Recursive MAGUS: scalable and accurate multiple sequence alignment

Vladimir Smirnov^{1,*}

¹ Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL, United States

* smirnov3@illinois.edu

Abstract

Multiple sequence alignment tools struggle to keep pace with rapidly growing sequence data, as few methods can handle large datasets while maintaining alignment accuracy. We recently introduced MAGUS, a new state-of-the-art method for aligning large numbers of sequences. In this paper, we present a comprehensive set of enhancements that allow MAGUS to align vastly larger datasets with greater speed. We compare MAGUS to other leading alignment methods on datasets of up to one million sequences. Our results demonstrate the advantages of MAGUS over other alignment software in both accuracy and speed. MAGUS is freely available in open-source form at <https://github.com/vlasmirnov/MAGUS>.

Introduction

One of the principal problems in computational biology is multiple sequence alignment (MSA), being necessary for a wide range of downstream applications. This challenge is well-studied, and a good number of strong methods have been developed [1–8]. Most of these leading methods follow the paradigm of “progressive alignment”, and are able to show reasonable accuracy and speed on datasets of modest size (a few hundred to a few thousand sequences).

Unfortunately, datasets with more sequences and greater evolutionary diameters require a different approach. Accurate progressive alignment methods rely on heuristics whose runtimes scale very poorly, and early mistakes are compounded over large numbers of pairwise alignments. As a consequence, a family of divide-and-conquer methods was developed to meet the demands of larger datasets [9–11].

MAGUS (Multiple Sequence Alignment using Graph Clustering) was recently introduced [12] as a new evolution of this family. MAGUS uses the GCM (Graph Clustering Merger) technique to combine an arbitrary number of subalignments, which allows MAGUS to align large numbers of sequences with highly competitive accuracy and speed. In its original form, MAGUS is able to align up to around 40,000 sequences.

In this paper, we extend MAGUS to handle datasets of much greater size, demonstrating alignments of up to one million sequences. The next section briefly explains how MAGUS operates, and presents our extensions to enable scalability. Next, we describe our experimental study and show our results, comparing MAGUS to other methods with regard to alignment accuracy and speed over ultra-large datasets. Finally, we discuss our findings and future work.

Design and Implementation

Overview of MAGUS

MAGUS is a recently developed divide-and-conquer alignment method that inherits the basic structure of the earlier PASTA [11] algorithm: MAGUS decomposes the dataset into subsets, aligns them piecewise, and merges these subalignments together. The basic algorithm is itemized below.

Fig 1. MAGUS overview. The unaligned sequences are decomposed into disjoint subsets, which are individually aligned and merged together with GCM.

Input: a set of unaligned sequences.

1. Construct a guide tree over the unaligned sequences. (Our default way of doing this is explained below.)
2. Use the guide tree to break the dataset into subsets. This is done by “centroid edge decomposition” [11], deleting edges to break the tree into sufficiently small, balanced pieces.
3. Align each subset with MAFFT -linsi [3].
4. Construct a set of backbone alignments spanning our subsets. Each backbone is composed of equal-sized random subsets from each subalignment and aligned with MAFFT -linsi.
5. Compile the backbones into an alignment graph. Each node represents a subalignment column, and the edges are weighted by how often they are matched by the backbone alignments.
6. Cluster the alignment graph with MCL [13].
7. Order the clusters into a valid alignment. We use a heuristic search to resolve conflicts with minimal changes.
8. Output the full alignment.

Please refer to the original paper [12] for more information. Steps 5-8 comprise GCM (Graph Clustering Merger, Fig S1), the method by which MAGUS merges subalignments and its biggest departure from previous divide-and-conquer methods. The pipeline was built to be flexible: the user can supply their own subalignments in lieu of steps 1-4, their own guide tree for step 2, and their own backbones for step 5. The number and size of subsets and backbones can also be controlled.

Motivation for MAGUS Enhancement

Despite its advantages, the original version of MAGUS (“MAGUS 1”) suffers from a number of constraints on its scalability. We motivate the need for improvement by glancing ahead to our experimental study, where MAGUS 1 is seen to struggle with increasing dataset sizes; Fig 6 shows MAGUS 1 taking over 20 hours to align 50,000 sequences, and Table 2 shows MAGUS 1 failing on larger datasets due to memory issues. In the next section, we explain the limitations of MAGUS 1 and present the improvements that comprise the paper.

MAGUS Improvements

Recursion

First, there is a soft limit on how many sequences MAGUS 1 can reasonably align. MAFFT -linsi [3], which is used for building subset and backbone alignments, starts to really slow down past around 200 sequences. Additionally, the cluster ordering step (step 8 above) tends to struggle with more than about 200 subsets. Therefore, assuming a practical limit of about 200 subsets of 200 sequences each, unmodified MAGUS can be expected to handle up to around $200 \times 200 = 40,000$ sequences.

We parry this limitation with a fairly straightforward recursive structure, shown in Figure 2. Instead of automatically aligning our subsets with MAFFT, subsets larger than a threshold are recursively aligned with MAGUS. This threshold can be set by the user and is, by default, the greater of the backbone size and the target subset size used for decomposition. Our subalignments are merged with GCM just as before, regardless of whether each subalignment was estimated with MAFFT or MAGUS.

Parallelism

The next issue is parallelism. MAGUS 1 already implements thread-parallelism: it runs on a single compute node, and it can use all available threads on that node to run MAFFT tasks in parallel. This is more than enough for a few tens of thousands of sequences on a decent machine. However, with ultra-large datasets, we definitely want to benefit from node-parallelism, when multiple compute nodes can collaborate. We implement node-parallelism by extending MAGUS 1’s task management code. MAGUS 1 maintains task files with MAFFT alignments and other self-contained tasks that are pending or running, which allows worker threads to divide the jobs and MAGUS to easily resume in case of failure. Reworking this system to allow for multiple compute nodes to use the same set of files effectively permits any number of nodes to join and take tasks to work on.

