

RONALD GRAHAM:
LAYING THE FOUNDATIONS OF ONLINE OPTIMIZATION

SUSANNE ALBERS

1 ABSTRACT. This chapter highlights fundamental contributions made
2 by Ron Graham in the area of online optimization. In an online
3 problem relevant input data is not completely known in advance but
4 instead arrives incrementally over time. In two seminal papers on
5 scheduling published in the 1960s, Ron Graham introduced the concept
6 of *worst-case approximation* that allows one to evaluate solutions
7 computed online. The concept became especially popular when the
8 term *competitive analysis* was coined about 20 years later. The frame-
9 work of approximation guarantees and competitive performance has
10 been used in thousands of research papers in order to analyze (online)
11 optimization problems in numerous applications.

12 2010 Mathematics Subject Classification: 68M20, 68Q25, 68R99,
13 90B35

14 Keywords and Phrases: Scheduling, makespan minimization, algo-
15 rithm, competitive analysis

16 AN ARCHITECT OF DISCRETE MATHEMATICS

17 Born in 1935, Ron Graham entered university at age 15. Already at that time
18 he was interested in a career in research. He first enrolled at the University
19 of Chicago but later transferred to the University of California at Berkeley,
20 where he majored in electrical engineering. During a four-year Air Force service
21 in Alaska he completed a B.S. degree in physics at the University of Alaska,
22 Fairbanks, in 1958. He moved back to the University of California at Berkeley
23 where he was awarded a M.S. and a Ph.D. degree in mathematics in 1961 and
24 1962, respectively.

25 Immediately after graduation Ron Graham joined Bell Labs. Some friends
26 were afraid that this could be the end of his research but, on the contrary,
27 he built the labs into a world-class center of research in discrete mathematics
28 and theoretical computer science. Ron Graham rose from Member of Technical
29 Staff to Department Head and finally to Director of the Mathematics Center



Figure 1: Ron Graham at work and at leisure. Pictures taken in New Jersey in the late 1060s and mid 1970s, respectively. Printed with the permission of Ron Graham.

30 at Bell Labs. After establishment of AT&T Labs Research he served as the
 31 first Vice President of the Information Sciences Research Lab and later became
 32 the first Chief Scientist of AT&T Labs. After 37 years of dedicated service he
 33 retired from AT&T in 1999. Since then he has held the Jacobs Endowed Chair
 34 of Computer and Information Science at the University of California at San
 35 Diego.

36 Ron Graham is a brilliant mathematician. He has done outstanding work in
 37 Ramsey Theory, quasi-randomness, the theory of scheduling and packing and,
 38 last not least, computational geometry. The “Graham scan” algorithm for
 39 computing the convex hull of a finite set of points in the plane is standard
 40 material in algorithms courses. His creativity and productivity are witnessed
 41 by more than 300 papers and five books. Ron Graham was a very close friend
 42 of Paul Erdős and allowed to look not only after his mathematical papers but
 43 also his income. Together they have published almost 30 articles. Ron Graham
 44 is listed in the *Guinness Book of Records* for the use of the highest number
 45 that ever appeared in a mathematical proof. He has many interests outside
 46 mathematics and, in particular, a passion for juggling. It is worth noting that
 47 he served not only as President of the American Mathematical Society but also
 48 as President of the International Jugglers Association.

49 Ron Graham has received numerous awards. He was one of the first recipients
 50 of the Pólya Prize awarded by the Society for Industrial and Applied Math-
 51 ematics. In 2003 he won the Steele Prize for Lifetime Achievement awarded
 52 by the American Mathematical Society. The citation credits Ron Graham as
 53 “one of the principle architects of the rapid development worldwide of discrete
 54 mathematics in recent years” [2].

