

Xilinx Sorting Application User Guide

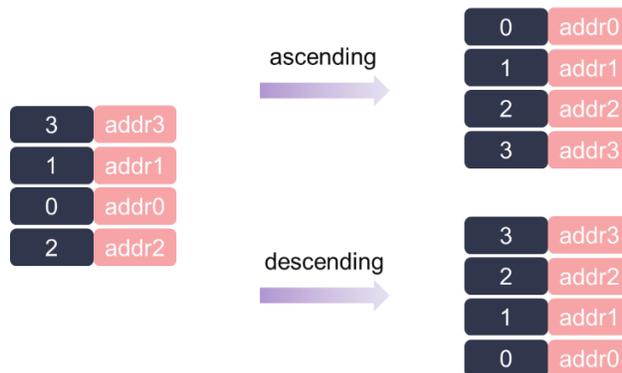
(Version 1.0)

1. Introduction

Xilinx Sorting Application is designed for a generalized (int32_t, int32_t) sorting solution, and is scalable for a large range. For large-scale data sorting, it can reach a high throughput.

Now the application can be run on both on-premises Alveo cards and in FPGA instances in the cloud, max supporting sort task with 1G 32-bit integer pairs.

2. Application Usage



The application is containerized and can be easily run in a few minutes in the Nimbox cloud or on premises. Details can be found at link <https://www.xilinx.com/products/acceleration-solutions/integer-sorting.html>.

2.1 Sample

Sample 1. Try the default demo (1M integer pairs sorting)

```
/opt/xilinx/apps/vt_database/sort/bin/vt_database_sort.exe --demo
```

By the end of the license message print, it will ask for user's input "yes/no" to acknowledge the agreement. Type "yes" to acknowledge the licence agreement. Parameter "--accept-EULA" can be used to bypass the license message.

Note: when run in a NUMA machine, suggest installing and using "numactl" command to make performance stable.

When the execution is done, it will print results shown in Figure 1.

```

Demo1: Single Mode Args: --in /home/nimbix/demo_data/input_1M_0.dat --out /home/nimbix/
demo_data/single_mode_out.dat
Single mode: input data loading...!
Single mode: input data load done!
Kernel has been created
Single mode: data sort done!
End-to-End time: 28.196 msec
Single mode: output data save done!
Single mode: start validating result...
Single mode, Status: Pass

Demo2: Batch Mode Args: --files_in /home/nimbix/demo_data/input.txt --files_out /home/n
imbix/demo_data/output.txt
Batch mode: input data loading...!
Batch mode: /home/nimbix/demo_data/input_1M_0.dat load done!
Batch mode: /home/nimbix/demo_data/input_1M_1.dat load done!
Batch mode: input data load done, total 2 files.
Kernel has been created
Kernel has been created
Batch mode: data sort done!
End-to-End time: 42.990 msec
Batch mode: /home/nimbix/demo_data/output_1M_0.dat saved, start validate...
Batch mode: /home/nimbix/demo_data/output_1M_0.dat Status: Pass.
Batch mode: /home/nimbix/demo_data/output_1M_1.dat saved, start validate...
Batch mode: /home/nimbix/demo_data/output_1M_1.dat Status: Pass.
Batch mode: output files saved, total: 2 files.

```

Figure 1. Output for demo

The demo case contains a single mode test and a batch mode test.

When single mode test passes, it will print **“Single mode, Status: Pass”**, else it prints **“Single mode, Status: Error”**

In batch mode, for each testing file, it will print **“Batch mode, file_name, Status: Pass”** when test passes, else it prints **“Batch mode, file_name, Status: Error”**

Sample 2. Generate datasets with larger size and run test cases

generate datasets

For example, generate two datasets with 64M integer pairs and run tests

```
/opt/xilinx/apps/vt_database/sort/bin/gendata.exe -ss 64 -ln 2 -out /home/nimbix/demo_data
```

Then we get “input_64M_0.dat” and “input_64M_1.dat” in directory /home/nimbix/demo_data

run sorting application (single mode)

```
/opt/xilinx/apps/vt_database/sort/bin/vt_database_sort.exe --accept-EULA
```

```
--in /home/nimbix/demo_data/input_64M_0.dat
```

```
--out /home/nimbix/demo_data/output_64M_0.dat
```

Final output is “Single mode, Status: Pass”

run sorting application (batch mode)

1) Add two lines in file /home/nimbix/demo_data/input_64.txt as below

```
/home/nimbix/demo_data/input_64M_0.dat
```

```
/home/nimbix/demo_data/input_64M_1.dat
```

2) Add two lines in file /home/nimbix/demo_data/output_64.txt as below

```
/home/nimbix/demo_data/output_64M_0.dat
```

```
/home/nimbix/demo_data/output_64M_1.dat
```

```
3) /opt/xilinx/apps/vt_database/sort/bin/vt_database_sort.exe --accept-EULA
```

```
--files-in /home/nimbix/demo_data/input_64.txt
```

```
--files-out /home/nimbix/demo_data/output_64.txt
```

Final output will print “**Status: Pass**” for each output.