Guide Tree

MAGUS decomposes the dataset into subsets by estimating a rough guide tree with FastTree [14], a fast maximum likelihood tree estimation method. Since FastTree requires an alignment, we first compile a rough alignment by aligning 300 random sequences with MAFFT and adding the remaining sequences with HMMER [15]. The

Fig 2. Recursive MAGUS overview. Instead of aligning all of our subsets with MAFFT, subsets larger than a given threshold are recursively aligned with MAGUS. Subsets below the threshold are aligned with MAFFT. As before, all subalignments are merged with GCM.

guide tree is recursively broken apart until the subsets are small enough. This is the same strategy used in PASTA, and seems very difficult to improve upon. On very large alignments, however, even FastTree becomes painfully slow (around 5 days on a million sequences, as will be shown below).

The new version of MAGUS presents a wider range of guide tree options, intended for situations where FastTree might not be fast enough or fails due to numerical issues. The guide tree can now also be generated with Clustal Omega's [2] initial tree method, MAFFT's PartTree [16] initial tree method, and FastTree's minimum evolution tree (i.e. limited to distance-based calculations without maximum likelihood). In extremis, the dataset can be decomposed randomly for maximum speed.

Memory Management and Alignment Compression

Memory management becomes a salient problem when handling very large datasets. For example, without modifications, MAGUS alignments on the full million-sequence RNASim dataset fall between 1 and 3 terabytes (Fig S2). Moreover, simply having too many subalignments loaded into memory at the same time can overrun the available RAM at such dataset sizes.

We solve the latter problem by reworking the code to ensure that at most one subalignment may be fully loaded into memory at any time. With large dataset sizes, this limits the memory complexity of MAGUS to the size of the largest subalignment.

The problem of excessively large alignments is addressed by introducing a method of conservative lossy compression. If MAGUS calculates that the size of the uncompressed alignment will exceed a threshold (100GB by default, may be set by the user), MAGUS will compress the alignment to the threshold size. The compression scheme is fairly straightforward and works by "dissolving" columns: the letters are set to lower-case and shunted to neighboring columns. If the neighboring columns already contain lower-case letters from the same sequences, these are also shunted away in a recursive domino effect. (If the neighboring columns already contain upper-case letters from the

same sequences, then the move is invalid.) Columns are dissolved one at a time, starting with those containing the fewest letters, until the threshold is reached or no more valid moves remain.

Results

Experimental Design

Our experimental design is outlined below.

The preliminary portion of our study that explores the effects of our MAGUS extensions described above, using MAGUS 1 as our baseline. We test the impact of compression on alignment error, the use of different guide trees, and the benefit of node-parallelism. Due to space limitations, these results are available in the Supplementary Materials.

Our subsequent experiments compare MAGUS against a range of competing methods across all of our datasets. This is the most important part of our study, intended to exercise the current state-of-the-art in the alignment of ultra-large nucleotide and protein datasets. We present our results below.

Datasets

Our study uses a number of simulated and biological datasets from previous publications [4, 11]. Please see Table 1 for dataset statistics. These datasets were selected to provide suitably large and varied alignment problems with reference alignments, containing both nucleotide and amino acid sequences.

- **RNASim:** [11] This is a simulated RNA dataset, generated under a non-homogeneous model of evolution that does not conform to the usual GTR model assumptions. We use subsamples ranging from 10,000 to the full one million sequences, with one replicate per size.
- **16S:** [17] We use three large biological nucleotide datasets from the Comparative Ribosomal Website: 16S.3, 16S.T, and 16S.B.ALL, with 6,323, 7,350, and 27,643 sequences, respectively.
- **HomFam:** [2] Finally, we include 19 amino acid HomFam datasets from, which have small Homstrad reference alignments on 5–20 sequences each. These datasets range from 10,099 to 93,681 sequences and allow us to evaluate our methods on large protein datasets. (Following the PASTA paper, we exclude the “rhv” dataset due to having a weak alignment.)

Methods

We compare the following methods in our study, taken from previous publications [4, 11]. To the best of our knowledge, these methods are presently the best-equipped to tackle very large multiple sequence alignments. Regressive T-Coffee [18] is another recent development, but we were unable to run it on Blue Waters.

- **MAGUS 1** We use the original MAGUS as a baseline. This version does not use recursion or compression, uses a FastTree decomposition, and can only run on a single node.

Dataset	# Seqs	Avg. p-dist.	Max p-dist.	% gaps	align. length	type
RNASim	10,000-1,000,000	0.41	0.61	93	18,268	sim NT
16S						
- 16S.3	6,323	0.32	0.83	82	8,716	bio NT
- 16S.T	7,350	0.35	0.90	87	11,856	bio NT
- 16S.B.ALL	27,643	0.21	0.77	80	6,857	bio NT
HomFam						
- gluts	10,099	0.60	0.81	8	235	bio AA
- myb-DNA-binding	10,398	0.59	0.77	12	61	bio AA
- tRNA-synt-2b	11,293	0.81	0.88	34	467	bio AA
- biotin-lipoyl	11,833	0.71	0.84	26	112	bio AA
- hom	12,037	0.64	0.84	35	98	bio AA
- ghf13	12,607	0.72	0.84	25	626	bio AA
- aldoserred	13,277	0.57	0.79	19	386	bio AA
- hla	13,465	0.24	0.33	0	178	bio AA
- Rhodanese	14,049	0.76	0.89	31	216	bio AA
- PDZ	14,950	0.69	0.84	15	110	bio AA
- blmb	17,200	0.79	0.90	30	344	bio AA
- p450	21,013	0.79	0.87	20	512	bio AA
- adh	21,331	0.36	0.47	0	375	bio AA
- aat	25,100	0.71	0.87	15	476	bio AA
- rrm	27,610	0.77	0.91	45	157	bio AA
- Acetyltransf	46,285	0.75	0.87	29	229	bio AA
- sdr	50,157	0.77	0.89	28	361	bio AA
- zf-CCHH	88,345	0.65	0.85	25	39	bio AA
- rvp	93,681	0.63	0.76	19	132	bio AA

Table 1. Dataset properties. Statistics taken from [11]. P-distance denotes the normalized Hamming distance, or the fraction of non-gap letter pairs that do not match. Alignment length shows the length of the reference alignment.

