UDK 519.688

THE MOBILE LINPACK: INITIAL EXPERIENCE OF THE PERFORMANCE RANKING FOR PORTABLE DEVICES

O. O. Iachnik [1], D. A. Nikitenko [2], and S. I. Sobolev [3]

Contemporary computing is strongly associated with parallelism. From personal tablets and phones and up to the top supercomputers, the scales and approaches differ, but the idea of work distribution is present everywhere. A study of portable platforms parallel computing capabilities was started at the Research Computing Center of Lomonosov Moscow State University several years ago, in 2014. There were Android and iOS implementations of popular benchmark developed. Four years of the test availability for regular users turned into almost 20 thousand of test results. This paper addresses the investigation of facts and trends, mined out of this data, proving the reasons for the further Mobile Linpack development.

Keywords: portable device performance analysis, mobile platform benchmark, mobile device performance, Android, iOS, mobile market share.

1. Introduction. The Linpack performance test based on the solution of a system of linear algebraic equations [1] is widely known for its extensive history of application in the field of supercomputer systems. So, on its principle basis, since 1993, the World Top500 Ranking has been formed [2], since 2004 — the Top50 Ranking of the most productive systems in Russia and the Commonwealth of Independent States (CIS) [3], and a number of other regional ratings [4, 5]. Despite all the shortcomings currently observed, first of all, a weak compliance with the algorithms inherent in most real applications, the results of the test expressed in floating point operations per second (Flop/s) remain the generally accepted measurement standard for the computer systems vendors, which is the main advantage of the test at the moment.

2. The Mobile Linpack. Modern mobile platforms inherit the same principles of parallel data processing data that are used in the construction of high-performance systems: multicore and multithreading, the use of specialized accelerators, and so on. In 2014, the team of authors proposed the first implementation of the Linpack test for mobile platforms in order to determine the trends and features of this segment with respect to the performance and efficiency in the same units as for traditional computing platforms [6]. At the first stage, a parallel version for the Android OS was implemented. For the second most common mobile operation system — iOS, just a basic sequential implementation was developed. Since 2015, mobile applications are available in the application stores of the respective platforms. For these 4 years of the project existence, a rich database with results has been accumulated with more than 18 000 entries, each of which corresponds to the successful run of the Linpack test on some mobile device.

This paper discovers the results of the analysis of accumulated data presented to determine the main trends in the segment of mobile devices as well as to determine the prospects for the development of the rating itself. The initial data is under the control of the MySQL database management system, the analysis is based on the formation of appropriate SQL queries and their subsequent processing for convenience of perception — the choice of diagrams, etc.

At the first stage of the analysis, an assessment was made concerning the suitability of the data for constructing the trials, during which some insignificant artifacts were identified that did not affect the feasibility of the analysis, such as some inaccuracies in the versions of the OS used. This did not significantly decrease the dataset size, because there had been just around several hundred of entities excluded. The analysis of the cleaned accumulated data was carried out in a number of directions and showed the following.

3. Number and regularity of rest runs. Since the first presentation of the rating at the largest in Europe conference series ISC (International Supercomputing Conference) in 2015, a fairly constant rate of

[1] Lomonosov Moscow State University, Faculty of Computational Mathematics and Cybernetics; Leninskie Gory, Moscow, 119991, Russia; Student, e-mail: olegator0604@gmail.com

[2] Lomonosov Moscow State University, Research Computing Center; Leninskie Gory, Moscow, 119992, Russia; Ph.D., Senior Scientist, e-mail: dan@parallel.ru

[3] Lomonosov Moscow State University, Research Computing Center; Leninskie Gory, Moscow, 119992, Russia; Ph.D., Senior Scientist, e-mail: sergeys@parallel.ru

launch has been established — from 400 to 600 per month without obvious failures (Fig. 1). There are periods of increased activity corresponding to the periods of supercomputer conferences and events, in which the audience was invited to participate in the ranking.

