[5] Jones, D.R., and M.A. Beltramo, Solving partitioning problems with genetic algorithms, in Belew, R.K., and L.B. Booker (Eds.), *Proceedings of the Fourth International Conference on Genetic Algorithms*, Morgan Kaufmann, San Mateo, Ca., 1991, pp. 442-449.

[6] Karmarkar, N., and R.M. Karp, The differencing method of set partitioning, Technical Report UCB/CSD 82/113, Computer Science Division, University of California, Berkeley, Ca., 1982.

[7] Karmarkar, N., R.M. Karp, G.S. Lueker, and A.M. Odlyzko, Probabilistic analysis of optimum partitions, Journal of Applied Probability, Vol. 23, 1986, pp. 626-645.

[8] Mitchell, D., B. Selman, and H. Levesque, Hard and easy distributions of SAT problems, *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-94)*, San Jose, Ca., July 1992, pp. 459-465.

[9] Ruml, W., J.T. Ngo, J. Marks, S. Shieber, Easily searched encodings for number partitioning, to appear, *Journal of Optimization Theory and Applications*, vol 89, number 2, 1996.

[10] Gent, I.P., and T. Walsh, The number partition phase transition, Technical Report RR-95-185, Department of Computer Science, University of Strathclyde, Glasgow G1 1XH, Scotland, May, 1995.

[11] Zhang, W., and R.E. Korf, A unified view of complexity transitions on the traveling salesman problem, to appear in *Artificial Intelligence*, 1996.

algorithm, and made it complete, so that the first solution found is the approximation, and then better solutions are found as long as the algorithm continues to run, eventually finding an optimal solution. Thirdly, it represents an example of another approach that may be more broadly applicable. Most algorithms for combinatorial problems construct a solution to a problem incrementally, adding one element at a time to a single partial solution. This is the case with the standard algorithm for number partitioning. The CKK algorithm, on the other hand, constructs a large number of partial solutions, and combines them together in all possible ways. In this case, this latter strategy is much more effective, and may be for other problems as well.

11 Acknowledgements

Thanks to Wheeler Ruml for introducing me to number partitioning, and the Karmarkar-Karp heuristic. Thanks to Wheeler, Ken Boese, Alex Fukunaga, and Andrew Kahng for helpful discussions concerning this research, and to Pierre Hasenfratz for comments on an earlier draft. This work was supported by NSF Grant IRI-9119825, and a grant from Rockwell International.

References

[1] Cheeseman, P., B. Kanefsky, and W.M. Taylor, Where the *really* hard problems are, *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-91)*, Sydney, Australia, Aug. 1991, pp. 331-337.

[2] Garey, M.R., and D.S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W.H. Freeman, San Francisco, 1979.

[3] Huberman, B., and T. Hogg, Phase transitions in artificial intelligence systems, *Artificial Intelligence*, Vol. 33, No. 2, Oct. 1987, pp. 155-171.

[4] Johnson, D.S., C.R. Aragon, L.A. McGeoch, and C. Schevon, Optimization by simulated annealing: An experimental evaluation; Part II, graph coloring and number partitioning, *Operations Research*, Vol. 39, No. 3, 1991, pp. 378-406.

9 0/1 Knapsack Problem

Given a set of integers, and a constant c, the 0/1 knapsack problem is to find a subset of the integers whose sum is exactly c. We can reduce this problem to the two-way partition problem, and hence apply CKK to this problem as well. Let s be the sum of all the integers. If c is less than $s/2$, use $s-c$ for c. Add a new integer d such that $(s+d)/2 = c$, or $d = 2c-s$. If this augmented set can be perfectly partitioned with a difference of zero, then the subset of the perfect partition that does not contain d is a subset of the original numbers whose sum is exactly c, and hence a solution to the 0/1 knapsack problem. Conversely, if this augmented set cannot be perfectly partitioned, then there is no subset of the original numbers that sum to exactly c, and hence no solution to the 0/1 knapsack problem.

10 Summary and Conclusions

The main contribution of this paper is to extend an elegant and effective polynomial-time approximation algorithm due to Karmarkar and Karp, to a complete algorithm (CKK). The first solution it finds is the KK solution, and as it continues to run it finds better solutions, until it eventually finds and verifies an optimal solution. For problems without perfect partitions, CKK appears to be asymptotically more efficient than the standard optimal algorithm. When a perfect partition exists, CKK outperforms the standard algorithm by orders of magnitude. We showed results for both two-way and three-way partitioning. In practice, two-way partitioning problems of arbitrary size can be solved if the numbers are restricted to no more than twelve significant digits of precision, while arbitrary-sized three-way partitioning problems can be optimally solved with six significant digits. For large problems with higher precision values, CKK can be run as long as time is available, returning the best solution found when time runs out. CKK outperforms every algorithm we could find in the literature for this problem.

What contribution does this work make beyond the specific problem of number partitioning? First, CKK is directly applicable to the 0/1 knapsack problem, and may apply to other related problems as well. Secondly, it presents an example of an approach that may be effective on other combinatorial problems. Namely, we took a good polynomial-time approximation