- **MAGUS** The latest version takes advantage of the new features detailed above. We enable recursion and compress alignments above 100GB. In addition to the default FastTree decomposition, we explore other guide trees: FastTree without Maximum Likelihood, MAFFT's PartTree, Clustal Omega's initial tree method, and a random decomposition. Henceforth, we indicate the guide tree and use of recursion in parentheses. For example, MAGUS(Recurse, Clustal) denotes MAGUS using Clustal Omega's guide tree and with recursion enabled.
- **PASTA** [11]
- **UPP** [4]
- **UPP(Fast)** We use the "Fast" mode described in the UPP paper.
- **Muscle** [1]
- **Clustal Omega** [2]
- **MAFFT -auto** [3]. The "auto" mode directs MAFFT to choose an appropriate alignment strategy based on the input dataset.

Error Metrics

We evaluate alignment accuracy using SPFP/SPFN (Sum-of-Pairs False Positives and Negatives) rates, computed using FastSP [19]. These values represent the fractions of missing and incorrect homologous pairs in the estimated alignment. For convenience, we show the average of SPFP and SPFN as a single "SP error" in the main paper; SPFP and SPFN are shown separately in the Supplementary Materials. Our estimated alignments are compared against the true alignment on RNASim and the curated

reference alignments on 16S. The HomFam datasets provide reference alignments over a small number of included sequences; we compute our alignment error over just these reference sequences.

Computing Resources

We used the NCSA Blue Waters supercomputer for our experiments. Our jobs were run on nodes with 32 cores, 64GB of RAM, and a maximum wall time of 7 days.

Experimental Results

The preliminary part of our study, which investigates the impacts of compression, guide tree selection, and node-parallelism, is available in the Supplementary Materials (due to space constraints). These results provide us with two natural guide tree choices for MAGUS: using FastTree (the default, described above) is the most accurate, while using Clustal Omega's initial tree is the faster alternative. Here, we present the principal part of our study, where we compare MAGUS to our other methods across all of our datasets.

HomFam

Our first set of results concern the HomFam protein datasets. The error rates are averaged in Fig 3, and the complete results for all datasets are available in Table S3. These results show more variability than the other datasets, but the general trends are as follows. Muscle and Clustal trail the others, averaging 46.6% and 27.2% error, respectively. MAFFT, UPP, and PASTA are all on par, averaging about 21-23% error. The MAGUS versions perform markedly better: MAGUS(Recurse, Clustal) yields 17.9% error, MAGUS(Recurse, FastTree) shows 16.5%, and MAGUS 1 leads with 15.5%. Furthermore, MAGUS 1 achieves the best result on 12 of the 19 datasets. Recursive MAGUS (both versions) accounts for 2 of the others, while Clustal and UPP each do best on 2.

Fig 3. Average SP error on HomFam datasets. Error is the average of SPFP and SPFN. Results are averaged over the datasets where all methods completed (Muscle segfaulted on two). Error bars show standard error. MAGUS was run with the default 25 subsets.

The HomFam runtime results are shown in Fig 4 and Fig S14. PASTA is visibly the slowest, taking about 2-5 hours on the smaller datasets and up to 20 hours on the larger ones. MAFFT, UPP(Fast), Clustal Omega, and MAGUS(Recurse, Clustal) are the fastest, generally finishing in a few minutes to an hour. Notably, we see MAGUS 1 begin to dramatically slow down without recursion, running longer than MAGUS(Recurse, FastTree) on the largest datasets.

Fig 4. Homfam (largest 10 datasets) runtime, all methods. MAGUS was run with the default 25 subsets. Muscle segfaulted on the two largest datasets.

16S

The next set of results pertain to the biological 16S datasets, shown in Fig 5. As above, Muscle and Clustal trail the other methods in accuracy. On the smallest dataset, 16S.3, the results are fairly close: UPP(Fast), PASTA, and all versions of MAGUS are at about 19% SP error. There is a larger difference on 16S.T, with PASTA at around 23%, UPP and UPP(Fast) around 21%, and all versions of MAGUS at about 20%. Lastly, UPP, PASTA, and MAGUS are again fairly close on 16S.B.ALL; PASTA shows about 11% error, while both versions of UPP and MAGUS have about 10.5% error.

In terms of runtime, we see that UPP, PASTA, and both versions of recursive MAGUS are the slowest methods on 16S.3 and 16S.T, running around 4-5 hours. The fastest method is MAFFT(auto) at about 2 minutes, while Muscle and UPP(Fast) take about half an hour. The picture is a little different on 16S.B.ALL, where Muscle, UPP, and PASTA seem to drastically slow down; they take about 11, 14, and 17 hours, respectively. MAGUS 1 also falters here, taking 18 hours, while recursive MAGUS with FastTree and Clustal only increases to 8 and 4 hours, respectively. MAFFT and UPP(Fast) remain the fastest, only taking 1-2 hours.

RNASim

In the final part of our study, we probe the limits of scalability on the RNASim datasets. Fig 6 shows us the error and runtime results, while Table 2 summarizes all method failures. Muscle is the worst performer here, with 65-70% error and segfaulting after 50,000 sequences. Clustal Omega does better, with errors between about 30% and 60%, running out of time after 200,000 sequences. Then comes MAFFT -auto, with a steady

Fig 5. 16S alignment error (top) and runtime (bottom), all methods. Error is the average of SPFP and SPFN. Runtime is shown in hours. MAGUS was run with the default 25 subsets. MAFFT -auto completed in a few minutes on 16S.3 and 16S.T.

error of 25-30% up to 100,000 sequences. Oddly, even though it is one of the fastest methods at 100,000 sequences (about 3.6 hours), it runs out of time at 200,000 sequences.