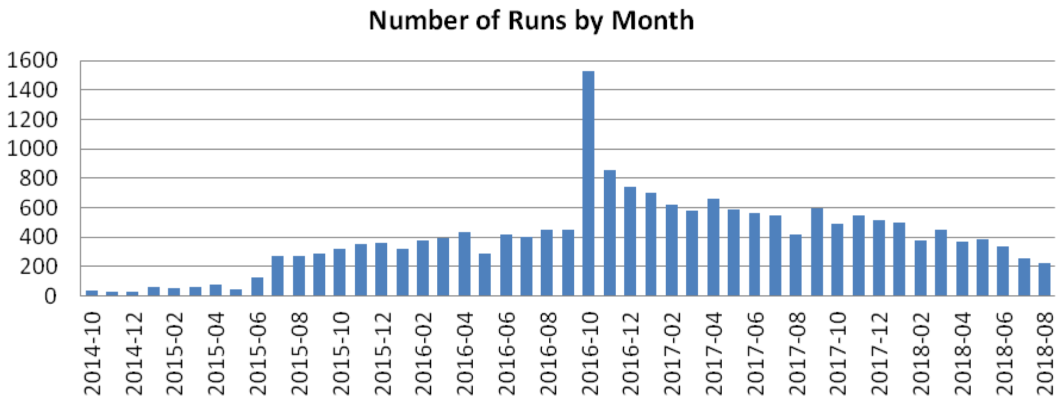

Fig. 1. Number of Runs by Month

These results allow us to continue the data study at the level of a monthly and less frequency.

4. Mobile platforms. The dominating majority of runs are made on Android platforms (Fig. 2). In many respects this is determined by the availability of a parallel Android implementation and just by a serial implementation for iOS. At the same time, this also corresponds to the hypothesis that the target advanced audience prefers more flexible devices.

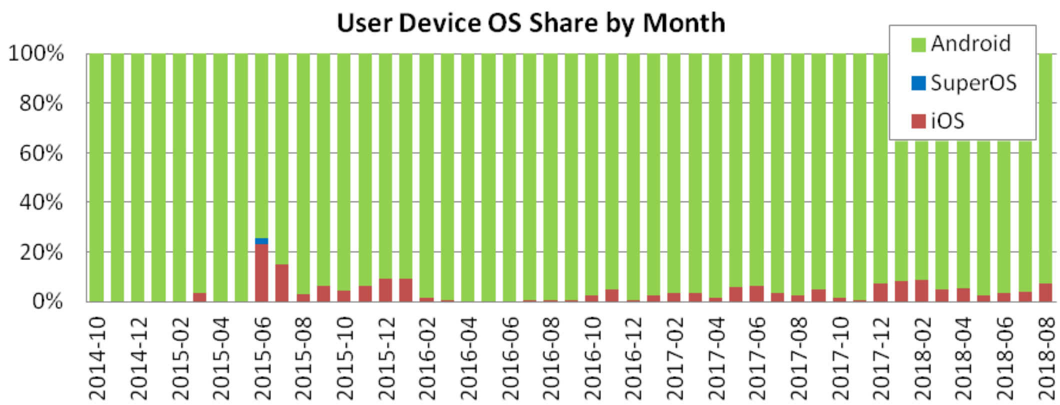

Fig. 2. User Device OS Share by Month

The results clearly show that we should exclude iOS devices from the further statistics analysis until we have a reasonable parallel implementation.

5. Performance growth. The rate of performance growth itself is a characteristic for the entire computing area, showing almost a triple increase every two years. By now, the average device performance has increased eleven times since the project start in late 2014 (Fig. 3).

Attention should be paid to the low value of the worst results showed by the lowest line. This can be explained by a significant number of launches with no run parameters optimization, accordingly, with no device features unfolding. Affirming this, we can also see the median values below averages.

These results demonstrate the necessity to develop a tool that would recommend run options for each device, or at least the user should be explicitly invited to try various options to achieve good performance.

6. Cores and memory. Just as in the whole computing, there is an increase in the number of processor cores in the world of portables. Top devices have about 8 cores per CPU at present. Of course, computational tasks grow as well, which requires an increase of the amount of memory, too.

But what is actually the most interesting revealed fact, the performance per core ratio increases a few times slower than the bytes per core ratio (Fig. 4). That is just opposite the trend in big computing. The reason for that is actually the focus on local data processing and strong power limitations for the mobiles. The portables do not need so much memory for typical applications, even games, but are very sensitive to energy and pricing issues. Ignoring such limitations would just throw the device away from the mass market.

