

# A Distributed System for SearchOnMath Based on the Microsoft BizSpark Program

Ricardo M. Oliveira  
Universidade Federal de Alfenas  
Rua Gabriel Monteiro da Silva, 700  
Alfenas, MG 37130-001, Brazil  
ricardo.oliveira@bcc.unifal-mg.edu.br

Valmir C. Barbosa  
Universidade Federal do Rio de Janeiro  
Caixa Postal 68511  
Rio de Janeiro, RJ 21941-972, Brazil  
valmir@cos.ufrj.br

Flavio B. Gonzaga  
Universidade Federal de Alfenas  
Rua Gabriel Monteiro da Silva, 700  
Alfenas, MG 37130-001, Brazil  
fbgonzaga@bcc.unifal-mg.edu.br

Geraldo B. Xexéo  
Universidade Federal do Rio de Janeiro  
Caixa Postal 68511  
Rio de Janeiro, RJ 21941-972, Brazil  
xexeo@cos.ufrj.br

## ABSTRACT

Mathematical information retrieval is a relatively new area, so the first search tools capable of retrieving mathematical formulas began to appear only a few years ago. The proposals made public so far mostly implement searches on internal university databases, small sets of scientific papers, or Wikipedia in English. As such, only modest computing power is required. In this context, SearchOnMath has emerged as a pioneering tool in that it indexes several different databases and is compatible with several mathematical representation languages. Given the significantly greater number of formulas it handles, a distributed system becomes necessary to support it. The present study is based on the Microsoft BizSpark program and has aimed, for 38 different distributed-system scenarios, to pinpoint the one affording the best response times when searching the SearchOnMath databases for a collection of 120 formulas.

## CCS CONCEPTS

•Information systems →Distributed retrieval;

## KEYWORDS

Mathematical information retrieval, SearchOnMath, Microsoft BizSpark

### ACM Reference format:

Ricardo M. Oliveira, Flavio B. Gonzaga, Valmir C. Barbosa, and Geraldo B. Xexéo. 2017. A Distributed System for SearchOnMath Based on the Microsoft BizSpark Program. In *Proceedings of ACM SIGIR Conference on Research and Development in Information Retrieval, Tokyo, Japan, August 2017 (SIGIR'2017)*, 4 pages.

DOI:

---

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

*SIGIR'2017, Tokyo, Japan*

© 2017 Copyright held by the owner/author(s).

DOI:

## 1 INTRODUCTION

Unlike textual information retrieval, for which there exist several techniques already widely studied and disseminated, as well as tools capable of tackling the required tasks while performing quite satisfactorily, the area of Mathematical Information Retrieval (MIR) is still in a much less developed stage. In fact, as summarized in Table 1, only in the past few years have techniques for MIR been introduced, usually focusing on very specific problems related to Wikipedia's mathematical pages and indexing around 500 000 formulas.

As with most niche-oriented forms of information retrieval, MIR has to contend with problems that are specific to the search for mathematical formulas. One of them is the large overhead caused by the various possible uses for the same symbol [10]. These possibilities constitute an important source of ambiguity in MIR, since completely different formulas can be written using essentially the same symbols [3]. Another problem is the fact that usually the formulas available on the

Table 1: Existing Tools for MIR

Reference	Search Domain	No. of Formulas
[4]	CONNEXIONS project, functions.wolfram	77 000 87 000
[1]	Coq proof assistant	40 000 theorems
[7]	Database created by authors	829
[9]	50 L <sup>A</sup> T <sub>E</sub> X documents	24 479
[3]	en.wikipedia.org, DLMF	611 210 252 148
[2]	en.wikipedia.org	495 958
[5]	en.wikipedia.org, CiteSeerX	521 782 9 482
[11]	en.wikipedia.org <sup>1</sup>	482 364
[13]	en.wikipedia.org <sup>2</sup>	387 947

<sup>1</sup>Used MREC (Math REtrieval Collection), a collection with approximately 324 000 academic publications, during the development phase.

<sup>2</sup>Includes some information about index size and response time when applied to an arXiv base with about 60 million formulas.

Web are represented in languages originally conceived with little or no concern for a formula's semantic aspects.

The search engine SearchOnMath<sup>3</sup> is one of the most recent arrivals to the field of MIR. Its first version was released in 2013, and by the end of 2015 it had become a start-up. It soon joined the Microsoft BizSpark program, with a modest but very effective monthly allowance, distributed among five email accounts, to hire computers (at most 20 processing cores). Until 2016, SearchOnMath was able to perform the search for formulas on four databases, namely, the English version of Wikipedia, Wolfram MathWorld, DLMF, and PlanetMath. In such a scenario, only one computer with the so-called A3 Basic configuration of the Microsoft Azure environment was sufficient. This configuration includes a four-core processor, 7 GiB of RAM, and a 120-GB HD.

