

NAG Numerical Functions for GPUs Function Document

nag_gpu_mrg32k3a_exp

1 Purpose

nag_gpu_mrg32k3a_exp generates n values x_i from an exponential distribution with mean 1.0.

The initialization function **nag_gpu_mrg32k3a_init** must be called prior to the first call to **nag_gpu_mrg32k3a_exp**.

2 Specification

```
#include <nag_gpu.h>
void nag_gpu_mrg32k3a_exp (int nb, int nt, int np, float *d_P)
void nag_gpu_mrg32k3a_exp (int nb, int nt, int np, double *d_P)
```

3 References

None.

4 Arguments

- 1: **nb** – int *Input*
On entry: the number of GPU blocks.
Constraint: $1 \leq \text{nb}$.
- 2: **nt** – int *Input*
On entry: the number of GPU threads per block. When $\text{nt} = 1$, the ordering of the output values x_i is identical to the standard sequential storage, and a default value of nt is chosen internally for efficient execution. For other choices of nt the ordering is chosen for efficient use of bandwidth and storage and it is recommended that nt is set to a multiple of 64.
Constraint: $1 \leq \text{nt}$.
- 3: **np** – int *Input*
On entry: the number of random values to be generated by each thread.
Constraint: $1 \leq \text{np}$.
- 4: **d_p[$\text{np} \times \text{nb} \times \text{nt}$]** – float * *Output*
 5: **d_p[$\text{np} \times \text{nb} \times \text{nt}$]** – double * *Output*

This overloaded parameter can take type double or float.

The array must be allocated in the *device memory space*.

When $\text{nt} = 1$, the ordering of the output values, $x_{n,\tau,\beta}$, is identical to the standard sequential storage, $n + \text{np} \times \tau + \text{nt} \times \text{np} \times \beta$, where the indices, $n = 0, 1, \dots, \text{np} - 1$, $\tau = 0, 1, \dots, \text{nt} - 1$ and $\beta = 0, 1, \dots, \text{nb} - 1$ denote the element, thread and block respectively.

Otherwise $x_{n,\tau,\beta}$, the n th element generated by thread τ in block β , is stored at location $\tau + \text{nt} \times n + \text{nt} \times \text{np} \times \beta$.

On exit: the $\text{np} \times \text{nb} \times \text{nt}$ pseudorandom numbers from an exponential distribution with mean 1.0.

5 Error Indicators and Warnings

No argument constraint checking is carried out by this function. You can insert a call to function `nag_gpu_utilCheckMsg` (contained in the header file `nag_gpu.h`) following the call to `nag_gpu_mrg32k3a_exp` to check for the last CUDA error message, if any, generated during execution.

6 Example

This example prints 240 pseudorandom numbers from an exponential distribution with mean 1.0 generated by `nag_gpu_mrg32k3a_exp`. The generator is first initialized by calling `nag_gpu_mrg32k3a_init`.

6.1 Program Text

```

/* nag_gpu_mrg32k3a_uniform
 * single-precision example program to compute random numbers
 *
 * Copyright 2009, Numerical Algorithms Group Ltd, Oxford, UK.
 *
 * Version 0.2, 2009.
 *
 */

/////////////////////////////////////////////////////////////////
// pre-processor includes
/////////////////////////////////////////////////////////////////
#include <stdio.h>

/////////////////////////////////////////////////////////////////
// include libraries
/////////////////////////////////////////////////////////////////
#include <nag_gpu.h>

/////////////////////////////////////////////////////////////////
// define precision
/////////////////////////////////////////////////////////////////
#define PRS float

/////////////////////////////////////////////////////////////////
// Main program
/////////////////////////////////////////////////////////////////

int main(int argc, char **argv)
{
    FILE *fpout=stdout;

    // Generate batchSize numbers at a time on the GPU. For large sample
    // sizes this can avoid out-of-memory errors
    const int batchSize = 2621440;

    // Total number of batches to generate
    const int nBatches = 10;

    PRS
        // Host (PC) storage for batch-generated GPU random numbers
        *h_GpuBatchNums = 0,
        // Device (GPU) storage for batch-generated random numbers
        *d_GpuBatchNums = 0;

    // Number of columns to print
    const int printCols = 6;
    // Number of random numbers to print
    int printRows = 30;
    // Make sure we have enough numbers