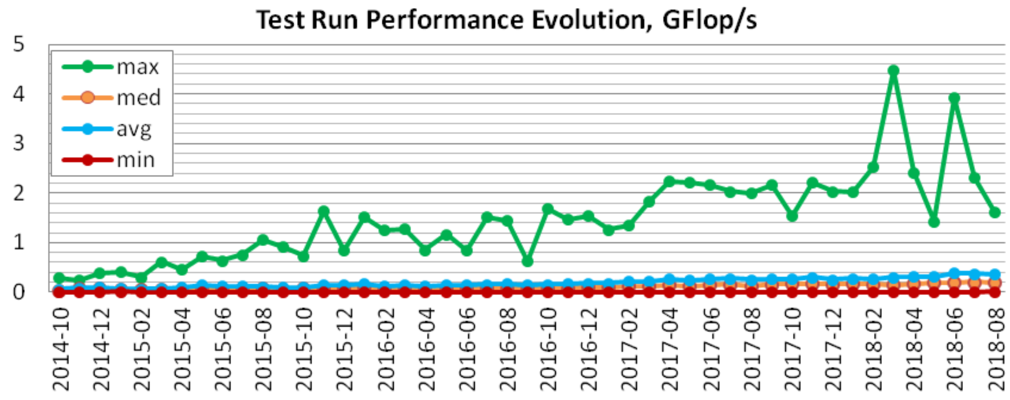

Fig. 3. Test Run Performance Evolution, GFlop/s

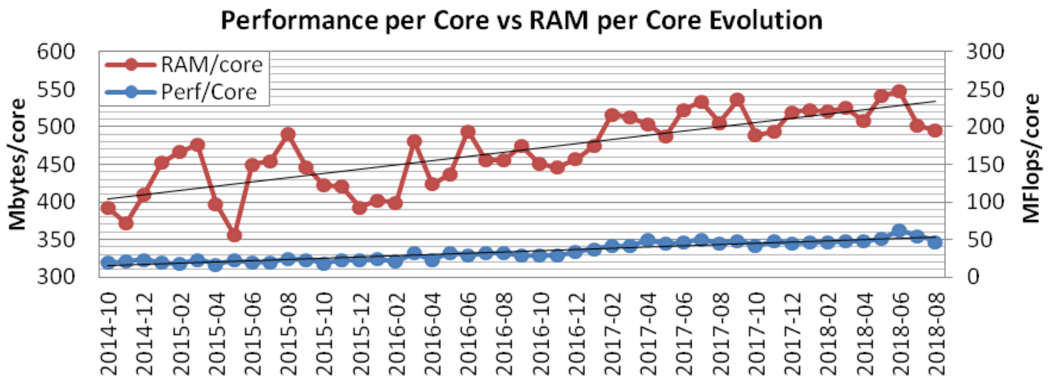

Fig. 4. Performance per Core vs RAM per core Evolution

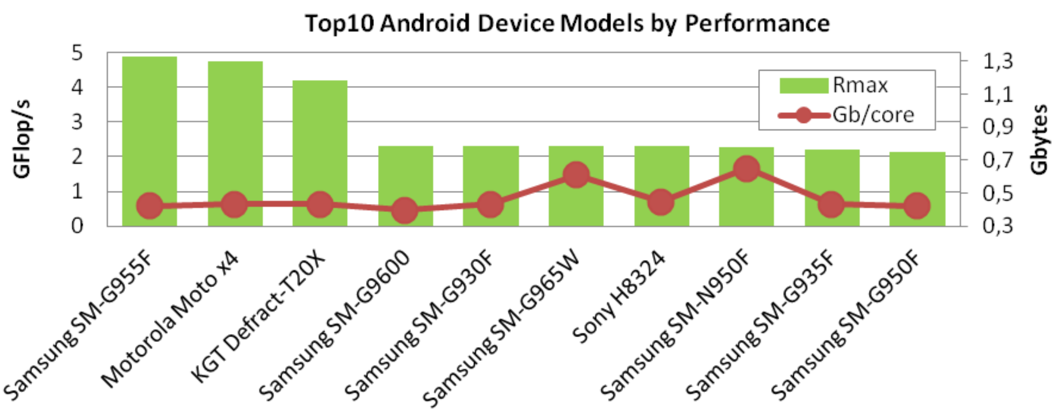

Fig. 5. Top10 Android Device Models by Performance

7. Top devices by performance. One can see a variety of Samsung devices demonstrating the performance that is enough to get into the top ten (Fig. 5), showing from 2.1 and up to 4.9 GFlop/s that is just 12 times slower then Top500 #1 system in the first world ranking (1993, Los Alamos SuperSparc-based CM-5/1024, having 59.7 GFlop/s). This proves a good idea to enrich the Mobile Linpack results interface with examples of historic and well-known Top500 systems for comparison [7]. With this, any user can feel having a past supercomputer in his present pocket. This is also a good illustration of the performance growth for the public.