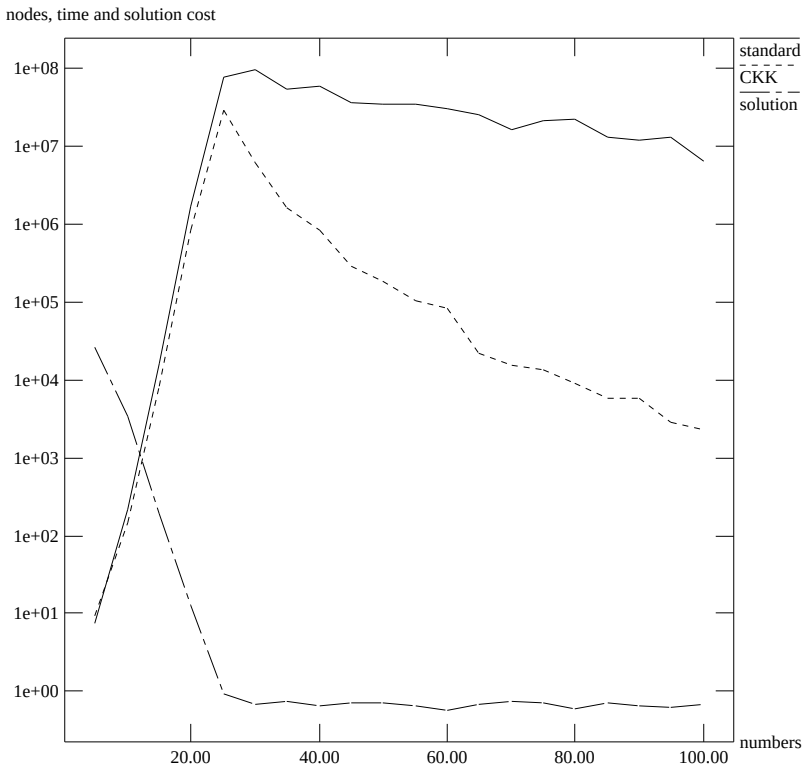

Figure 8: Nodes generated to partition 5-digit numbers into three subsets

8 Stochastic Approaches

There have been at least three studies applying stochastic algorithms to number partitioning, none of which can guarantee optimal solutions. Johnson et al.[4] applied simulated annealing to the problem, but found that it was not competitive with even the Karmarkar-Karp heuristic solution. Ruml et al.[9] applied various stochastic algorithms to some novel encodings of the problem, but their best results outperform the KK solution by only three orders of magnitude, compared to the six orders of magnitude CKK achieves in less than an hour. Jones and Beltramo[5] applied genetic algorithms to the problem, but don't mention the Karmarkar-Karp heuristic. Their technique fails to find an optimal solution to the single problem instance they ran, while the KK solution to this instance is optimal.

and the sum of the remaining values is $8+5+4 = 17$. If we divide 17 among the two remaining subsets, the best we could do is to have 9 in one of the subsets, and 8 in the other. The best possible partition difference would then be $13 - 8 = 5$. Since the leftmost leaf node has a partition difference of only 3, we can prune this node. Finally, a complete partition with a difference of zero or one is perfect, and terminates the search.

7.5 Experimental Results

We implemented the above algorithms for three-way partitioning, using numbers uniformly distributed from zero to $100,000$. Figure 8 shows the results, in a format identical to that of Figure 5. Each data point is an average of 100 random trials. The horizontal axis is the number of values being partitioned, and the vertical axis for the standard and CKK algorithms is the number of nodes generated. We also show the value of the optimal subset difference on the same scale, indicating that the hardest problems occur where the probability of a perfect partition is about one-half. The results are very similar to those for two-way partitioning. Namely, CKK appears asymptotically more efficient than the standard algorithm when no perfect partition exists, and is orders of magnitude more efficient when there are perfect partitions.

While the constant factors for CKK and the standard algorithm are similar for two-way partitioning, the three-way version of CKK is more complex. Our three-way implementation of CKK runs about 33% slower per node generation than the standard algorithm. While this reduces the absolute performance of CKK, it still appears to be asymptotically more efficient than the standard algorithm, and runs faster in practice.

In order to run large numbers of three-way partitioning problems of different sizes, we used numbers with five significant decimal digits. Single instances of arbitrary size with six digits of precision can be solved in practice, however. While two-way partitioning problems with up to twelve digits can be optimally solved, three-way partitioning is computationally more difficult, since the number of k-way partitions is $O(k^n/k!)$.

keep the numbers in different subsets apart in every partition arising from it. A subset sum of zero represents an empty set, since we don't normalize. Each node takes the two subpartitions with the largest subset sums, and combines them together in all possible ways. For example, the root node in Figure 7 combines the two subpartitions containing just the 8 and the 7 into a single subpartition in two different ways. Either the 8 and 7 will go into different subsets, or in the same subset. Since all the original numbers must eventually be combined into a single partition of k subsets, combining the subpartitions pairwise in all possible ways guarantees that every partition will eventually be generated at some leaf node of the tree.

To see that no partition is generated more than once, note that each node combines only two subpartitions, the ones with the largest subset sums. The new subpartitions that result are all different, representing different choices about assigning the numbers in the combined subpartitions to the same or different subsets. For example, consider the root of the tree in Figure 7. Its left child puts the 8 and 7 in different subsets, and they will stay in different subsets in every partition generated below that node, since numbers in different subsets of the same subpartition are never combined. Conversely, the right child of the root puts the 8 and 7 in the same subset, and they will stay in the same subset in every partition below that node. Thus, the partitions generated below the left and right children are completely disjoint. This is true in general, and each distinct partition appears only once.

Normalization can reduce the size of the tree. For example, if the two smallest subsets in a subpartition have the same sum, after normalization a zero will occur in the combined subpartition, which does not represent an empty subset. When another value is combined with this subpartition, it will only be placed in one of the smallest subsets. Since the two subsets have the same sum, which subset a new value is placed in has no affect on the final partition difference, even though it may generate different partitions. This normalization savings also applies to the standard algorithm as well.

Pruning the CKK tree is similar to pruning in the standard algorithm. We add up all the subset sums except for the largest one, and evenly divide this total among $k-1$ subsets. If the difference between the largest subset sum and this quotient is greater than or equal to the best complete partition difference found so far, we can prune this branch, since this is the best we could possibly do below that node. For example, consider the subpartition $(13,8)(5,0)(4,0)$, near the middle of Figure 7. The largest subset sum is 13,

the two subpartitions with the largest normalized subset sums are combined in each possible way, and the resulting subpartitions are normalized and inserted in the sorted order of remaining subpartitions. The resulting child nodes are then searched in increasing order of the largest normalized subset sum. Thus, the first solution found is the KK solution.

