

TR-2015-04

Investigation of Charm++ SPMV performance using CSR format

Dominik Heller

March 18, 2015

Contents

| | | |
|----------|-----------------------|----------|
| 1 | Introduction | 2 |
| 2 | Simulation | 2 |
| 2.1 | Environment | 2 |
| 2.2 | Results | 3 |

1 Introduction

CHARM++ [1] is a portable parallel programming platform running on single-core machines or large clusters alike. The programming environment builds upon C++ with additional library calls and an interface description syntax for publishing Charm++ objects. It utilizes asynchronous message passing for communication and promises many features such as latency hiding by overlaying communication with calculation, effective dynamic load balancing and fault tolerance, while offering a high-level OOP-style approach to parallel programming. In this report the performance of Charm++ for CSR-format sparse matrix-vector multiplication SPMV is investigated for a collection of test matrices.

2 Simulation

2.1 Environment

Hardware. All tests were run on each of the following nodes:

- a) 2 socket, 16-core Intel Xeon E5520 (Nehalem 2009) node
- b) 4 socket, 64-core AMD Opteron 6274 (Interlagos 2011) node
- c) 2 socket, 40-core Intel Xeon E5-2690 v2 (Ivy Bridge 2013) node

Setups b) and c) are described in more detail in [5].

Software. Benchmarks are based on a custom SPMV code using Charm++ version 6.6.1. The code is available to UW-SBEL members on Github [2]. Charm++ was compiled using gcc version 4.4.6 with `--with-production` statement. The SPMV benchmark code itself was compiled with `-O3 -optimize -production` options.

Despite having 16, 40 or 64 cores at disposal, the program was run on only 8, 20 or 32 cores, respectively, using the `+pN` option for the `charmrun` environment to deploy the program on N cores. Tests have shown that running the program on all 16, 40 or 64 physical and virtual cores in fact decreases performance significantly in all configurations, as additional overhead is introduced without having access to more floating-point resources on the CPU.

Input data. Results were generated for 129 different test matrices as described in [3], mostly originating from the University of Florida Matrix Collection [4] with some additions originating from multibody dynamics problems from the Simulation Based Engineering Lab SBEL. Problem dimensions vary between 7,000 (`eurqsa`, `fp`) and 4,700,000 (`rajat31`) with up to 24 million nonzero entries (`ldoor`).

Chare setup. Generally, Chares were created to hold the highest possible number of 1024, 2048, 4096 or 8192 nonzero entries each, depending on the memory requirements of a specific problem. In cases where this would lead to the number of chares n_{ch} not being much larger than the number of cores n_{cpu} ; $n_{ch} \leq 10 n_{cpu}$, the number of chares was increased up to the next integer multiple of the number of cores. For instance, 96 chares would be created when running with 32 cores, even if data would fit into a lower number of e.g. 87 chares.

2.2 Results

See Tables 1-3 for input file dimensions and MFLOPS on the different nodes. Results are plotted in Figure 1 for AMD Opteron 6274, Figure 2 for Intel Xeon E5520 and Figure 3 for Intel Xeon E5-2690 v2.

Unsurprisingly, the much newer E5-2690 v2 chip significantly outperforms its competitors with a peak flop rate of 18.5 GFLOPS, an average (arithmetic mean) of 7.6 GFLOPS and a worst-case performance of 1.3 GFLOPS.

The AMD Opteron 6274 node achieves between 0.27 and 5.5 GFLOPS with an average of 2.20 GFLOPS. Intel's Xeon E5520 chip has about the same average performance with 2.16 GFLOPS, peaks at 4.8 GFLOPS and reaches a better worst-case rate of 0.74 GFLOPS.

A direct comparison of the E5520 and 6274 nodes is shown in Figure 4, where the ratio of flop rates is plotted. Despite the slightly higher average performance of the Opteron node, values are greater than 1 in most cases (81 out of 129), indicating that the Xeon node outperforms the Opteron node.

The Xeon node is up to 4.15 times as fast as Opteron (*eurqsa*, *dw8192*), while in some cases the Opteron node beats Xeon, being up to 3 times as fast (*rma10*).

References

- [1] Charm++ documentation. <http://charm.cs.illinois.edu/help>.
- [2] Github UW-SBEL. <https://github.com/uwsbel/benchmarks/tree/master/spmv/charm/code>.
- [3] R. Serban A. Li, H. Mazhar and D. Negrut. TR-2015-02 Comparison of SPMV performance on matrices with different matrix format using CUSP, cuSPARSE and ViennaCL. Technical report, Simulation Based Engineering Lab, University of Wisconsin-Madison, 2015.
- [4] T.A. Davis and Y. Hu. The University of Florida Sparse Matrix Collection. *ACM Trans. Math. Softw.*, 38(1):1:1–1:25, 2011.
- [5] H. Mazhar and D. Negrut. TR-2015-01 Comparison of OpenCL performance on different platforms using VexCL and Blaze. Technical report, Simulation Based Engineering Lab, University of Wisconsin-Madison, 2015.




