float a;
169     cl_float b;
170     cl_float c;
171     cl_float s;
172     cl_float r;
173     cl_float Rnew1[2048];
174     cl_float Rnew2[2048];
175     cl_float Qnew1[2048];
176     cl_float Qnew2[2048];
177     for (int j = 0; j < arrayWidth; j++)
178     {
179         for (int i = arrayHeight - 1; i > j; i--)
180         {
181             // Calculate Givens rotations
182             a = R[arrayWidth * (i - 1) + j];
183             b = R[arrayWidth * i + j];
184             r = sqrt(a * a + b * b);
185             c = a / r;
186             s = -b / r;
187             // Zero out elements in R matrix
188             for (int k = j; k < arrayWidth; k++)
189             {
190                 Rnew1[k] = R[arrayWidth * (i - 1) + k] * c - R[arrayWidth * i + k] * s;
191                 Rnew2[k] = R[arrayWidth * (i - 1) + k] * s + R[arrayWidth * i + k] * c;
192             }
193             // Copy new values back to R matrix
194             for (int k = j; k < arrayWidth; k++)
195             {
196                 R[arrayWidth * (i - 1) + k] = Rnew1[k];
197                 R[arrayWidth * i + k] = Rnew2[k];
198             }
199             // Update Q matrix
200             for (int k = 0; k < arrayHeight; k++)
201             {
202                 Qnew1[k] = Q[arrayHeight * (i - 1) + k] * c + Q[arrayHeight * i + k] * s;
203                 Qnew2[k] = -Q[arrayHeight * (i - 1) + k] * s + Q[arrayHeight * i + k] * c;
204             }
205             for (int k = 0; k < arrayHeight; k++)
206             {

```

```

207         Q[arrayHeight * (i - 1) + k] = Qnew1[k];
208         Q[arrayHeight * i + k] = Qnew2[k];
209     }
210 }
211 }
212 }
213 }
214 }
215 int exSeq_QRD(ResultsStruct* results)
216 {
217     const cl_uint arrayWidth = 3;
218     const cl_uint arrayHeight = 5;
219     cl_uint numIter = 10000; // Number of iterations for runtime averaging
220
221         // allocate working buffers.
222         // the buffer should be aligned with 4K page and size should fit 64-byte cached line
223     cl_uint optimizedSize = ((sizeof(cl_float) * arrayWidth * arrayHeight - 1) / 64 + 1) * 64;
224     cl_float* A = (cl_float*)_aligned_malloc(optimizedSize, 4096);
225
226     optimizedSize = ((sizeof(cl_float) * arrayHeight * arrayHeight - 1) / 64 + 1) * 64;
227     cl_float* Q = (cl_float*)_aligned_malloc(optimizedSize, 4096);
228
229     cl_float Atmp[] = { 0.8147, 0.0975, 0.1576,
230                       0.9058, 0.2785, 0.9706,
231                       0.1270, 0.5469, 0.9572,
232                       0.9134, 0.9575, 0.4854,
233                       0.6324, 0.9649, 0.8003 };
234
235     cl_float Qtmp[] = { 1.0, 0, 0, 0, 0,
236                       0, 1.0, 0, 0, 0,
237                       0, 0, 1.0, 0, 0,
238                       0, 0, 0, 1.0, 0,
239                       0, 0, 0, 0, 1.0 };
240
241     if (NULL == A)
242     {
243         LogError("Error: _aligned_malloc failed to allocate buffers.\n");
244         return -1;
245     }
246
247     // Initialize A
248     for (int i = 0; i < arrayWidth * arrayHeight; i++)
249     {
250         A[i] = Atmp[i];
251     }
252
253     // Initialize Q
254     for (int i = 0; i < arrayHeight * arrayHeight; i++)
255     {
256         Q[i] = Qtmp[i];
257     }
258

```

```
259     // add
260     ProfilerStruct profiler;
261     profiler.Start();
262
263     QR(A, Q, arrayWidth, arrayHeight);
264
265     profiler.Stop();
266     float runTime = profiler.Log();
267
268     _aligned_free(A);
269
270     results->WindowsRunTime = (double)runTime;
271     results->HasWindowsRunTime = true;
272     return 0;
273
274     return 0;
275 }
```

References

- [1] S. C. Kim and S. S. Bhattacharyya. “Implementation of a low-complexity low-latency arbitrary resampler on GPUs”. In: *2014 IEEE Dallas Circuits and Systems Conference (DCAS)*. 2014, pp. 1–4. DOI: [10.1109/DCAS.2014.6965333](https://doi.org/10.1109/DCAS.2014.6965333).
- [2] Mitra. *Digital Signal Processing: A Computer Based Approach*. 4th International edition. McGraw Hill Higher Education, 2010. ISBN: 9780071289467. URL: <http://amazon.com/o/ASIN/0071289461/>.
- [3] Eric W. Weisstein. *Least Squares Fitting–Polynomial*. From *MathWorld—A Wolfram Web Resource*. Last visited on 21/2/2017. URL: <http://mathworld.wolfram.com/LeastSquaresFittingPolynomial.html>.