Figure 7 shows the tree that is generated to optimally partition the numbers $(4, 5, 6, 7, 8)$ into three subsets. At the root, each original number is in its own subpartition. Each subpartition contains only two subset sums, since the third is normalized to zero. At each node, the two subpartitions with the largest subset sums are combined in all possible ways. The nodes at depth $n-1$ are complete partitions, and their subset difference is the largest subset sum, since the smallest is zero. Pruning is discussed below.

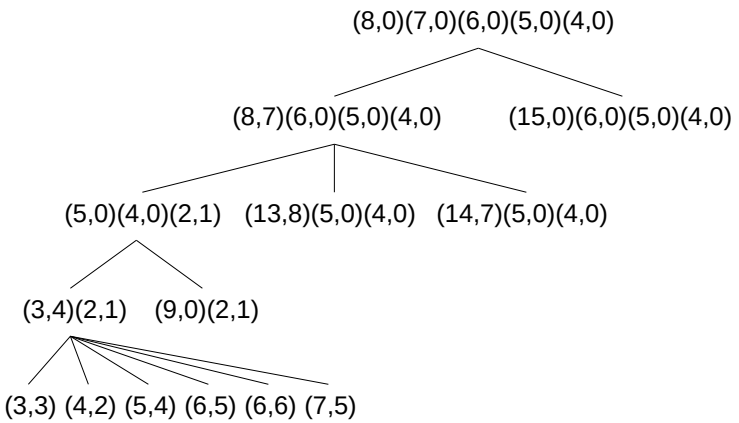

Figure 7: Tree generated to partition 4,5,6,7,8 into three subsets

The resulting tree has depth n, and for three-way partitioning, nodes with branching factors two, three, and six, depending on whether the two subpartitions being combined have two, three, or four non-zero subset sums, respectively. The number of leaf nodes, however, is the same as in the standard tree, assuming no normalization or pruning. In other words, if all k subset sums are maintained, and every branch proceeds to depth n, there is one leaf node for every distinct partition of the original numbers.

To see this, note that every node in the tree represents a complete collection of subpartitions. A single subpartition is a division of some of the original numbers into k different subsets, and represents a commitment to

these last two subpartitions are combined by merging the largest subset sum with the smallest $(5 + 0)$, the smallest with the largest $(0 + 2)$, and the two medium subset sums together $(4 + 1)$, yielding $(5, 5, 2)$. Subtracting the smallest from all three subset sums results in the final subpartition of $(3, 3, 0)$, which has a difference of 3, and happens to be optimal in this case. While we have shown all three subset sums for clarity, our implementation only maintains the two normalized non-zero subset sums for each subpartition.

The actual partition is reconstructed as follows. Each subset sum in each subpartition represents a set of original numbers. Whenever we combine two subset sums together, we merge the corresponding sets. For example, in the state $((5, 4, 0)(2, 1, 0))$ above, the 5 represents the original 5, the 4 represents the 4, the 2 represents the 8, the 1 represents the 7, and the 0 in $(2, 1, 0)$ represents the 6, since we subtracted 6 from each of the subset sums in this subpartition. At the last step, we combine the 5 with the 0, resulting in the set $(5, 6)$, the 4 with the 1, resulting in the set $(4, 7)$, and the 2 with the 0, resulting in the singleton set (8). Thus, the final partition is $((8)(7, 4)(6, 5))$, with subset sums of 8, 11, and 11, and a final difference of 3.

7.4 Complete Karmarkar-Karp Algorithm

The CKK algorithm also generalizes to multi-way partitioning. Instead of combining subpartitions in only one way, to make the algorithm complete we must combine them together in all possible ways. Again consider three-way partitioning. A particular subpartition represents a commitment to keep the elements in the different subsets separate. There are three cases to consider in combining a pair of subpartitions. In the first case, both subpartitions have only a single non-zero subset sum, say $(A, 0, 0)$ and $(X, 0, 0)$. We can combine these in only two different ways, either putting the non-zero elements together, $(X + A, 0, 0)$, or keeping them apart, $(X, A, 0)$. In the second case, one subpartition has one non-zero subset sum and the other has two, say $(A, 0, 0)$ and $(X, Y, 0)$. In this case we can combine them in three different ways, putting the single non-zero element in any of the three subsets of the other subpartition, resulting in the subpartitions (X, Y, A), $(X, Y + A, 0)$, and $(X + A, Y, 0)$. Finally, both subpartitions can have two non-zero subset sums, say $(A, B, 0)$ and $(X, Y, 0)$. In this case, there are six different ways to combine them: $(X, Y + B, A)$, $(X, Y + A, B)$, $(X + B, Y, A)$, $(X + A, Y, B)$, $(X + B, Y + A, 0)$, and $(X + A, Y + B, 0)$. At each step of the algorithm,

than the largest subset sum, then the difference between them is the best possible final difference we could achieve. There is no guarantee that we can actually achieve this, since it represents a perfect solution to a $k-1$-way partitioning problem, but it is a lower bound. If the resulting difference is greater than the best complete partition difference found so far, we can prune this branch of the tree, since we can't improve on the existing partition.

Formally, let $s_1, s_2, \ldots, s_k$ be the current subset sums, let s_1 be the largest of these, and let r be the sum of the remaining unassigned numbers. If

$$s_1 - \frac{\sum_{i=2}^{k} s_i + r}{k-1}$$

is greater than or equal to zero, then this is the best possible completion of this partial partition. If it is greater than or equal to the best complete partition difference found so far, then the corresponding branch is pruned.

Finally, a perfect partition will have a difference of zero if the sum of the original numbers is divisible by k, and a difference of one otherwise. Once a perfect partition is found, the algorithm terminates.