Figure 1: GFLOPS for AMD Opteron 6274 node.




Figure 2: GFLOPS for Intel Xeon E5520 node.




Figure 3: GFLOPS for Intel Xeon E5-2690 v2 node.




Figure 4: GFLOPS ratio Xeon E5520 over Opteron 6274. Values > 1 indicate the Xeon node outperforming Opteron.

| File | N | nnz | 6274 | E5520 | 2690v2 |
|-----------------|----------|------------|-------------|--------------|---------------|
| 2cubes_sphere | 101492 | 874378 | 1864.8 | 1934.9 | 6580.9 |
| 2D_54019_highK | 54019 | 996414 | 3355.5 | 1877.1 | 11563.3 |
| 31770-lhs | 31770 | 183540 | 1286.9 | 2248.7 | 4248.3 |
| 88950-lhs | 88950 | 513900 | 1585.1 | 2287.8 | 5329.3 |
| a2nnsnsl | 80016 | 196115 | 1667.0 | 2792.2 | 5896.2 |
| a5esindl | 60008 | 145004 | 2020.9 | 2668.7 | 4918.9 |
| ABACUS_shell_ud | 23412 | 218484 | 1800.8 | 2865.0 | 5851.0 |
| af_5_k101 | 503625 | 9027150 | 3148.9 | 2525.2 | 4431.8 |
| af23560 | 23560 | 484256 | 3146.3 | 4485.4 | 10908.6 |
| ancfBigDan | 63603 | 569262 | 1996.8 | 2625.4 | 7369.9 |
| Andrews | 60000 | 410077 | 1377.8 | 1658.1 | 4093.6 |
| apache1 | 80800 | 311492 | 986.4 | 1792.9 | 3588.3 |
| apache2 | 715176 | 2766523 | 831.0 | 939.3 | 2418.6 |
| appu | 14000 | 1853104 | 5523.9 | 2323.7 | 14153.8 |
| ASIC_100k | 99340 | 954163 | 2139.7 | 2175.9 | 7674.6 |
| ASIC_100ks | 99190 | 578890 | 1468.2 | 2069.9 | 5124.8 |
| av41092 | 41092 | 1683902 | 4631.8 | 1831.0 | 14520.7 |
| bauru5727 | 40366 | 145019 | 835.7 | 1446.3 | 2857.7 |
| bayer01 | 57735 | 277774 | 1137.4 | 1941.1 | 4229.0 |
| bcircuit | 68902 | 375558 | 1302.0 | 1939.9 | 4684.9 |
| bcstk39 | 46772 | 1068033 | 3428.1 | 2970.6 | 12856.3 |
| blockqp1 | 60012 | 340022 | 1338.8 | 1860.5 | 4498.9 |
| bmw3_2 | 227362 | 5757996 | 4135.6 | 2607.6 | 5080.3 |
| bmwera_1 | 148770 | 5396386 | 1539.0 | 2460.5 | 6220.8 |
| boyd1 | 93279 | 652246 | 1805.6 | 1774.5 | 4762.9 |
| bratu3d | 27792 | 88627 | 1630.8 | 2653.9 | 4350.3 |
| bundle1 | 10581 | 390741 | 2739.7 | 4718.4 | 12031.4 |
| c-59 | 41282 | 260909 | 1570.4 | 1892.7 | 4142.3 |
| c-61 | 43618 | 176817 | 992.1 | 1454.4 | 2984.9 |
| c-62 | 41731 | 300537 | 1569.6 | 1945.8 | 4631.6 |
| cant | 62451 | 2034917 | 4617.4 | 2049.8 | 15353.8 |
| case39 | 40216 | 526139 | 2614.8 | 3999.9 | 12886.8 |
| case39_A_01 | 40216 | 526139 | 3015.6 | 4066.6 | 13406.0 |
| c-big | 345241 | 1343126 | 1077.8 | 934.9 | 3444.4 |
| cf1 | 70656 | 949510 | 2681.5 | 1846.3 | 8544.0 |
| cf2 | 123440 | 1605669 | 2698.1 | 1448.7 | 10047.7 |
| circuit_4 | 80209 | 307604 | 991.5 | 1605.3 | 3265.2 |
| ckt11752_tr_0 | 49702 | 333029 | 1505.1 | 2346.9 | 5500.4 |
| cond-mat-2005 | 40421 | 175693 | 2759.8 | 3501.5 | 7736.9 |
| cont-201 | 80595 | 239596 | 901.1 | 1263.3 | 2674.3 |
| cont-300 | 180895 | 539396 | 858.3 | 1146.3 | 2846.2 |
| copter2 | 55476 | 407714 | 1502.0 | 2365.3 | 5581.3 |
| crankseg_1 | 52804 | 5333507 | 1700.8 | 2480.6 | 7832.8 |