55 SCHEDULING AND PERFORMANCE GUARANTEES

56 The technical results presented in this chapter arose from extensive research
 57 on scheduling theory conducted at Bell Labs in the mid 1960s. Even today
 58 they exhibit some remarkable features: (1) They can be perfectly used to teach
 59 the concepts of provably good algorithms and performance guarantees to non-
 60 specialists, e.g., high school students or scientists from other disciplines. (2)
 61 The specific scheduling strategies are frequently used as subroutines to solve
 62 related scheduling problems. (3) The results stimulate ongoing research; some
 63 major problems are still unresolved.

64 Consider a sequence $\sigma = J_1, \dots, J_n$ of jobs that must be scheduled on m
 65 identical machines operating in parallel. Job J_i has a processing time of p_i ,
 66 $1 \leq i \leq n$. The jobs of σ arrive one by one. Each job J_i has to be assigned
 67 immediately and irrevocably to one of the machines without knowledge of any
 68 future jobs J_k , $k > i$. Machines process jobs non-preemptively: Once a machine
 69 starts a job, this job is executed without interruption. The goal is to minimize
 70 the makespan, i.e. the maximum completion time of any job in the schedule
 71 constructed for σ .

72 The scheduling problem defined above is an online optimization problem. The
 73 relevant input arrives incrementally. For each job J_i an algorithm has to make
 74 scheduling decisions not knowing any future jobs J_k with $k > i$. Despite this
 75 handicap, a strategy should construct good solutions. Graham [5] proposed a
 76 simple greedy algorithm. The algorithm is also called *List* scheduling, which
 77 refers to the fact that σ is a list of jobs.

78 ALGORITHM LIST: Schedule each job J_i on a machine that currently
 79 has the smallest load.

80 The load of a machine is the sum of the processing times of the jobs presently
 81 assigned to it.

82 A natural question is, what is the quality of the solutions computed by *List*.
 83 Here Graham introduced the concept of worst-case approximation. For any job
 84 sequence σ , compare the makespan of the schedule constructed by *List* to that
 85 of an optimal schedule for σ . How large can this ratio grow, for any σ ? For-
 86 mally, let $List(\sigma)$ denote the makespan of *List*'s schedule for σ . Furthermore,
 87 let $OPT(\sigma)$ be the makespan of an optimal schedule for σ . We would like to
 88 determine

$$c = \sup_{\sigma} \frac{List(\sigma)}{OPT(\sigma)},$$

89 which gives a worst-case performance guarantee for *List*. In online optimization
 90 such a guarantee is called *competitive ratio*. Following Sleator and Tarjan [8],
 91 an online algorithm A is c -competitive if, for any input, the cost of the solution
 92 computed by A is at most c times that of an optimal solution for that input.
 93 Graham [5] gave an elegant proof that *List* is $(2 - 1/m)$ -competitive, i.e. re-
 94 markably *List* achieves a small constant performance ratio. For the proof, fix an
 95 arbitrary job sequence σ and consider the schedule computed by *List*. Without

96 loss of generality, number the machines in order of non-increasing load. Hence
 97 machine 1 is one having the highest load and defines the makespan. Figure 2
 98 depicts an example. In the time interval $[0, List(\sigma))$ machine 1 continuously
 99 processes jobs. Any other machine j , $2 \leq j \leq m$, first processes jobs and then
 100 may be idle for some time. Let J_{i_0} be the job scheduled last on machine 1. We
 101 observe that in *List*'s schedule J_{i_0} does not start later than the finishing time
 102 of any machine j , $2 \leq j \leq m$, because *List* assigns each job to a least loaded
 103 machine. This implies that the idle time on any machine j , $2 \leq j \leq m$, cannot
 104 be higher than p_{i_0} , the processing time of J_{i_0} . Considering the time interval
 105 $[0, List(\sigma))$ on all the m machines we obtain

$$mList(\sigma) \leq \sum_{i=1}^n p_i + (m-1)p_{i_0}.$$

106 Dividing by m and taking into account that $p_{i_0} \leq \max_{1 \leq i \leq n} p_i$, we obtain

$$List(\sigma) \leq \frac{1}{m} \sum_{i=1}^n p_i + \left(1 - \frac{1}{m}\right) \max_{1 \leq i \leq n} p_i.$$

107 A final argument is that the optimum makespan $OPT(\sigma)$ cannot be smaller
 108 than $\frac{1}{m} \sum_{i=1}^n p_i$, which is the average load on the m machines. Moreover,
 109 obviously $OPT(\sigma) \geq \max_{1 \leq i \leq n} p_i$. We conclude that $List(\sigma) \leq OPT(\sigma) +$
 110 $(1 - 1/m)OPT(\sigma) = (2 - 1/m)OPT(\sigma)$.