The accuracy of our remaining methods is shown more clearly in Fig 7. UPP(Fast) trails the other methods in accuracy, with about 2% higher error than PASTA and UPP. PASTA and UPP are about the same at around 10% error. MAGUS 1 and recursive MAGUS (both versions) have the best accuracy. MAGUS 1 is the most accurate at 10,000-50,000 sequences (8.2-7.8% error), but can't proceed beyond that. MAGUS(Recurse, FastTree) is second-best at about 8.5-8%. MAGUS(Recurse, Clustal) consistently trails MAGUS(Recurse, FastTree) by about 0.5% below 200,000 sequences, and declines to about 8.3% on 1,000,000 sequences.

Aside from MAGUS(Recurse, Clustal), UPP(Fast) is the only other method that aligned all 1,000,000 sequences in a week; UPP took about 77 hours to align all 1,000,000 sequences, while MAGUS(Recurse, Clustal) took about 128 hours. PASTA

Method	Highest # Aligned	Failure
Muscle	50,000	"segmentation fault"
Clustal Omega	200,000	Max runtime elapsed (7 days)
MAFFT(auto)	100,000	Max runtime elapsed (7 days)
UPP	200,000	Max runtime elapsed (7 days)
PASTA	50,000	"Error detected during page fault processing. Process terminated via bus error."
MAGUS 1	50,000	"OOM killer terminated this process."

Table 2. Method failures on RNASim. PASTA and MAGUS 1 failed due to excessive memory usage; compute nodes had 64GB of memory.

Fig 6. RNASim alignment error (top) and runtime (bottom), all methods. Error is the average of SPFP and SPFN. Runtime is show in hours. 'X' markers indicate that compression was used (MAGUS alignments above 100GB). MAGUS was run with 100 subsets on RNASim to reduce load on Blue Waters. Compute nodes had 64GB of RAM and a maximum wall time of 7 days.

encountered memory issues, while UPP and MAGUS(Recurse, FastTree) ran out of time. Notably, UPP, Clustal Omega, and MAGUS(Recurse, FastTree) showed comparable runtime scaling, all three just meeting the 1 week time limit at 200,000 sequences. MAGUS 1 initially scales better than recursive MAGUS on a single node,

but only reaches 50,000.

Fig 7. RNASim alignment error, best methods. Error is the average of SPFP and SPFN. 'X' markers indicate that compression was used (MAGUS alignments above 100GB). MAGUS was run with 100 subsets on RNASim to reduce load on Blue Waters. Compute nodes had 64GB of RAM and a maximum wall time of 7 days.

Discussion

The accuracy of MAGUS convincingly exceeds the other methods we tried on the datasets in our study. As shown in Figs 6, 5, and 3, this is true regardless of whether recursion is used, and whether FastTree or Clustal is used for decomposition. The more difficult question we need to tease apart concerns the different ways of running MAGUS, and how they affect scalability and accuracy. We do this by considering recursion, guide tree, and node-parallelism in turn.

On one hand, recursion actually slows MAGUS down on smaller datasets. On the other hand, this is rapidly reversed as MAGUS chokes on larger datasets without recursion. This can be seen from our 16S results, where MAGUS is much faster without recursion on 6,000-7,000 sequences, but much slower on 27,000. This reversal can also be seen on the HomFam datasets. On RNASim, MAGUS without recursion is faster on 10,000-50,000 sequences, but simply fails after that.

The nature of this limitation is fairly clear: given N sequences and S subsets, MAGUS without recursion must run MAFFT -linsi on chunks of $\frac{N}{S}$ sequences. Thus, MAGUS without recursion is only viable for as long as MAFFT -linsi can handle these chunks. Our results suggest that subsets approaching around 1,000 sequences really become a problem: this is about where RNASim fails and 16S.B.ALL takes an inordinate amount of time. There is less of a problem on HomFam, where the amino acid sequences are much shorter.

Moreover, recursion does not improve accuracy; MAGUS without recursion is noticeably more accurate on HomFam, about the same on 16S, and slightly better on 10,000 sequences of RNASim. These observations suggest that recursion should be avoided if possible, and only engaged when the dataset becomes too large for the subsets to be reasonably aligned with the base method.

As far as decomposition strategy is concerned, the FastTree method remains the most accurate. The runtime becomes an issue on the largest datasets, where the tree takes

about 5 days to compute on 1 million sequences. The best alternative, as suggested by our results, is to use the Clustal Omega guide tree. This gives the best compromise between accuracy and runtime, and only takes 14 hours on 1 million sequences.

Taking advantage of our newfound node-parallelism has a considerable impact on runtime. If we exclude the FastTree computation from the MAGUS runtime on 1 million sequences, the actual alignment stage takes about 9 days on a single node, but only about 17 hours on 10 nodes and 2.5 hours on 100 nodes. Thus, given enough compute nodes, the total runtime is mostly dominated by the guide tree method, rather than the alignment itself; this is the motivation for considering Clustal as a FastTree alternative.

Conclusions

We presented a powerful set of improvements to our MAGUS method, allowing it to scale from 50,000 to a full million sequences. Moreover, MAGUS is able to align such vast datasets more accurately than the other methods we compared against.

UPP(Fast) remains the fastest way to effectively align a million sequences on a single compute node, but suffers from consistently worse alignment accuracy. Other methods are able to finish quickly on smaller datasets, but struggle to complete on larger numbers of sequences, while also trailing MAGUS in accuracy.

We conclude by distilling our results into a number of concrete recommendations for interested practitioners.