7.3 Karmarkar-Karp Heuristic

Karmarkar and Karp also generalized their set differencing method to k-way partitioning. A state is represented by a collection of subpartitions, each with k subset sums. The initial numbers are each represented by a subpartition with the number itself in one subset, and the remaining subsets empty. For example, a three-way partitioning of the set $(4,5,6,7,8)$ would initially be represented by the subpartitions $((8,0,0),(7,0,0),(6,0,0),(5,0,0),(4,0,0))$, which are sorted in decreasing order. The two largest numbers are combined into a single subpartition by putting them in different subsets, resulting in the list $((8,7,0),(6,0,0),(5,0,0),(4,0,0))$. The new subpartition still has the largest subset sum, and hence the next smaller subpartition, $(6,0,0)$, is combined with it by placing the 6 in the smallest subset, resulting in the subpartition $(8,7,6)$. Since we are only interested in the difference between the largest and the smallest subset sums, we normalize by subtracting the smallest sum, 6, from each of the subsets, yielding the subpartition $(2,1,0)$. This subpartition is then inserted into the remaining sorted list in decreasing order of largest subset sum, resulting in $((5,0,0)(4,0,0)(2,1,0))$. Again, the two largest subpartitions are combined, yielding $((5,4,0)(2,1,0))$. Finally,

7 Multi-Way Partitioning

So far, we have considered partitioning into two subsets. Here we discuss the generalization of these techniques to partitioning into multiple subsets. The task is to partition a set of numbers into k mutually exclusive and collectively exhaustive subsets, so that the difference between the largest and smallest subset sums is minimized. This also minimizes the largest subset sum.

7.1 Greedy Algorithm

The generalization of the greedy heuristic to k-way partitioning is straightforward. We sort the numbers in decreasing order, and always place the next unassigned number in the subset with the smallest sum so far, until all the numbers have been assigned. Since we can always subtract the smallest subset sum from all the others without affecting the final partition difference, we only have to maintain $k-1$ values, the normalized differences between each subset sum and the smallest one.

7.2 Standard Complete Algorithm

The standard complete algorithm also readily generalizes to k-way partitioning. Again we sort the numbers in decreasing order, and each number in turn is placed in each of k different subsets, generating a k-ary tree. The left-most branch places the next number in the subset with the smallest sum so far, the next branch places it in the next larger subset, etc. Thus, the first solution found is the greedy solution. By never placing a number in more than one empty subset, we avoid generating duplicate partitions that differ only by a permutation of the subsets, and produce all $O(k^n/k!)$ distinct k-way partitions of n elements. Furthermore, by never placing a number in two different subsets with the same subset sum, we avoid generating different partitions that have the same partition difference.

To prune the tree, we use branch-and-bound, and maintain the smallest difference found so far for a complete partition. Given the largest current subset sum, the best we can do is to bring each of the remaining subset sums up to the value of the largest. To see if this is possible, we add all the subset sums except the largest together, add to this the sum of the remaining unassigned numbers, and divide the result by $k-1$. If this quotient is less

For large problems with very high precision values, we must settle for approximate solutions. In that case, we can run CKK for as long as time allows, and return the best solution found. The first solution found is the KK solution, and as CKK continues to run, it finds better solutions. This technique is very effective, with much of the improvement in solution quality occurring early in the run. Figure 6 shows the average improvement as a function of running time for 100 trials of partitioning 50 48-bit numbers. The horizontal axis is the number of nodes generated, and the vertical axis is the ratio of the initial KK solution to the best solution found in the given number of node generations, both on a logarithmic scale. The entire horizontal scale represents about an hour of real time, and shows over six orders of magnitude improvement, relative to the initial KK solution. Three orders of magnitude improvement is obtained in less than a second of running time.

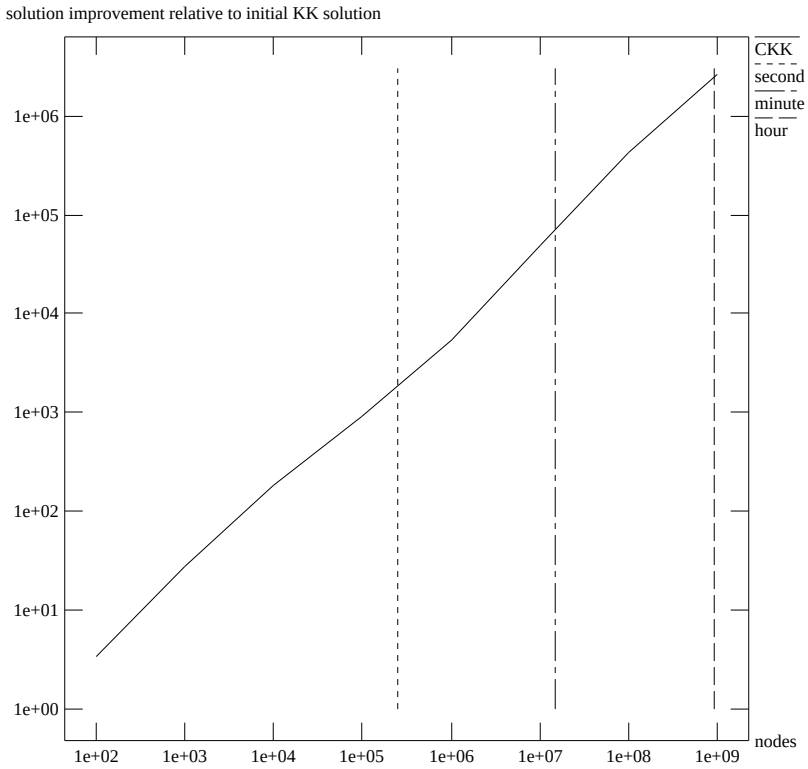

Figure 6: Solution quality relative to KK solution for 50 48-bit numbers

This complexity transition also appears in optimization problems[1, 11], as long as there exist optimal solutions that can be recognized as such without comparison to any other solutions. This is the case with number partitioning, where a subset difference of zero or one is always an optimal solution. As another example, when searching a problem space with non-negative costs, a solution of cost zero is always optimal, independent of any other solutions.