```

```

printRows =
    (printRows*printCols > batchSize*nBatches
     ? batchSize*nBatches/printCols + 1
     : printRows);
// Storage for numbers we'll print
PRS *h_print = 0;
int
    // Current offset into h_print
    printIdx = 0,
    // Final print counter
    idx = 0,
    r = 0,
    c = 0;

unsigned int
    // Seed vector 1
    v1[] = {1, 2, 3},
    // Seed vector 2
    v2[] = {1, 2, 3};
unsigned long long
    // Offset into random number sequence
    offset = 0;

// Greeting
fprintf(fpout, "\n\n\tNAG GPU Mrg32k3a Example Program: "
        "\n\texponential numbers, %s precision\n\n",
        (sizeof(PRS)==sizeof(float) ? "SINGLE" : "DOUBLE"));

// Set device with highest Gflops/s
cudaSetDevice( nag_gpu_GetMaxGflopsDeviceId() );

// Allocate memory
if (!(h_GpuBatchNums = (PRS*)malloc(sizeof(PRS)*batchSize)) ) {
    fprintf(stderr, "h_GpuBatchNums allocation failure\n");
    goto END;
}
if (!(cudaSuccess == cudaMalloc((void **)&d_GpuBatchNums,
    sizeof(PRS)*batchSize))) {
    fprintf(stderr, "d_GpuBatchNums allocation failure\n");
    goto END;
}
if (!(h_print = (PRS*)malloc(sizeof(PRS)*printCols*printRows)) ) {
    fprintf(stderr, "h_print allocation failure\n");
    goto END;
}

for(int batchCount = 0; batchCount < nBatches; batchCount++) {
    // Due to the way mrg works, we have to initialise each time
    // in order to update the offset.
    nag_gpu_mrg32k3a_init(v1, v2, offset);
    // Launch GPU computation asynchronously. Feel free to experiment
    // with nb=#thread blocks, nt=#threads/block & np=#points/thread
    const int np = 1024;
    const int nt = 64;
    const int nb = batchSize/(np*nt);
    const int size = nb*nt*np;
    nag_gpu_mrg32k3a_exp(nb, nt, np, d_GpuBatchNums);
// Update the offset
    offset += size;

    /*
     * One can now launch other kernels to operate on the random
     * numbers, or copy the numbers to the host and operate on
     * them there. Here we simply copy them to the host in order
     * to print them
     */

    cudaMemcpy(h_GpuBatchNums, d_GpuBatchNums, sizeof(PRS)*size,
               cudaMemcpyDeviceToHost);

```

```

        // Copy GPU values to print storage
        for (int i = 0; i < size; i++) {
if (printIdx < printRows*printCols)
    h_print[printIdx++] = h_GpuBatchNums[i];
else break;
        } // END print storage loop

    } // END batching for loop

// Print random numbers
fprintf(fpout, "\nThe first %d GPU random numbers: "
        "read from left to right\n", printIdx);
for(r = 0; r < printRows; r++) {
    for(c = 0; c < printCols; c++) {
if (idx < printIdx) fprintf(fpout, "% .6f    ", h_print[idx++]);
else goto END;
        }
        fprintf(fpout, "\n");
    }
}

END:
fprintf(fpout, "\n\n");

// Release GPU and CPU memory
if (d_GpuBatchNums)
    nag_gpu_utilSafeCall( cudaFree(d_GpuBatchNums) );

if (h_GpuBatchNums)
    free(h_GpuBatchNums);

if (h_print)
    free(h_print);

fflush(stdout);
fflush(stderr);

if (fpout!=stdout)
    fclose(fpout);

cudaThreadExit();
return 0;
}

```

6.2 Program Data

None.

6.3 Program Results

NAG GPU Mrg32k3a Example Program:
exponential numbers, SINGLE precision

```

The first 180 GPU random numbers: read from left to right
7.786238    3.133761    0.390628    0.047322    0.619221    2.099508
1.452981    1.105778    1.488542    0.408835    0.254377    0.029574
0.232089    0.015527    0.984127    0.822564    0.871695    0.551717
0.033704    0.941827    0.558833    0.954600    3.461384    0.095214
1.816231    1.181292    2.115539    1.387164    3.292795    1.269703
1.401660    0.477105    0.020860    0.595189    4.028894    0.792625
0.314731    0.129319    0.335028    0.182206    1.681445    0.930176
0.164281    0.400494    2.604764    2.037144    0.260264    0.653435
0.443069    0.295203    0.006340    0.354065    0.392179    0.952175
3.693213    0.317748    0.588724    0.646826    0.028410    0.207162
0.377054    0.393141    0.472447    1.022450    0.074960    0.567766
0.825466    1.132510    0.131138    0.037198    0.596066    0.948409

```

0.234985	1.983983	1.040410	3.704617	1.737241	0.070950
0.123908	0.541282	0.888283	0.024281	0.530314	0.257832
0.119228	0.773380	2.853325	2.061939	0.035372	1.014673
0.729872	0.541850	0.808346	0.153797	0.539600	0.125768
1.828103	0.211178	0.083666	0.069153	0.813538	1.616621
0.574085	1.324631	1.002174	1.235641	0.709327	1.085632
0.758845	1.401383	0.001939	1.245684	0.212345	1.736751
0.224785	0.340602	0.989803	1.642448	0.977084	2.385972
0.106393	1.530689	0.592335	0.716975	0.146883	0.251216
0.186531	0.186935	0.615338	0.350338	3.111253	1.391711
0.163588	0.648762	2.791454	0.023737	0.571506	0.865380
0.017962	1.167064	1.666676	1.048909	1.179549	0.513803
0.897970	0.319500	0.066555	1.273219	1.824736	1.118284
0.036303	1.456271	0.184298	0.355326	0.059224	1.264546
0.521997	2.629050	0.796220	0.201281	0.406908	0.687556
0.003317	0.267058	3.389269	2.093461	1.551370	0.299624
0.232478	0.942925	0.975317	0.440281	0.415388	0.004879
1.070120	1.617652	0.692264	3.106889	0.664650	2.891359