Recursion is harmful on smaller datasets, but necessary on larger datasets. If the dataset is small enough, MAGUS will run considerably faster without recursion and might have slightly better accuracy. On larger datasets, MAGUS will rapidly grind to a halt without recursion. Thus, it is advised to avoid recursion if the dataset permits this. This threshold is dictated by subset size ($\frac{\# \text{ sequences}}{\# \text{ subsets}}$). Given our data, we found the “threshold” subset size to be somewhere around 1,000 sequences of a few thousand nucleotides, or somewhere above 4,000 sequences of a few hundred amino acids.

The importance of node-parallelism and guide tree. The default FastTree-based subset decomposition gives the best accuracy, and is fast enough for most purposes. For huge datasets of half a million or more, the Clustal Omega-based decomposition runs much faster and is nearly as accurate. As one might expect, using as many compute nodes as possible will improve the runtime. However, using more nodes than subsets will decrease the added gains from node-parallelism.

Running MAGUS. Putting all of the above together, the most accurate way of running MAGUS is to use the default FastTree-based decomposition without recursion, preferably on as many compute nodes as are available. If the dataset is too large to allow the subsets to align in a reasonable amount of time, recursion should be enabled. Finally, if the dataset is too large to allow FastTree to finish in a reasonable amount of time, the Clustal-based decomposition should be used.

Availability and Future Directions

MAGUS is open-source and freely available at <https://github.com/vlasmirnov/MAGUS>. The datasets used in this study can be downloaded from the Illinois Data Bank at https://doi.org/10.13012/B2IDB-1048258_V1.

We plan to explore several future directions towards further improving MAGUS. The first is to comprehensively investigate the performance of MAGUS on fragmentary data.

Fragmentary sequences can potentially confound effective methods, and we will extend MAGUS to reliably handle such scenarios.

The second avenue of improvement is to consider alternative procedures for assembling backbone alignments, and is intended to further increase alignment accuracy. Currently, MAGUS uses the simple expedient of building backbones with equal, random samples from each subset. We will develop and evaluate ways to build more compact (and, thus, more accurate) backbone sets that still sufficiently span the subsets.

Thirdly, we have mostly developed MAGUS to be able to align vast numbers of sequences accurately. In the future, we hope to also extend MAGUS “in the other direction” - to handle datasets with arbitrarily long, even genome-scale sequences.

Funding

This work was funded by the Ira & Debra Cohen Graduate Fellowship to VS and NSF grant ABI-1458652 to Dr. Tandy Warnow.

Commands Used

MAGUS

```
python3 magus.py -d tempdir -o result.txt -i unalign.txt
-t <guide tree option or path> --recurse <true|false>
--maxnumsubsets <25|100>
```

PASTA 1.8.3

```
python3 run\_pasta.py -i unalign.txt -o result.txt
--temporaries tempdir -d <dna|rna|protein> --keeptemp
```

UPP 4.3.10

```
python3 run\_upp.py -s unalign.txt -p result.txt -m rna
```

UPP(Fast) 4.3.10

```
python3 run\_upp.py -s unalign.txt -p result.txt -B 100 -m rna
```

Muscle 3.8.425

```
muscle -maxiters 2 -in unalign.txt -out result.txt
```

Clustal Omega 1.2.4

```
clustalo -i unalign.txt -o result.txt --threads=32
```

MAFFT 7.450 --auto

```
mafft --auto --ep 0.123 --quiet --thread 32 --anysymbol
unalign.txt > result.txt
```

FastSP 1.6.0 (Computing alignment error)

```
java -Xmx256G -jar FastSP\_1.6.0.jar -r reference\_align.txt
-e estimated\_align.txt -ml
```

References

1. Edgar RC. MUSCLE: a multiple sequence alignment method with reduced time and space complexity. *BMC bioinformatics*. 2004;5(1):113.
2. Sievers F, Wilm A, Dineen D, Gibson TJ, Karplus K, Li W, et al. Fast, scalable generation of high-quality protein multiple sequence alignments using Clustal Omega. *Molecular systems biology*. 2011;7(1):539.
3. Katoh K, Kuma Ki, Toh H, Miyata T. MAFFT version 5: improvement in accuracy of multiple sequence alignment. *Nucleic acids research*. 2005;33(2):511–518.
4. Nguyen NpD, Mirarab S, Kumar K, Warnow T. Ultra-large alignments using phylogeny-aware profiles. *Genome Biology*. 2015;16(1):124.
5. Lassmann T. Kalign 3: multiple sequence alignment of large datasets. *Bioinf*. 2019;36(6):1928–1929. doi:10.1093/bioinformatics/btz795.
6. Notredame C, Higgins DG, Heringa J. T-Coffee: A novel method for fast and accurate multiple sequence alignment. *Journal of molecular biology*. 2000;302(1):205–217.
7. Do CB, Mahabhashyam MS, Brudno M, Batzoglou S. ProbCons: Probabilistic consistency-based multiple sequence alignment. *Genome research*. 2005;15(2):330–340.
8. Pei J, Grishin NV. PROMALS: towards accurate multiple sequence alignments of distantly related proteins. *Bioinf*. 2007;23(7):802–808.
9. Liu K, Raghavan S, Nelesen S, Linder CR, Warnow T. Rapid and accurate large-scale coestimation of sequence alignments and phylogenetic trees. *Science*. 2009;324(5934):1561–1564.
10. Liu K, Warnow TJ, Holder MT, Nelesen SM, Yu J, Stamatakis AP, et al. SATE-II: very fast and accurate simultaneous estimation of multiple sequence alignments and phylogenetic trees. *Systematic biology*. 2012;61(1):90.
11. Mirarab S, Nguyen N, Guo S, Wang LS, Kim J, Warnow T. PASTA: ultra-large multiple sequence alignment for nucleotide and amino-acid sequences. *Journal of Computational Biology*. 2015;22(5):377–386.
12. Smirnov V, Warnow T. MAGUS: Multiple Sequence Alignment using Graph Clustering. *Bioinformatics*. 2020;doi:10.1093/bioinformatics/btaa992.
13. Van Dongen SM. A cluster algorithm for graphs. Amsterdam: National Research Institute for Mathematics and Computer Science in the Netherlands; 2000. Available from: <https://ir.cwi.nl/pub/4463>.
14. Price MN, Dehal PS, Arkin AP. FastTree 2—approximately maximum-likelihood trees for large alignments. *PloS one*. 2010;5(3):e9490.
15. Eddy SR. HMMER website; 2020. Available from: <http://hmmer.org>.
16. Katoh K, Toh H. Recent developments in the MAFFT multiple sequence alignment program. *Briefings in bioinformatics*. 2008;9(4):286–298.