In most of these problems, the hardest problem instances occur where the probability of an exact or perfect solution is one-half. In our experiments, 38% of random sets of 35 10-digit numbers had a perfect partition, and 63% of problem instances of size 36 could be partitioned perfectly, suggesting that these should be the hardest problems. In fact, the problem instances that generate the largest median number of nodes are those of size 36. If we look at mean node generations instead, the hardest problems are of size 38, since the outliers have a larger effect on the mean than the median. See [10] for more detail on this complexity transition in number partitioning.

Ideally, we would like to predict where the hardest problems are, for a given precision of values. To do this, we need to know the value of the optimal subset difference for a given problem class. Karmarkar and Karp et al[7] showed that for a set of independent trials of partitioning n real numbers from zero to one, the median value of the minimum subset difference is $\Theta(\sqrt{n}/2^n)$. Since we are using integers, we simply multiply this by the maximum value m of a number. For example, $m = 10^{10}$ for the experiments in Figure 5. If a majority of a set of problem instances yield a perfect partition, then the median value of the optimal subset differences will almost certainly be 1. Thus, to find the hard problems for a given range m, we solve for n in the equation $\sqrt{n}m/2^n = 1$. For example, if $m = 10^{10}$ and $n = 35$, $\sqrt{n}m/2^n = 1.72$, and for $n = 36$, $\sqrt{n}m/2^n = .87$. Thus, for 10 digit integers, the hardest problems should be of size 35 or 36, which is what we found.

6 Finding Approximate Solutions

While this formula tells us where to find the hardest problems for a given precision of values, the easiest way to generate hard problems is to increase the precision. Most of the work on number partitioning has focussed on problems without perfect partitions. To generate large such problem instances, numbers with up to 36 decimal digits have been used[9].

The intuition behind the observation that large problems are easier to solve than those of intermediate size is quite simple. Given n integers, the number of different subsets of those integers is 2^n. If the integers range from 0 to m, the number of different possible subset sums is less than nm, since nm is the maximum possible subset sum. If m is held constant while n is increased, the number of different subsets grows exponentially, quickly exceeding the number of different subset sums, which grow only linearly. Thus, there are likely to be many subsets with any given subset sum. If there is a subset whose sum equals half the sum of all the elements, then we have a perfect partition, and the density of perfect partitions increases with increasing n, making them easier to find.

The data in Figure 5 are for numbers with ten decimal digits of precision, to allow running many trials with different numbers of values. Arbitrary-size single problem instances with up to twelve digits of precision can be solved by CKK in less than a day on average. For example, this would represent an accuracy of one second in over 30,000 years. Since no physical quantities are known with higher precision, any two-way partitioning problems that arise in practice can be optimally solved, regardless of problem size. While all our experiments were run on uniformly distributed values, we believe that the same results will apply to other naturally occurring distributions as well.

5 Where the Hardest Problems Are

Figure 5 shows that for numbers of fixed precision, increasing the problem size makes the problem more difficult up to a point, and then easier beyond that point. The reason for the decrease in problem difficulty with increasing problem size is that perfect partitions become more common with increasing problem size. Once a perfect partition is found, the search is terminated.

This phenomenon has been observed in a number of different constraint-satisfaction problems, such as graph coloring and boolean satisfiability, and has been called a *phase transition*[3, 1, 8]. In a constraint-satisfaction problem, the difficulty increases with problem size as long as no solution exists, since the entire problem space must be searched. For some problems however, as problem size increases further, solutions begin to appear and become increasingly common. In that case, the problem gets easier with increasing size, since once any solution is found, the search is terminated.

Both algorithms were coded in C, and generate approximately 15 million nodes per minute on a SUN SPARC 2. Thus, the entire vertical axis respresents about an hour of computation time. While CKK would seem to require more time per node to insert the computed difference in the remaining list, this amounts to only a small constant factor, since most of the nodes in the tree are near the bottom, where the lists are very short. In fact, the average height of a node in a complete binary tree approaches one as the height of the tree goes to infinity. This constant factor is made up for by the fact that when there are only four numbers remaining, the final partition difference can be computed more efficiently by CKK, since they are in sorted order. For the standard algorithm, the current subset difference is not in sorted order with respect to the remaining three numbers.

There are clearly two different regions of this graph, depending on how many values are partitioned. With less than 30 numbers, no perfect partitions were found, while with 40 or more numbers, a perfect partition was found in every case. The optimal subset difference averages .5 beyond 40 numbers, since there are roughly equal numbers of differences of zero and one. Figure 5 shows that CKK outperforms the standard algorithm over the entire range.

Without a perfect partition, the ratio of the number of nodes generated by the standard algorithm to those generated by CKK grows linearly with the number of values partitioned. This suggests that CKK is asymptotically more efficient than the standard algorithm. In this region, the performance of both algorithms is largely independent of the precision of the numbers being partitioned, except for a constant based on whether single-precision or double-precision arithmetic is used. To optimally partition 40 48-bit double-precision numbers with the standard algorithm requires an average of 9 hours and 16 minutes, while CKK requires an average of only 3 hours and 42 minutes on the same problems, a factor of 2.5 improvement.

The performance improvement is much more dramatic when a perfect partition exists. In that case, CKK finds the perfect partition much sooner than the standard algorithm, and hence terminates the search earlier. As the problem size increases, the running time of the standard algorithm drops gradually, but the running time of CKK drops precipitously, resulting in orders of magnitude improvement. We have run CKK on 10-digit problems up to size 300, and the trend seen in Figure 5 continues, bottoming out at the point where the KK heuristic solution is almost always optimal, and the running time is dominated by the $O(n \log n)$ time to find this first solution.

shown in Figure 5. We chose random numbers uniformly distributed from 0 to 10 billion, which have ten significant decimal digits of precision. Each data point is the average of 100 random problem instances. To make the algorithms more efficient, CKK directly computes the optimal partition when there are four numbers left, since the KK heuristic is optimal in that case. The equivalent stopping point for the standard algorithm is when there are three unassigned numbers remaining, with the current subset difference playing the role of the fourth number. The horizontal axis shows the number of values partitioned, with data points for sets of size 5 to 100, in increments of 5, and from 30 to 40 in increments of 1. The vertical axis shows the number of nodes generated by the two algorithms. The descending line shows the average optimal partition difference on the vertical axis, fortuitously representable on the same scale.