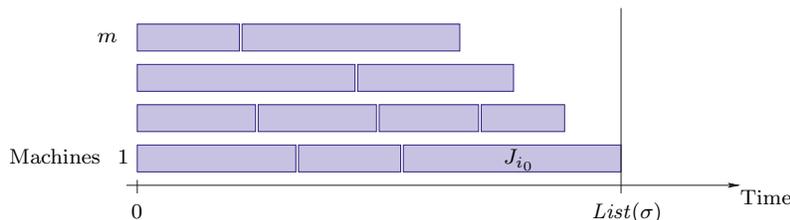


Figure 2: Analysis of *List*

111 Graham [5] also showed that the above analysis is tight. *List* does not achieve
 112 a competitive ratio smaller than $2 - 1/m$. Consider the specific job sequence σ
 113 consisting of $m(m-1)$ jobs of processing time 1 followed by a large job having
 114 a processing time of m . *List* distributes the small jobs evenly among the m
 115 machines so that the final job cause a makespan of $m-1 + m = 2m-1$.
 116 On the other hand the optimum makespan is m because an optimal schedule
 117 will reserve one machine for the large job and distribute the small jobs evenly
 118 among the remaining $m-1$ machines. Figure 3 shows the schedules by *List*
 119 and *OPT*.

120 The above nemesis job sequence motivated Graham to formulate a second algo-
 121 rithm. Obviously *List*'s performance can degrade if large jobs arrive at the end
 122 of the input sequence. Why not sort the jobs initially? Graham [6] proposed

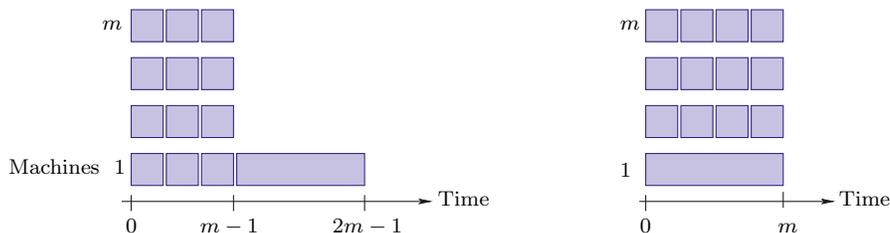


Figure 3: The worst-case performance of *List*. Online schedule (left) and an optimal schedule (right).

123 a *Sorted List* algorithm that first sorts the jobs in order of non-increasing pro-
 124 cessing time and then applies *List* scheduling. Of course *Sorted List* is not an
 125 online algorithm because the entire job sequence must be known and rearranged
 126 in advance.

127 Graham [6] proved that *Sorted List* achieves a worst-case approximation ratio of
 128 $4/3 - 1/(3m)$. The analysis is more involved than that of *List* but the global idea
 129 can be described in one paragraph: Consider an arbitrary sorted job sequence σ
 130 and assume without loss of generality that the last job of σ defines *Sorted List*'s
 131 makespan. If this is not the case, then one can consider the prefix sequence
 132 σ' such that the last job of σ' defines *Sorted List*'s makespan for σ' and σ .
 133 It suffices to consider two cases. (1) If the last job J_n of σ has a processing
 134 time p_n of at most $OPT(\sigma)/3$, then using the same arguments as above one
 135 can establish a performance factor of $4/3 - 1/(3m)$. (2) If $p_n > OPT(\sigma)/3$,
 136 then all jobs of σ have a processing time greater than $OPT(\sigma)/3$. Hence in
 137 an optimal schedule each machine can contain at most two jobs and $n \leq 2m$.
 138 Assume for simplicity $n = 2m$. One can show that there exists an optimal
 139 schedule that pairs the largest with the smallest job, the second largest with
 140 the second smallest job and so on. That is, the pairing on the m machines is
 141 $(J_1, J_{2m}), (J_2, J_{2m-1}), \dots, (J_m, J_{m+1})$. If $n = 2m - k$, for some $k \geq 1$, then
 142 there is an optimal schedule that is identical to the latter pairing except that
 143 J_1, \dots, J_k are not combined with any other job. *Sorted List* produces a schedule
 144 that is no worse than this optimal assignment, i.e., in this case the performance
 145 ratio is equal to 1.