17. Cannone JJ, Subramanian S, Schnare MN, Collett JR, D'Souza LM, Du Y, et al. The comparative RNA web (CRW) site: an online database of comparative sequence and structure information for ribosomal, intron, and other RNAs. *BMC Bioinf.* 2002;3(1):2.
18. Garriga E, Di Tommaso P, Magis C, Erb I, Mansouri L, Baltzis A, et al. Large multiple sequence alignments with a root-to-leaf regressive method. *Nature Biotech.* 2019;37(12):1466–1470.
19. Mirarab S, Warnow T. FastSP: linear time calculation of alignment accuracy. *Bioinf.* 2011;27(23):3250–3258.

Recursive MAGUS: scalable and accurate multiple sequence alignment

Vladimir Smirnov^{1,*}

¹ Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL, United States

* smirnov3@illinois.edu

Supplementary Materials

GCM Overview

Fig S1. GCM overview. Given a set of disjoint subalignments, we first compile (or are given) a set of backbone alignments that span our subsets. These are used to build a weighted alignment graph, where each node is a subalignment column, and the edge weights represent our confidence that these columns should be aligned. Then, we use MCL to cluster the graph. The clusters are further refined into a “trace”, which corresponds to a valid multiple sequence alignment. Optionally, the trace can be further optimized with respect to the Maximum Weight Trace criterion.

Preliminary Study: Comparing MAGUS Variants

The Effect of Compression

This part of our study addresses the issues alignment sizes and compression. Fig. S2 and Table S1 show the rapid growth of estimated uncompressed MAGUS alignment sizes with increasing dataset size. The size of the true alignment on the full

million-sequence dataset is 21.4 gigabytes, compared to about 3.4 terabytes produced by recursive MAGUS with PartTree and random decompositions, and about 1 terabyte with the other guide trees.

Fig S2. RNASim log-scale alignment sizes (in GB), MAGUS variants only. MAGUS was run with 100 subsets on RNASim to reduce load on Blue Waters.

#	M(R,Random)	M(R,PT)	M(R,FT)	M(R,FT-NoML)	M(R,Clustal)	M1
10,000	0.6	0.5	0.4	0.4	0.4	0.4
20,000	1.8	2.2	1.3	1.4	1.5	1.1
50,000	15.0	14.9	7.5	7.5	8.6	5.0
100,000	55.6	55.4	25.3	25.3	29.3	
200,000	208.9	207.6	83.1	83.1	97.5	
500,000	1,085.4	1,076.9	365.5	364.9	432.9	
1,000,000	3,441.8	3,424.4	1,037.3	1,036.5	1,228.5	

Table S1. RNASim alignment sizes (in GB) of MAGUS variants. “R” indicates recursion, PT denotes PartTree, FT denotes FastTree. MAGUS was run with 100 subsets on RNASim to reduce load on Blue Waters.

Consequently, we measure the increase in SP error from subjecting our alignments to the MAGUS compression. Table S2 shows the delta SP error of recursive MAGUS (using the FastTree guide tree), with different dataset sizes being compressed down to varying size thresholds. Even under the most aggressive compression policies (e.g. compressing 83 gigabytes to 5), the error increases by minuscule amounts. Looking at the full dataset, we see that compressing the full 1-terabyte alignment to 25 gigabytes only increases the SP error by $1.2e-7$ compared to the 100 gigabyte compression. Therefore, we can safely say that our compression scheme for managing huge alignment sizes has negligible effect on alignment accuracy and can be safely used in subsequent experiments.

The Effects of Different Guide Trees

Next, we compare MAGUS 1 to MAGUS with recursion and a selection of different guide trees. Fig S3 shows the resulting SP error. The runtime analysis is slightly more complicated. We note that the guide tree portion of MAGUS currently always runs on a

	10,000	20,000	50,000	100,000	200,000	500,000	1,000,000
Uncompressed GB	0.4	1.3	7.5	25.3	83.1	365.5	1,037.3
100 GB							
50 GB					0.0	4.3e-09	3.6e-08
25 GB				0.0	1.3e-09	3.6e-08	1.2e-07
10 GB				8.5e-10	1.6e-08	1.8e-07	
5 GB			0.0	7.1e-09	1.2e-07		
1 GB		0.0	5.9e-08				

Table S2. Delta error from lossy compression, MAGUS(Recurse, FastTree) alignments on RNASim. Values show the increase in SP error over the uncompressed alignment. 500K and 1M show increase over the 100GB-compressed alignment. Missing values indicate that the original alignment could not be compressed to that threshold, or was already below it.

single compute node, and node parallelism only affects the remaining portion of MAGUS. Thus, we first show all of the guide tree runtimes separately in Fig S4, and the full end-to-end runtime comparison with FastTree and Clustal guide trees in Fig S5.

Fig S3. RNASim alignment error, MAGUS variants only. Error is the average of SPFP and SPFN. 'X' markers indicate that compression was used (MAGUS alignments above 100GB). 'RFT' denotes that FastTree guide trees were used in recursive subalignments. MAGUS was run with 100 subsets on RNASim to reduce load on Blue Waters.