In the course of 2016, as SearchOnMath began preparations for expansion, a distributed system was developed and tested on Microsoft Azure with the goal of assessing each of the 38 possible configurations afforded by our constraints within BizSpark. This investigation was based on a set of 120 preselected formulas that were to be worked on by SearchOnMath within a domain of almost 2 million formulas, and aimed at discovering which of the candidate configurations was capable of delivering the best response times. Our results and conclusions are presented in this paper.

Our study contributes to the field of MIR in two different ways. The first of them is more of a confirmation of the path we have selected for SearchOnMath. It is therefore of an immediate nature, with short-term applicability by other entrepreneurs. As companies that develop search engines for mathematical formulas begin to appear, mainly as start-ups, it may be reassuring to know that the Microsoft BizSpark program is a very viable opportunity, since it already supports more than 100 000 start-ups worldwide and continues to expand. In this regard, information about the infrastructure and operation of SearchOnMath on Azure can be more widely useful. The second contribution is the performance assessment we carried out itself, including the set of 120 formulas that we put together in order to measure response time, but which can be used for other purposes as well.

## 2 METHODOLOGY

For the present study we considered five databases, all obtained throughout the year 2016. Each of these databases is identified in Table 2, along with the respective number of mathematical formulas extracted from it, disregarding repetitions.

After the bases were obtained individually, a final database was constructed as the union of all five, still disregarding repetitions. The resulting database contains 1 905 358 indexed formulas. SearchOnMath was then configured as in Figure 1.

<sup>3</sup><http://searchonmath.com/>

<sup>4</sup><http://en.wikipedia.org/>

<sup>5</sup><http://mathworld.wolfram.com/>

<sup>6</sup><http://dlmf.nist.gov/>

<sup>7</sup><http://planetmath.org/>


<sup>8</sup><http://socratic.org/>

**Table 2: Database List**

Database	Number of Formulas
Wikipedia, English version <sup>4</sup>	590 417
Wolfram MathWorld <sup>5</sup>	79 677
DLMF <sup>6</sup>	33 219
PlanetMath <sup>7</sup>	159 944
Socratic <sup>8</sup>	1 063 754

For operation, a client submits a formula to be searched to the master machine, which runs the engine's front-end. After reception by the master, the formula is sent to the slave machines, which do all the necessary processing to find out which formulas in the database are similar to the query formula. The database is distributed across the slaves so that, for example, if we have 10 of them, then each one has approximately 10% of the formulas. After processing, each slave returns a list containing the most similar formulas it found. The master receives all the lists and then performs the final ordering of the results, returning the consolidated list to the client.

All machines run Linux, and in all cases we configured the master machine with four cores, 7 GiB of RAM, and a 120-GB HD (this configuration is called A3 Basic in Azure). The number of slave machines was obtained based on the remaining allowance resources. We first estimated the amount of our monthly allowance that would correspond to an hour, and then took into account the fact that Azure prices the allocation of machines differently, depending on geographic location. We always chose the region that offered the lowest possible cost in the United States, considering as reference value the one quoted at the date of the beginning of the experiments (Nov. 23, 2016). In these circumstances, discounting the A3 Basic cost per hour allowed us the allocation of up to 16 cores to work as slave machines, arranged according to 38 (out of 65) different configurations available in the Azure environment.



**Figure 1: The SearchOnMath architecture.**

**Table 3: Slave-Machine Configurations and Azure Groups**

Configuration	Cores	RAM	HD	Machines
A0 Basic	1	0.75 GiB	20 GB	16
A1 Basic	1	1.75 GiB	40 GB	16
A2 Basic	2	3.50 GiB	60 GB	8
A3 Basic	4	7.00 GiB	120 GB	4
A4 Basic	8	14.00 GiB	240 GB	2
A1 v2	1	2.00 GiB	10 GB	16
A2 v2	2	4.00 GiB	20 GB	8
A4 v2	4	8.00 GiB	40 GB	4
A8 v2	8	16.00 GiB	80 GB	2
A2m v2	2	16.00 GiB	20 GB	8
A4m v2	4	32.00 GiB	40 GB	3
A8m v2	8	64.00 GiB	80 GB	1
D1 v1	1	3.50 GiB	50 GB	12
D2 v1	2	7.00 GiB	100 GB	6
D3 v1	4	14.00 GiB	200 GB	3
D4 v1	8	28.00 GiB	400 GB	1
D1 v2*	1	3.50 GiB	50 GB	14
D2 v2*	2	7.00 GiB	100 GB	7
D3 v2*	4	14.00 GiB	200 GB	3
D4 v2*	8	28.00 GiB	400 GB	1
A0 Standard	1	0.75 GiB	20 GB	16
A1 Standard	1	1.75 GiB	70 GB	16
A2 Standard	2	3.50 GiB	135 GB	8
A3 Standard	4	7.00 GiB	285 GB	4
A4 Standard	8	14.00 GiB	605 GB	2
A5 Standard	2	14.00 GiB	135 GB	3
A6 Standard	4	28.00 GiB	285 GB	1
F1*	1	2.00 GiB	16 GB	16
F2*	2	4.00 GiB	32 GB	8
F4*	4	8.00 GiB	64 GB	4
F8*	8	16.00 GiB	128 GB	2
D11 v1	2	14.00 GiB	100 GB	4
D12 v1	4	28.00 GiB	200 GB	2
D13 v1	8	56.00 GiB	400 GB	1
D11 v2*	2	14.00 GiB	100 GB	5
D12 v2*	4	28.00 GiB	200 GB	2
D13 v2*	8	56.00 GiB	400 GB	1
G1*	2	28.00 GiB	384 GB	1