Table 1: Input file, rows **N**, nonzeros **nnz** and MFLOPS per node - 1/3

| File | N | nnz | 6274 | E5520 | 2690v2 |
|-----------------|----------|------------|-------------|--------------|---------------|
| d_pretok | 182730 | 885416 | 1312.9 | 1188.5 | 4508.8 |
| dawson5 | 51537 | 531157 | 2050.6 | 2971.3 | 7901.9 |
| dc1 | 116835 | 766396 | 1779.5 | 1419.3 | 6092.1 |
| dixmaanl | 60000 | 179999 | 818.3 | 1410.6 | 2533.3 |
| Dubcova2 | 65025 | 547625 | 1779.2 | 2277.5 | 6315.4 |
| dw8192 | 8192 | 41746 | 421.6 | 1736.3 | 1910.6 |
| ecl32 | 51993 | 380415 | 1655.9 | 2554.1 | 5804.2 |
| epb3 | 84617 | 463625 | 1444.0 | 2173.3 | 5098.9 |
| eurqsa | 7245 | 25633 | 271.6 | 1128.2 | 1289.0 |
| ex11 | 16614 | 1096948 | 4736.3 | 3420.8 | 17604.1 |
| ex19 | 12005 | 259879 | 2863.1 | 3761.3 | 8471.5 |
| F2 | 71505 | 2682895 | 1615.2 | 1689.8 | 15016.4 |
| FEM_3D_thermal1 | 17880 | 430740 | 3063.5 | 4242.5 | 11217.0 |
| filter3D | 106437 | 1406808 | 2593.2 | 1858.1 | 9032.6 |
| finan512 | 74752 | 335872 | 1148.3 | 1945.9 | 3944.8 |
| fp | 7548 | 848553 | 4408.7 | 2388.3 | 16812.1 |
| G3_circuit | 1585478 | 4623152 | 661.8 | 738.9 | 1377.2 |
| g7jac140 | 41490 | 565956 | 2433.4 | 3400.1 | 8999.2 |
| Ga3As3H12 | 61349 | 3016148 | 1654.1 | 2198.8 | 16231.7 |
| GaAsH6 | 61349 | 1721579 | 4070.5 | 2211.2 | 13456.1 |
| garon2 | 13535 | 390607 | 3028.2 | 4790.7 | 11086.2 |
| gas_sensor | 66917 | 885141 | 2645.8 | 1914.1 | 9529.1 |
| gridgena | 48962 | 280523 | 1246.2 | 2310.1 | 4593.3 |
| gsm_106857 | 589446 | 11174185 | 1386.4 | 1449.1 | 2808.5 |
| H2O | 67024 | 1141880 | 3069.4 | 2767.5 | 11193.8 |
| hcircuit | 105676 | 513072 | 1311.0 | 1959.3 | 4372.0 |
| HTC_336_4438 | 226340 | 565431 | 742.8 | 902.9 | 2524.5 |
| ibm_matrix_2 | 51448 | 1056610 | 3263.1 | 3519.6 | 12562.9 |
| inline_1 | 503712 | 18660027 | 4342.8 | 2479.4 | 5611.9 |
| jan99jac120 | 41374 | 260202 | 1465.6 | 2311.9 | 4573.7 |
| juba40k | 40337 | 144945 | 852.6 | 1419.0 | 2879.5 |
| k1_san | 67759 | 303364 | 1101.3 | 1778.0 | 4006.0 |
| laminar_duct3D | 67173 | 3833077 | 3039.6 | 2432.3 | 14137.4 |
| ldoor | 952203 | 23737339 | 2627.6 | 1865.6 | 3848.8 |
| lhr10c | 10672 | 232633 | 2691.1 | 3946.0 | 7972.9 |
| lhr71 | 70304 | 1528092 | 3805.5 | 1649.7 | 13353.9 |
| Lin | 256000 | 1011200 | 1121.8 | 1084.6 | 4127.8 |
| lung2 | 109460 | 492564 | 1227.9 | 1954.4 | 4385.7 |
| mario002 | 389874 | 1167685 | 853.2 | 904.1 | 3034.3 |
| mark3jac100 | 45769 | 285215 | 1297.6 | 2093.9 | 4973.9 |
| mark3jac140 | 64089 | 399735 | 1449.0 | 2058.9 | 5416.4 |
| matrix_9 | 103430 | 2121550 | 3535.7 | 2333.7 | 13503.5 |
| minsurfo | 40806 | 122214 | 721.5 | 1420.7 | 2462.5 |