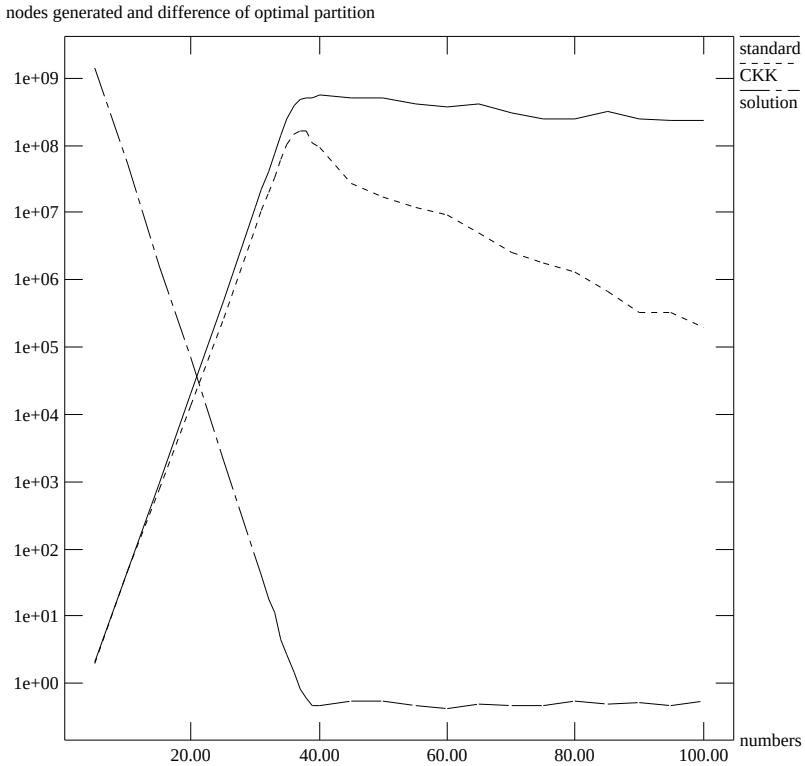

Figure 5: Nodes generated to optimally partition 10-digit numbers

even though both find optimal solutions to the same problem instance.

There are two reasons why CKK is more efficient than the standard complete algorithm. If there is no perfect partition, then both algorithms must search the whole tree. Consider the left subtrees in Figures 3 and 4, where both algorithms place the 8 and 7 in different subsets. This state is represented by $1(6,5,4)$ in Figure 3, where 1 is the current subset difference, and by $(6,5,4,1)$ in Figure 4. The distinction between these two nodes is that in the latter case, the difference of 1 is treated like any other number, and inserted in the sorted order, instead of having the special status of the current subset difference. Thus, at the next level of the tree, represented by nodes $(4,1,1)$ and $(11,4,1)$ in Figure 4, the largest number is greater than the sum of the remaining numbers, and these branches are pruned. In the standard algorithm, however, the two children of the left subtree, $5(5,4)$ and $7(5,4)$ in Figure 3, have to be expanded further. Thus, CKK allows more pruning than the standard algorithm.

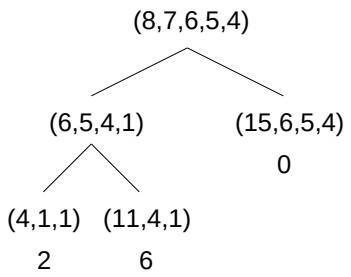

Figure 4: Tree generated by CKK algorithm to partition (4,5,6,7,8)

The second reason that CKK is more efficient occurs when a perfect partition exists. In that case, since the KK heuristic produces better solutions than the greedy heuristic, the CKK algorithm finds better solutions sooner, including the perfect solution. This allows it to terminate the search much earlier than the standard algorithm, on average.

4 Experimental Results: Optimal Solutions

We implemented both the standard complete algorithm and the CKK algorithm, which find optimal solutions. The results for two-way partitioning are

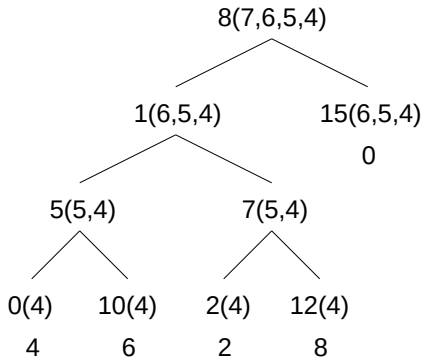

Figure 3: Tree generated by standard algorithm to partition (4,5,6,7,8)

3 Complete Karmarkar-Karp (CKK)

Similar to the extension of the greedy heuristic to a complete algorithm, the main contribution of this paper is to extend the KK heuristic to a complete algorithm. While the idea is extremely simple, it doesn't appear in Karmarkar and Karp's paper[6], and apparently escaped several other researchers who worked on the problem subsequently[9, 4].

At each cycle, the KK heuristic commits to placing the two largest numbers in different subsets, by replacing them with their difference. The only other option is to place them in the same subset, by replacing them by their sum. The resulting algorithm, which we call Complete Karmarkar-Karp (CKK), searches a binary tree, where each node replaces the two largest remaining numbers. The left branch replaces them by their difference, while the right branch replaces them by their sum. The difference is inserted in sorted order in the remaining list, while the sum is simply added to the head of the list, since it will be the largest element. Thus, the first solution found by CKK is the KK solution, but as it continues to run it finds better solutions, until an optimal solution is eventually found and verified.