Table 3 shows all the configurations analyzed for the slave machines. The Configuration column indicates the name of the configuration, its resources detailed in the Cores, RAM, and HD columns. The Machines column indicates the number of machines with this configuration that could be instantiated as slaves. This number is equal to either  $\lfloor h/p \rfloor$  or  $\lfloor 16/c \rfloor$ , whichever is smaller, where  $h$  is the available budget per hour,  $p$  is the cost per hour of instantiating one machine, and  $c$  is the number of cores one machine has.

Azure groups similar machine configurations [6]. In Table 3, a white backdrop indicates machines classified as “General Purpose—Balanced CPU to memory ratio.” A light-gray backdrop indicates “Compute Optimized—High CPU

to memory ratio” machines. Those on a dark-gray backdrop, finally, are “Memory Optimized—High memory to core ratio” machines. Configurations with an asterisk (\*) by their denominations comprise machines that Azure offers with or without an SSD. Thus, for these machines, both possibilities were evaluated.

### 3 RESULTS

All tests were executed on a set of 120 formulas<sup>9</sup> from [3, 7, 8, 11, 12].

The overall testing scheme for each line of Table 3 (each configuration of the slave machines in Figure 1) was the following. The first of the 120 formulas was submitted for search to the master machine (of type A3 Basic), which then passed it on to the slave machines (of types dependent upon the configuration in question, as per Table 3) and awaited their results. Having received these, the master machine put together and sorted the final list of results and proceeded to submitting the second formula in the set. This was repeated until all 120 formulas were searched.

This full search pass over all 120 formulas was repeated 41 times for each of the configurations of Table 3. The time spent on each pass was recorded and, at the end, the average time of all 41 executions was found and its confidence interval estimated (at the 99% level). We note that each time measurement disregards every communication delay between the client and the master (cf. Figure 1). As a result, all time figures we report are search-related, referring to processing time at the master or at a slave, or to internal network delays of the distributed system. We give results in Figures 2 and 3, respectively for the machines of Azure type General Purpose and for those of the other two types (Compute Optimized and Memory Optimized).

Each plot in these figures refers to a group of slave-machine configurations, as implied by the horizontal rules in Table 3, and positions each of the group’s configurations on the abscissa axis in the order given in the table. So, for example, configurations A0–4 Basic are grouped together, with A0 Basic appearing leftmost in Figure 2, followed by A1 Basic, and so on. It is also worth noting that the confidence intervals are often negligible and therefore hard to discern in the figures.

As it turns out, the best scenario for the SearchOnMath system is the slave-machine configuration F1 with SSD, which comprises 16 identical single-core machines, each with 2 GiB of RAM and a 16-GB SSD. With this configuration, the time needed to search for the 120 formulas was about 120 seconds on average (roughly 1 second per formula), with a confidence interval of approximately  $\pm 0.86$  seconds.

Notwithstanding this, we note that in general the SSD-based configurations did not result in a large difference when compared to their HD-based counterparts. This was expected, given that SearchOnMath carries the formulas in memory while running, thus considerably reducing the need for access to secondary storage. Two exceptions to this note occurred for configurations F1 and F4, in which case time differences

<sup>9</sup><http://searchonmath.com/formulas>, accessed: Feb. 18, 2017.




Figure 2: Time spent on General Purpose machines.




Figure 3: Time spent on Compute Optimized and Memory Optimized machines.

were indeed significant. Nevertheless, we are unable to explain such differences on grounds of the SearchOnMath algorithms, and must therefore speculate that they have to do with factors internal to Azure.