There are just three of ten devices of other manufacturers in the top — Motorola, KGT and Sony. It is surprising to see no Xiaomi, Huawei or other top vendors among these ten.

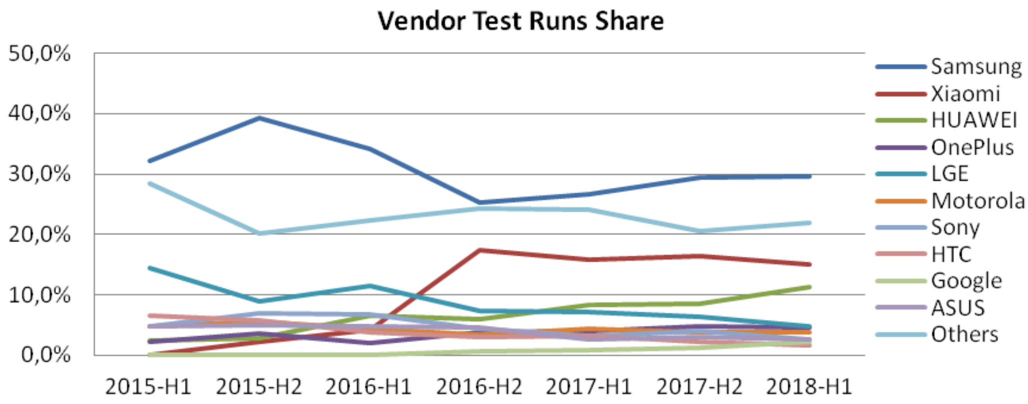

Fig. 6. Vendor Performance Share

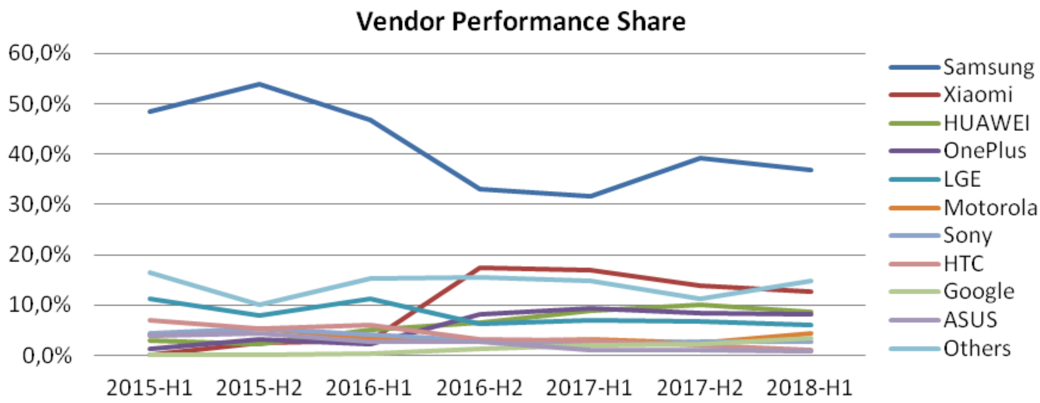

Fig. 7. Vendor Test Runs Share

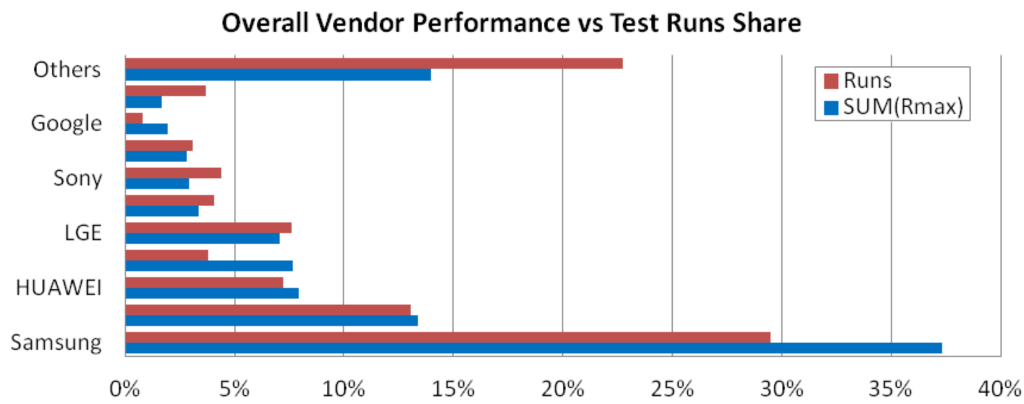

Fig. 8. Overall Vendor Performance vs Test Runs Share

8. Top vendors. Excluding virtual machines from consideration at the data preparation stage, as well as iOS Apple devices due to inaccessibility of the serial version of the test, there is a Samsung leadership found both in terms of test runs number and in total performance, but slightly giving its performance share away. This can be seen in Figs. 6 and 7, which demonstrate two complimentary visions of the market share — performance share, and the number of test runs share, that is popularity. The graphs look similar at a glance, so Fig. 8 provides the comparison.