The worst-case running time of CKK is $O(2^n)$. However, similar pruning rules apply as in the standard algorithm, with the largest element playing the role of the current subset difference. In other words, a branch is terminated when the largest element is greater than or equal to the sum of all the remaining elements. Figure 4 shows the resulting binary tree for the numbers $(4, 5, 6, 7, 8)$. Note that the tree in Figure 4 is smaller than that in Figure 3,

15, so the best we can do is to put all the remaining numbers in the smaller subset. This pruning doesn't depend on any solutions found so far, thus the size of the resulting tree is independent of the order in which it is searched.

If we reach a terminal node whose subset difference is zero or one, representing a perfect partition, then we can terminate the entire search. The above example illustrates this as well, since once we put the remaining numbers in the smaller subset, the resulting complete partition has a difference of zero. If a perfect partition exists, then the search order matters, since the sooner we find it, the sooner we can quit. The obvious way to order the search is to always put the next number in the smaller subset so far, before putting it in the larger subset. This algorithm produces the greedy solution first, and continues to search for better solutions, until an optimal solution is eventually found and verified.

Several additional optimizations deserve mention. One is that the first number should only be assigned to one subset, cutting the search space in half. The second is that whenever the current subset sums are equal, the next number should only be assigned to one subset, cutting the remaining subtree in half. The first optimization is in fact a special case of the second. Finally, when only one unassigned number remains, it should be assigned only to the smaller subset. Figure 3 shows the resulting binary tree for the numbers $(4, 5, 6, 7, 8)$, where the number in front of the parentheses is the difference between the current subset sums, and the numbers below the leaf nodes are the corresponding final partition differences.

2.4 Dynamic Programming Algorithm

There is also a dynamic programming algorithm for finding optimal solutions[2]. It assumes that the original numbers are integers, or can be mapped to integers, such as rational numbers. It requires an array whose size is half the sum of all the numbers, divided by their lowest common denominator. The algorithm enumerates all possible subset sums, as opposed to all possible subsets, to determine the achievable subset sum closest to half the total sum. See [2] for more detail. Although this algorithm runs in pseudo-polynomial time, its space complexity makes it impractical for numbers of high precision.

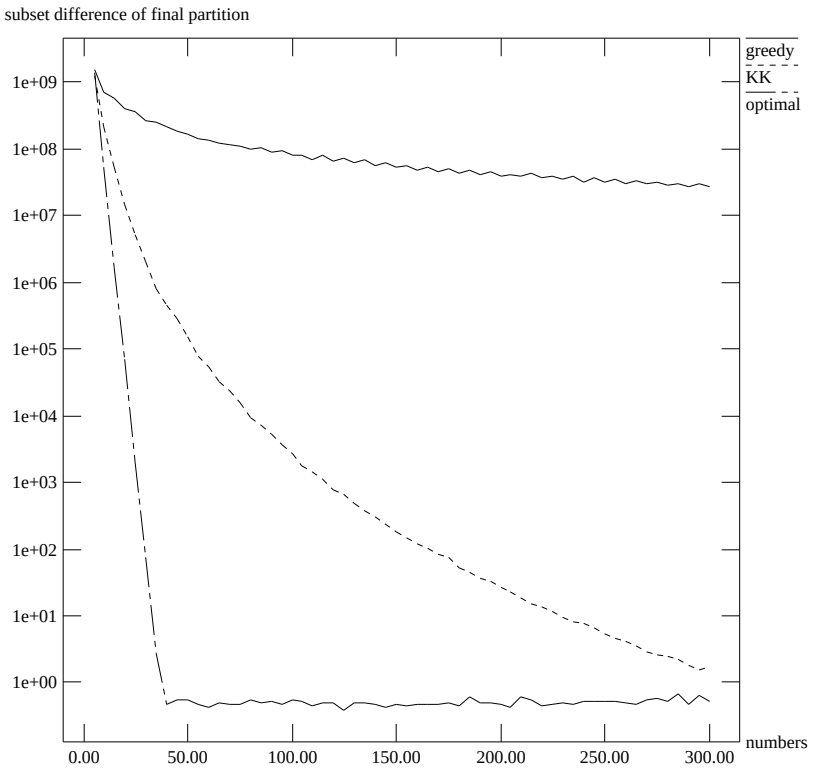

Figure 2: Greedy, KK, and optimal solution quality for 10-digit numbers

2.3 Finding Optimal Solutions

Both algorithms run in $O(n \log n)$ time, but only find approximate solutions. If we want optimal solutions, the obvious algorithm is to search a binary tree, where at each node the left branch assigns the next unassigned number to one subset, and the right branch assigns it to the other subset. We return the best final difference found during the search, along with the actual subsets.

The running time of this algorithm is $O(2^n)$, since we are searching a binary tree of depth n, and its space complexity is $O(n)$, since we can search it depth-first. There are two ways to prune the tree, however. At any node where the difference between the current subset sums is greater than or equal to the sum of all the remaining unassigned numbers, all the remaining numbers are placed in the smaller subset. For example, in the state $15, 0(6, 5, 4)$, the sum of 6, 5, and 4 is no greater than the current subset difference of

since all the numbers must eventually be combined, and $n-1$ edges are created, one for each differencing operation. We then color the nodes of the graph with two colors, so that no two adjacent nodes receive the same color, to get the final partition itself. To two-color a tree, color one node arbitrarily, and then color any node adjacent to a colored node the opposite color. Two-coloring the above graph results in the subsets $(7, 4, 5)$, and $(8, 6)$, whose subset sums are 16 and 14, respectively, for a final partition difference of 2.

The running time of this algorithm is $O(n \log n)$ to sort the n numbers, $O(n \log n)$ for the differencing, since each difference must be inserted into the sorted order, and finally $O(n)$ to two-color the graph, for an overall time complexity of $O(n \log n)$.