146 The above results led to a considerable body of further research. It was open
 147 for quite some time if online algorithms for makespan minimization can attain
 148 a competitive ratio smaller than $2 - 1/m$. It turned out that this is indeed pos-
 149 sible. Over the past 20 years the best competitiveness of deterministic online
 150 strategies was narrowed down to $[1.88, 1.9201]$. More specifically, there exists a
 151 deterministic online algorithm that is 1.9201-competitive, and no deterministic
 152 online strategy can attain a competitive ratio smaller than 1.88. If job pre-
 153 emption is allowed, i.e., the processing of a job may be stopped and resumed
 154 later, the best competitiveness drops to $e/(e-1) \approx 1.58$. The book chapter [7]
 155 contains a good survey of results.

156 During the last few years researchers have explored settings where an online
157 algorithm is given extra information or ability to serve the job sequence. For in-
158 stance, an online algorithm might be able to migrate a limited number of jobs or
159 alternatively might know the total processing time of all jobs in σ . In these sce-
160 narios significantly improved performance guarantees can be achieved. Using
161 limited job migration, the competitiveness reduces to approximately 1.46. The
162 recent manuscript [1] points to literature for these extended problem settings.
163 Nonetheless a major question is still unresolved. What is the best competi-
164 tive ratio that can be achieved by randomized online algorithms? It is known
165 that no randomized online strategy can attain a competitiveness smaller than
166 $e/(e - 1)$. However, despite considerable research interest, no randomized on-
167 line algorithm that provably beats deterministic ones, for general m , has been
168 devised so far.

169 Finally, as mentioned above, the design and analysis of online algorithms has
170 become a very active area of research. We refer the reader to two classical
171 books [3, 4] in this field.

172 REFERENCES

- 173 [1] S. Albers and M. Hellwig. On the value of job migration in online makespan
174 minimization. CoRR abs/1111.0773, 2012.
- 175 [2] AMS document about the 2003 Steele Prize. Accessible at [http://en.
176 wikipedia.org/wiki/Ronald_Graham](http://en.wikipedia.org/wiki/Ronald_Graham).
- 177 [3] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Anal-*
178 *ysis*. Cambridge University Press, 1998.
- 179 [4] A. Fiat and G.J. Woeginger (eds). *Online Algorithms: The State of the*
180 *Art*. Springer LNCS 1442, 1998.
- 181 [5] R.L. Graham. Bounds for certain multi-processing anomalies. *Bell System*
182 *Technical Journal*, 45:1563–1581, 1966.
- 183 [6] R.L. Graham. Bounds on multiprocessing timing anomalies. *SIAM Journal*
184 *of Applied Mathematics*, 17(2):416–429, 1969.
- 185 [7] K. Pruhs, J. Sgall and E. Torng. Online scheduling. *Handbook on Schedul-*
186 *ing*, edited by J. Y-T. Leung. Chapman & Hall / CRC. Chapter 15, 2004.
- 187 [8] D.D. Sleator and R.E. Tarjan. Amortized efficiency of list update and
188 paging rules. *Communications of the ACM*, 28:202–208, 1985.

Susanne Albers
Department of Computer Science
Humboldt-Universität zu Berlin
Unter den Linden 6
10099 Berlin
Germany
`albers@informatik.hu-berlin.de`