The second place is taken by Xiaomi by a jump since middle 2016. The “Others“ group is clearly seen as numerous devices with a rather low achieved performance. A steady increase in the number of launches is observed on Huawei devices while LGE drops its share from 14% to 5%.

Figure 9 illustrates the results of the above-mentioned shares comparison in dynamics. We can see also that Samsung has a better share regarding performance in 4 of 5 half a year periods. It proves that the vendor has been always targeting top-level portables segment, what we have just seen in Fig. 5 for the year 2018.

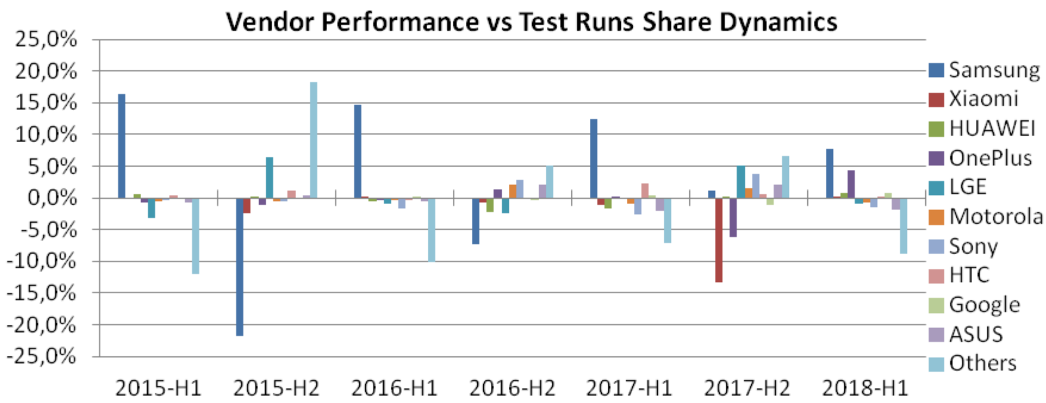

Fig. 9. Vendor Performance vs Test Runs Share Dynamics

9. Further work. With a high and constant rate of application run we see the user interest in respect to the ranking, and it is proved by the statistical data analysis results, so the ranking should be evaluated. The further development splits into these two main branches.

The obtained results represent interesting trends and facts. The most interesting trend is that the ratio of performance per core increases a few times slower than the bytes per core ratio, that is just opposite the trend in big computing.

First, it is a development of benchmark application that should support the features, which the present implementation definitely lacks:

— iOS parallel implementation,

— support for onboard accelerators,

— user-friendly option adjustment interface.

Second, it is a ranking engine development. As of today, the authors suggest moving to the engine based on the renewed Top50 supercomputer ranking.

The authors are grateful for any feedback and contribution of the users, any ideas or bug reports are warmly welcome.

The research is carried out using the equipment of the shared research facilities of HPC computing resources at Lomonosov Moscow State University.

References

1. J. J. Dongarra, P. Luszczek, and A. Petitet, “The LINPACK Benchmark: Past, Present and Future,” Concurrency Computat.: Pract. Exper. **15** (9), 803–820 (2003).

2. The Top500 List of the World’s Most Powerful Supercomputers. https://www.top500.org/. Cited October 31, 2018.

3. The Top50 List of Supercomputers in the Russian Federation. http://top50.supercomputers.ru/?page=rating. Cited October 31, 2018.

4. The Irish Supercomputer List. http://www.irishsupercomputerlist.org. Cited October 31, 2018.

5. Top Supercomputers — India. http://topsc.cdacb.in. Cited October 31, 2018.

6. The Mobile Linpack Official Website. http://linpack.hpc.msu.ru. Cited October 31, 2018.

7. The Mobile Linpack Results Table. http://linpack.hpc.msu.ru/index.php. Cited October 31, 2018.

Received October 19, 2018