Guide Tree Runtimes

Firstly, we note that MAGUS 1 wasn't able to proceed past 50,000 sequences, due to out-of-memory issues. MAGUS 1's accuracy is about 0.5% better than MAGUS(Recurse, FastTree) on 10,000 sequences. Notably, where it did finish on 10-50,000 sequences, MAGUS 1 is much faster than MAGUS(Recurse, FastTree) on a single node - about 23 hours vs. about 74 hours.

Comparing recursive MAGUS with different decompositions, we see that the accuracy of using PartTree is about the same as decomposing randomly, with both being predictably the fastest; the PartTree decomposition takes about 4 hours on a million sequences, while the random decomposition is nearly instant. Next, we see that

Fig S4. RNASim runtimes (in hours) of MAGUS guide trees only.

using FastTree(no ML) yields about the same accuracy as using Clustal (about 1% better than random), but takes more time on larger datasets: on the full million sequences, it takes about 76 hours to compute the FastTree(no ML) tree, vs. about 14 hours to compute the Clustal Omega tree. Finally, we see that using FastTree seems to give the best accuracy overall, usually giving a small advantage (0-0.5%) over Clustal/FastTree(no ML). Obviously, this is also the most expensive option: the tree takes about 5 days to compute on a million sequences.

From these results, we glean two natural guide tree choices for MAGUS. Using FastTree will constitute our “slowest-but-most-accurate” option, while using Clustal Omega becomes our “fast” option. We choose Clustal Omega, because the resulting accuracy is considerably better than PartTree/random, while the runtime is considerably better than FastTree(no ML), giving us the most practical compromise between speed and accuracy. Consequently, we carry these two variants of MAGUS into our next experiment.

The Impact of Node-Parallelism

Fig S5 compares the runtime behavior of these two variants, also demonstrating the effects of node parallelism on MAGUS(Recurse, FastTree). On a million sequences, computing a FastTree guide tree takes about 5 days, and the alignment stage on a single node takes about 9 more days. This reduces to about 17 hours on 10 nodes, and about 2.5 hours on 100 nodes. Using Clustal takes about 14 hours for the guide tree, and then about 5 more days to align on a single node.

Fig S5. RNASim runtimes (in hours), MAGUS variants only. MAGUS was run with 100 subsets on RNASim to reduce load on Blue Waters. Compute nodes were restricted to a maximum walltime of 7 days, but MAGUS(Recurse, FastTree, single node) was run to completion by restarting when time ran out.

Additional Figures

59

Fig S6. RNASim SPFN alignment error, MAGUS variants only. Dashed lines indicate that lossy compression was used (MAGUS alignments above 100GB). MAGUS was run with 100 subsets on RNASim to reduce load on Blue Waters.

Fig S7. RNASim SPFP alignment error, MAGUS variants only. Dashed lines indicate that lossy compression was used (MAGUS alignments above 100GB). MAGUS was run with 100 subsets on RNASim to reduce load on Blue Waters.

Fig S8. RNASim SPFN alignment error, all methods. Dashed lines indicate that lossy compression was used (MAGUS alignments above 100GB). MAGUS was run with 100 subsets on RNASim to reduce load on Blue Waters.

Fig S9. RNASim SPFP alignment error, all methods. Dashed lines indicate that lossy compression was used (MAGUS alignments above 100GB). MAGUS was run with 100 subsets on RNASim to reduce load on Blue Waters.

Fig S10. 16S SPFN alignment error, all methods. MAGUS was run with the default 25 subsets.

Fig S11. 16S SPFP alignment error, all methods. MAGUS was run with the default 25 subsets.

Fig S12. HomFam (smallest 9 datasets) alignment error, all methods. Error is the average of SPFP and SPFN. MAGUS was run with the default 25 subsets. (Errors on “hla” are near-zero.)

Fig S13. Homfam (largest 10 datasets) alignment error, all methods. Error is the average of SPFP and SPFN. MAGUS was run with the default 25 subsets. Muscle segfaulted on the two largest datasets.

Fig S14. Homfam (smallest 9 datasets) runtime, all methods. MAGUS was run with the default 25 subsets.

	Muscle	Clustal	MAFFT	UPP	UPP-F	PASTA	MAGUS 1	MAGUS(R,F)	MAGUS(R,C)
gluts	0.369	0.078	0.174	0.268	0.326	0.179	0.049	0.059	0.057
myb-DNA-binding	0.138	0.097	0.081	0.061	0.102	0.089	0.089	0.077	0.068
tRNA-synt-2b	0.529	0.552	0.454	0.379	0.420	0.417	0.359	0.369	0.384
biotin-lipoyl	0.062	0.060	0.061	0.101	0.085	0.075	0.075	0.073	0.076
hom	0.069	0.053	0.081	0.049	0.109	0.035	0.034	0.043	0.038
ghf13	0.732	0.379	0.374	0.282	0.303	0.304	0.244	0.256	0.330
aldoserred	0.480	0.151	0.089	0.087	0.110	0.112	0.069	0.072	0.085
hla	0.031	0.000	0.000	0.000	0.003	0.000	0.000	0.000	0.000
Rhodanese	0.458	0.509	0.371	0.298	0.295	0.305	0.210	0.239	0.238
PDZ	0.534	0.168	0.173	0.176	0.129	0.178	0.117	0.101	0.139
blmb	0.817	0.518	0.375	0.238	0.349	0.305	0.211	0.219	0.255
p450	0.909	0.422	0.340	0.385	0.300	0.369	0.205	0.217	0.246
adh	0.191	0.013	0.017	0.321	0.324	0.156	0.148	0.152	0.018
aat	0.738	0.221	0.199	0.150	0.142	0.243	0.124	0.122	0.205
rrm	0.525	0.421	0.215	0.199	0.202	0.269	0.184	0.207	0.213
Acetyltransf	0.652	0.491	0.461	0.410	0.441	0.325	0.281	0.355	0.409
sdr	0.692	0.497	0.413	0.268	0.279	0.271	0.240	0.246	0.281
zf-CCHH	segfault	0.314	0.116	0.113	0.123	0.137	0.139	0.134	0.137
rvp	segfault	0.337	0.233	0.217	0.190	0.235	0.184	0.234	0.214
Average	0.466	0.272	0.228	0.216	0.231	0.214	0.155	0.165	0.179