Table 2: Input file, rows **N**, nonzeroes **nnz** and MFLOPS per node - 2/3

| File | N | nnz | 6274 | E5520 | 2690v2 |
|------------------|----------|------------|-------------|--------------|---------------|
| msc23052 | 23052 | 588933 | 3104.9 | 4004.7 | 11509.4 |
| ncvxbqp1 | 50000 | 199984 | 1056.9 | 1621.5 | 3323.8 |
| nd24k | 72000 | 14393817 | 2762.9 | 3503.0 | 5912.1 |
| offshore | 259789 | 2251231 | 1993.3 | 1157.8 | 6925.4 |
| oilpan | 73752 | 1835470 | 3687.0 | 1710.1 | 14430.3 |
| olesnik0 | 88263 | 402623 | 1205.8 | 1961.7 | 4313.1 |
| OPF_10000 | 43887 | 255799 | 1330.6 | 2023.7 | 4443.1 |
| parabolic_fem | 525825 | 2100225 | 1032.7 | 868.3 | 3261.3 |
| pdb1HYS | 36417 | 2190591 | 5204.8 | 2626.2 | 18513.7 |
| poisson3Db | 85623 | 2374949 | 3962.5 | 2087.4 | 9608.5 |
| pwtk | 217918 | 5926171 | 3011.7 | 2074.4 | 6564.6 |
| qa8fk | 66127 | 863353 | 2622.4 | 1935.3 | 9464.5 |
| qa8fm | 66127 | 863353 | 2587.4 | 1814.1 | 9600.9 |
| raefsky4 | 19779 | 674195 | 3729.7 | 1908.5 | 12803.2 |
| rail_79841 | 79841 | 316881 | 1026.9 | 1732.5 | 3649.9 |
| rajat26 | 51032 | 249302 | 1223.0 | 1857.9 | 4077.4 |
| rajat30 | 643994 | 6175377 | 1736.0 | 1610.3 | 3198.9 |
| rajat31 | 4690002 | 20316253 | 901.3 | 1021.6 | 1479.4 |
| rma10 | 46835 | 2374001 | 4994.8 | 1707.0 | 15685.8 |
| s3dkq4m2 | 90449 | 2455670 | 4164.4 | 1605.8 | 15643.1 |
| shallow_water1 | 81920 | 204800 | 703.1 | 1132.5 | 2330.8 |
| shallow_water2 | 81920 | 204800 | 704.9 | 1126.6 | 2412.8 |
| ship_003 | 121728 | 4103881 | 4713.7 | 2254.5 | 12519.3 |
| shipsec1 | 140874 | 3977139 | 2965.6 | 2334.1 | 11945.8 |
| shipsec5 | 179860 | 5146478 | 1528.7 | 2267.8 | 7394.4 |
| Si34H36 | 97569 | 2626974 | 1546.3 | 2137.8 | 11234.2 |
| SiO2 | 155331 | 5719417 | 1500.2 | 2506.9 | 7178.4 |
| sparsine | 50000 | 799494 | 2273.0 | 2471.2 | 7580.2 |
| stokes128 | 49666 | 295938 | 1293.9 | 2402.6 | 4941.1 |
| stomach | 213360 | 3021648 | 1498.3 | 1918.2 | 10974.7 |
| t3dh | 79171 | 2215638 | 4117.4 | 1987.5 | 14686.8 |
| t3dh_a | 79171 | 2215638 | 3955.4 | 2288.0 | 14053.7 |
| thermal1 | 82654 | 328556 | 1002.7 | 1518.8 | 3486.0 |
| thermal2 | 1228045 | 4904179 | 772.2 | 738.9 | 1858.3 |
| torso3 | 259156 | 4429042 | 1509.1 | 1876.2 | 5540.3 |
| TSOPF_FS_b162_c4 | 40798 | 1204322 | 4526.1 | 2155.0 | 16048.0 |
| TSOPF_FS_b39_c19 | 76216 | 998359 | 4156.1 | 2155.3 | 13764.0 |
| turon_m | 189924 | 912345 | 1303.6 | 1160.4 | 4763.4 |
| vanbody | 47072 | 1191985 | 3712.0 | 3328.7 | 13219.2 |
| venkat25 | 62424 | 1717792 | 4130.1 | 1990.5 | 14348.1 |
| vibrobox | 12328 | 177578 | 1873.4 | 3024.7 | 5717.7 |
| xenon1 | 48600 | 1181120 | 3755.4 | 2829.1 | 13355.5 |
| xenon2 | 157464 | 3866688 | 1511.0 | 1850.6 | 9646.8 |

Table 3: Input file, rows **N**, nonzeros **nnz** and MFLOPS per node - 3/3