2.2 Prerequisites

2.2.1 Device and Software

This application supports Xilinx FPGA Alveo U200 card at this moment. To run this application on users' machines, please make sure:

- Xilinx FPGA Alveo U200 (shell xilinx_u200_xdma_201830_2) card is installed correctly.
- XRT 2020.1

2.2.2 Datasets

The application provides 1M (key, value) integer pairs inside the docker images.

Besides, the application provides user a function to generate more large-scale datasets:

Each key is generated randomly, and value = key + 1.

Usage:

```
/opt/xilinx/apps/vt_database/sort/bin/gendata.exe -ss dataset-length(million) -ln dataset-number -out  
output-dataset-location
```

e.g.

```
/opt/xilinx/apps/vt_database/sort/bin/gendata.exe -ss 32 -ln 2 -out /home/nimbix/demo_data/
```

Run command above, it will generate two datasets, input_32M_0.dat and input_32M_1.dat in output directory /home/nimbix/demo_data/, each dataset has a length of 32 million (integer pairs).

2.3 Run Application

Try the default demo (1M integer pairs sorting)

```
/opt/xilinx/apps/vt_database/sort/bin/vt_database_sort.exe --demo
```

The application provides two task modes, single and batch mode.

In single mode, input is a single unsorted file (application command is “-i/--in”), and output is corresponding sorted file (application command is “-o/--out”)

In batch mode, users input input.txt (-I/--files-in) and output.txt (-O/--files-out). In input.txt, each line means an absolute path for unsorted file, and in output.txt, each line means sorted file path for the same location in input.txt.

Table 1. lists the entries user can use in this application.

Table 1. Command list

Command	Default value	Function
--accept-EULA	False	acknowledge the license agreement and skip printing license acknowledgement. If false, it will print out license file to user console and ask for user's input "yes/no" to acknowledge the agreement
--demo	False	If true, run demo case
-l --files-in	""	Batch mode, a text file contains all unsorted file paths
-O --files-out	""	Batch mode, a text file contains all sorted file paths
-i --in	""	Single mode, an unsorted data file
-o --out	""	Single mode, a sorted data file
-a --asc	1, ascending	Sorting order (ascending-1/descending-0)

Performance Spec

Latency (test in single mode)

round	1M	4M	16M	64M	128M	256M	512M	1024
1	32.18	138.293	597.332	2446.029	4788.504	8753.532	12682.47	23298.87
2	33.065	141.891	597.666	2437.177	4904.821	7579.786	12573.5	23285.17
3	36.218	145.382	601.243	2667.793	4879.099	7456.273	12584.36	23301.12
4	36.473	143.098	597.487	2306.118	4806.695	7427.101	12580.81	23336.42
5	36.628	146.426	602.998	2435.3	4878.045	7436.795	12568.02	23309.08
6	36.34	144.05	605.225	2305.119	4885.258	7365.113	12541.97	23229.87
7	36.396	142.041	603.001	2436.571	4790.51	7458.068	12504.98	23169.8
8	32.545	143.804	612.948	2449.066	4812.213	7399.806	12580.62	23038.9
9	36.404	140.39	601.488	2437.42	4753.544	7407.933	12622.59	23163.97
10	36.326	147.223	588.894	2300.761	4813.385	7399.927	12585.63	23190.07
latency/ms (one run)	35.2575	143.2598	600.8282	2422.135	4831.207	7568.433	12582.5	23232.33

Throughput (test in batch mode)

	1M	4M	16M	64M	128M	256M	512M	1024M
time(ms) 20runs for 1~128M 10runs for 256M~512M 5runs for 1024M	489.812	1336.099	4390.317	14338.22	35972.24	36355.21	82627.12	94696.92
	480.35	1297.498	4320.751	14449.2	36046.53	36540.6	81458.88	88895.84
	493.52	1269.302	4322.465	13975.11	36230.26	36670.4	81225.19	79275.61
	457.474	1272.13	4275.223	14298.62	35865.99	36419.86	80940.29	86440.87
	487.427	1313.467	4193.212	14081.12	36069.27	36454.03	80043.09	90345.19
	429.112	1289.216	4434.132	14312.84	35523.85	36700.49	69104.45	88323.86
	456.023	1318.694	4505.401	14091.23	35920.08	41279.94	81031.41	87281.27
	420.751	1330.806	4230.702	14529.21	35600.43	36264.85	81570.16	86698.76
	497.653	1286.847	4322.115	14190.82	35842.97	36220.46	80204.5	88758.06
	412.688	1363.947	4799.562	14614.98	35784.57	36319.21	73502.73	86927.95
462.481	1307.801	4379.388	14288.14	35885.62	36922.5	79170.78	87764.43	
E2E time (per run)	23.12405	65.39003	218.9694	714.4069	1794.281	3692.25	7917.078	17552.89
Throughput (MB/s)	345.9602	489.3712	584.5566	716.6785	570.7021	554.6753	517.3626	466.7039