Table S3. SP Error over all homfam datasets. MAGUS(R,F) and MAGUS(R,C) indicate recursive MAGUS with FastTree and Clustal, respectively. Error is the average of SPFP and SPFN. MAGUS was run with the default 25 subsets. Values in bold show the best performer on each dataset. The average is shown over the datasets where all methods completed.

	Muscle	Clustal	MAFFT	UPP	UPP-F	PASTA	MAGUS 1	MAGUS(R,F)	MAGUS(R,C)
gluts	0.399	0.093	0.194	0.475	0.524	0.304	0.048	0.061	0.059
myb-DNA-binding	0.164	0.140	0.076	0.096	0.142	0.118	0.130	0.118	0.106
tRNA-synt-2b	0.522	0.626	0.509	0.543	0.578	0.552	0.498	0.511	0.524
biotin-lipoyl	0.054	0.064	0.064	0.138	0.124	0.106	0.107	0.115	0.099
hom	0.088	0.070	0.080	0.071	0.193	0.052	0.050	0.064	0.054
ghf13	0.785	0.433	0.405	0.382	0.434	0.402	0.358	0.384	0.494
aldoserred	0.588	0.213	0.091	0.129	0.157	0.146	0.086	0.099	0.127
hla	0.042	0.000	0.000	0.000	0.006	0.000	0.000	0.000	0.000
Rhodanese	0.544	0.607	0.393	0.386	0.404	0.407	0.298	0.350	0.360
PDZ	0.610	0.196	0.184	0.220	0.189	0.207	0.169	0.157	0.181
blmb	0.876	0.637	0.411	0.319	0.441	0.357	0.292	0.311	0.364
p450	0.947	0.540	0.388	0.619	0.357	0.576	0.244	0.274	0.295
adh	0.192	0.013	0.017	0.639	0.642	0.221	0.277	0.286	0.024
aat	0.809	0.304	0.218	0.184	0.176	0.291	0.168	0.167	0.255
rrm	0.631	0.577	0.227	0.255	0.266	0.325	0.226	0.251	0.266
Acetyltransf	0.745	0.634	0.484	0.632	0.662	0.481	0.462	0.561	0.621
sdr	0.750	0.623	0.457	0.425	0.438	0.401	0.366	0.382	0.429
zf-CCHH	segfault	0.399	0.143	0.186	0.192	0.205	0.192	0.197	0.205
rvp	segfault	0.388	0.250	0.242	0.244	0.241	0.190	0.254	0.226
Average	0.515	0.339	0.247	0.324	0.337	0.291	0.222	0.241	0.251

Table S4. SPFN Error over all homfam datasets. MAGUS(R,F) and MAGUS(R,C) indicate recursive MAGUS with FastTree and Clustal, respectively. MAGUS was run with the default 25 subsets. Values in bold show the best performer on each dataset. The average is shown over the datasets where all methods completed.

	Muscle	Clustal	MAFFT	UPP	UPP-F	PASTA	MAGUS 1	MAGUS(R,F)	MAGUS(R,C)
gluts	0.338	0.062	0.153	0.061	0.128	0.055	0.050	0.057	0.056
myb-DNA-binding	0.111	0.053	0.085	0.026	0.061	0.060	0.048	0.035	0.030
tRNA-synt-2b	0.537	0.478	0.399	0.215	0.262	0.281	0.219	0.227	0.243
biotin-lipoyl	0.070	0.057	0.058	0.064	0.046	0.044	0.043	0.031	0.052
hom	0.049	0.037	0.082	0.027	0.026	0.019	0.018	0.022	0.021
ghfl3	0.679	0.325	0.344	0.182	0.172	0.206	0.130	0.128	0.165
aldosered	0.372	0.090	0.087	0.045	0.064	0.077	0.052	0.044	0.042
hla	0.020	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Rhodanese	0.371	0.411	0.348	0.210	0.185	0.204	0.123	0.128	0.116
PDZ	0.458	0.140	0.162	0.133	0.069	0.150	0.065	0.045	0.096
blmb	0.757	0.400	0.339	0.157	0.257	0.252	0.129	0.126	0.145
p450	0.871	0.305	0.292	0.151	0.243	0.162	0.167	0.159	0.197
adh	0.191	0.012	0.017	0.003	0.006	0.091	0.019	0.018	0.012
aat	0.666	0.138	0.180	0.117	0.109	0.195	0.080	0.077	0.155
rrm	0.420	0.264	0.204	0.143	0.138	0.212	0.141	0.163	0.160
Acetyltransf	0.558	0.349	0.438	0.188	0.221	0.169	0.099	0.148	0.197
sdr	0.634	0.371	0.369	0.112	0.120	0.142	0.114	0.109	0.132
zf-CCHH	segfault	0.230	0.090	0.041	0.054	0.069	0.086	0.072	0.068
rvp	segfault	0.286	0.216	0.192	0.136	0.229	0.178	0.214	0.201
Average	0.418	0.205	0.209	0.108	0.124	0.136	0.088	0.089	0.107

Table S5. SPFP Error over all homfam datasets. MAGUS(R,F) and MAGUS(R,C) indicate recursive MAGUS with FastTree and Clustal, respectively. MAGUS was run with the default 25 subsets. Values in bold show the best performer on each dataset. The average is shown over the datasets where all methods completed.