## 4 CONCLUSIONS

Carrying out the experiments described in this paper has allowed us to observe the functioning of SearchOnMath on a variety of configurations of the Microsoft Azure cloud environment. We experimented with all configurations compatible with our BizSpark status and, within these limits, identified a configuration capable of supporting 1-second searches for 120 (out of just over 1 900 000) formulas. At the relatively modest cost currently afforded us by the Microsoft BizSpark program, these experiments will help us envisage plans to scale up operations. We note, finally, that the 120 formulas selected for the experiments will remain available from

SearchOnMath for possible future use in further comparative studies.

## ACKNOWLEDGMENTS

We thank the Federal University of Alfnas and NidusTec Business Incubator, as well as CNPq, CAPES, and a FAPERJ BBP grant for financial support. We also thank Microsoft for the opportunity to participate in their BizSpark program.

## REFERENCES

- [1] A. Asperti, F. Guidi, C. S. Coen, E. Tassi, and S. Zacchiroli. 2006. A Content Based Mathematical Search Engine: Whelp. In *Types for Proofs and Programs*, J.-C. Filliâtre, C. Paulin-Mohring, and B. Werner (Eds.). Lecture Notes in Computer Science, Vol. 3839. Springer, Berlin, Germany, 17–32. DOI:http://dx.doi.org/10.1007/11617990\_2
- [2] X. Hu, L. Gao, X. Lin, Z. Tang, X. Lin, and J. B. Baker. 2013. WikiMirs: A Mathematical Information Retrieval System for Wikipedia. In *Proceedings of the 13th ACM/IEEE-CS Joint Conference on Digital Libraries*. ACM, New York, 11–20. DOI: http://dx.doi.org/10.1145/2467696.2467699
- [3] S. Kamali and F. Wm. Tompa. 2013. Retrieving Documents with Mathematical Content. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, New York, 353–362. DOI: http://dx.doi.org/10.1145/2484028.2484083
- [4] M. Kohlhasse and I. Sucan. 2006. A Search Engine for Mathematical Formulae. In *Artificial Intelligence and Symbolic Computation*, J. Calmet, T. Ida, and D. Wang (Eds.). Lecture Notes in Computer Science, Vol. 4120. Springer, Berlin, Germany, 241–253. DOI:http://dx.doi.org/10.1007/11856290\_21
- [5] X. Lin, L. Gao, X. Hu, Z. Tang, Y. Xiao, and X. Liu. 2014. A Mathematics Retrieval System for Formulae in Layout Presentations. In *Proceedings of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, New York, 697–706. DOI:http://dx.doi.org/10.1145/2600428.2609611
- [6] Microsoft. 2017. Azure Price Calculator. (2017). https://azure.microsoft.com/en-us/pricing/details/virtual-machines/linux/ Accessed: Nov. 23, 2016.
- [7] P. Pavan Kumar, A. Agarwal, and C. Bhagvati. 2012. A Structure Based Approach for Mathematical Expression Retrieval. In *Multi-Disciplinary Trends in Artificial Intelligence*, C. Sombaththeera, N. K. Loi, R. Wankar, and T. Quan (Eds.). Lecture Notes in Computer Science, Vol. 7694. Springer, Berlin, Germany, 23–34. DOI:http://dx.doi.org/10.1007/978-3-642-35455-7\_3
- [8] L. Salem, F. Testard, and C. Salem. 1992. *The Most Beautiful Mathematical Formulas*. Wiley, New York.
- [9] T. Schellenberg, B. Yuan, and R. Zanibbi. 2012. Layout-Based Substitution Tree Indexing and Retrieval for Mathematical Expressions. In *Document Recognition and Retrieval XIX*, C. Viard-Gaudin and R. Zanibbi (Eds.). Proceedings of SPIE, Vol. 8297. SPIE, Bellingham, Washington, 829701. DOI: http://dx.doi.org/10.1117/12.912502
- [10] M. Schubotz, A. Grigorev, M. Leich, H. S. Cohl, N. Meuschke, B. Gipp, A. S. Youssef, and V. Markl. 2016. Semantification of Identifiers in Mathematics for Better Math Information Retrieval. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, New York, 135–144. DOI:http://dx.doi.org/10.1145/2911451.2911503
- [11] D. Stalnakar and R. Zanibbi. 2015. Math Expression Retrieval Using an Inverted Index Over Symbol Pairs. In *Document Recognition and Retrieval XXII*, E. K. Ringger and B. Lamiroy (Eds.). Proceedings of SPIE, Vol. 9402. SPIE, Bellingham, Washington, 940207. DOI:http://dx.doi.org/10.1117/12.2074084
- [12] I. Stewart. 2012. *In Pursuit of the Unknown: 17 Equations That Changed the World*. Basic Books, New York.
- [13] R. Zanibbi, K. Davila, A. Kane, and F. Wm. Tompa. 2016. Multi-Stage Math Formula Search: Using Appearance-Based Similarity Metrics at Scale. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, New York, 145–154. DOI: http://dx.doi.org/10.1145/2911451.2911512