The KK heuristic finds much better solutions on average than the greedy heuristic. Figure 2 compares the two algorithms, partitioning random integers uniformly distributed from 0 to 10 billion. The horizontal axis is the number of values partitioned, and the vertical axis is the difference of the final subset sums, on a logarithmic scale. Each data point in the top two lines is an average of 1000 random problem instances, while those in the bottom line are averages of 100 problems. As the number of values increases, the final difference found by the KK heuristic becomes orders of magnitude smaller than for the greedy heuristic. We also show the optimal solution quality. With 40 or more 10-digit numbers, a perfect partition difference of zero or one was found in every case. By about 300 numbers, the KK line nearly joins the optimal line, finding a perfect partition almost every time. The greedy line, however, remains relatively flat.

The explanation for the difference between the quality of the greedy and KK solutions is quite simple. The difference of the final partition is on the order of the size of the last number to be assigned. For the greedy heuristic, this is the size of the smallest original number. This explains the small improvement with increasing problem size, since the more values we start with, the smaller the smallest of them is likely to be. For n numbers uniformly distributed between zero and one, the greedy method produces a final difference of $O(1/n)$. For the KK method, however, repeated differencing operations dramatically reduce the size of the remaining numbers. The more numbers we start with, the more differencing operations, and hence the smaller the size of the last number. Karmarkar and Karp showed that the value of the final difference is $O(1/n^{\alpha \log n})$, for some constant α[6].

from both subsets without affecting the final difference. Similarly, placing 8 in the right subset and 7 in the left subset is equivalent to placing 1 in the right subset. The algorithm removes the two largest numbers, computes their difference, and then treats the difference just like any other number to be assigned, inserting it in sorted order in the remaining list of numbers. The algorithm continues removing the two largest numbers, replacing them by their difference in the sorted list, until there is only one number left, which is the value of the final subset difference.

For example, given the sorted numbers $(8, 7, 6, 5, 4)$, the 8 and 7 are replaced by their difference of 1, which is inserted in the remaining list, resulting in $(6, 5, 4, 1)$. Next, the 6 and 5 are replaced by their difference of 1, yielding $(4, 1, 1)$. The 4 and 1 are replaced by their difference of 3, giving $(3, 1)$, and finally the difference of these last two numbers is the final subset difference of 2. The KK heuristic also fails to find the optimal partition in this case, but does better than the greedy heuristic.

To compute the actual partition, the algorithm builds a tree, with one node for each original number. Each differencing operation adds an edge between two nodes, to signify that the corresponding numbers must go in different subsets. The resulting graph forms a spanning tree of the original nodes, which is then two-colored to determine the actual subsets, with all the numbers of one color going in one subset.

For example, Figure 1 shows the final tree for the example above. First, replacing 8 and 7 by their difference creates an edge between their nodes. The larger of the two, node 8, represents their difference of 1. Next, replacing 6 and 5 by their difference adds an edge between their nodes, with node 6 representing their difference of 1. We then take the difference of 4 and 1, representing the difference between 7 and 8, and add an edge between node 4 and node 8, since node 8 represents the difference of 1. Since 4 is larger than 1, node 4 represents their difference of 3. Finally, an edge is added between node 4 and node 6, representing 3 and 1 respectively.

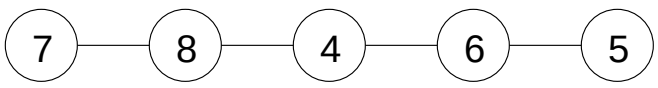

Figure 1: Tree from KK partitioning of (4,5,6,7,8)

In general, the resulting graph forms a spanning tree of the original nodes,

CKK is the best existing algorithm for number partitioning, outperforming even stochastic approaches. It also can be directly applied to the 0/1 knapsack problem, which can be reduced to number partitioning. Instead of incrementally building a single partition, CKK constructs a large number of subpartitions, and combines them together in all possible ways. This general strategy may be applicable to other NP-complete problems as well.

2 Previous Work

2.1 Greedy Heuristic

The obvious greedy heuristic for this problem is to first sort the numbers in decreasing order, and arbitrarily place the largest number in one of the two subsets. Then, place each remaining number in the subset with the smaller total sum thus far, until all the numbers are assigned.

For example, given the sorted numbers $(8, 7, 6, 5, 4)$, the greedy algorithm would proceed through the following states, where the numbers outside the parentheses are the current subset sums: $8, 0(7, 6, 5, 4)$, $8, 7(6, 5, 4)$, $8, 13(5, 4)$, $13, 13(4)$, $13, 17()$, for a final subset difference of 4. Note that the greedy algorithm does not find the optimal solution in this case. The above notation maintains both subset sums, but to find the value of the final difference, we only need to store the difference of the two subset sums. Thus we can rewrite the above trace as: $8(7, 6, 5, 4)$, $1(6, 5, 4)$, $5(5, 4)$, $0(4)$, $4()$. In practice, we would keep track of the actual subsets as well.

This algorithm requires $O(n \log n)$ time to sort the n numbers, and then $O(n)$ time to assign them, for an overall time complexity of $O(n \log n)$.

2.2 Set Differencing (Karmarkar-Karp)

The set differencing method of Karmarkar and Karp[6], also called the KK heuristic, is another polynomial-time approximation algorithm. It also begins by sorting the numbers in decreasing order. At each step, the algorithm commits to placing the two largest numbers in different subsets, while deferring the decision about which subset each will go in. In the above example, if we place 8 in the left subset, and 7 in the right subset, this is equivalent to placing their difference of 1 in the left subset, since we can subtract 7

1 Introduction and Overview

Consider the following very simple scheduling problem. Given two identical machines, a set of jobs, and the time required to process each job on either machine, assign each job to one of the machines, in order to complete all the jobs in the shortest elapsed time. In other words, divide the job processing times into two subsets, so that the sum of the times in each subset are as nearly equal as possible. This is the two-way number partitioning problem, which is NP-complete[2]. The generalization to k-way partitioning with k machines is straightforward, with the cost function being the difference between the largest and smallest subset sums. This basic problem is likely to occur as a subproblem in many practical scheduling applications.

For example, consider the set of numbers $(4, 5, 6, 7, 8)$. If we divide it into the two subsets $(7, 8)$ and $(4, 5, 6)$, the sum of each subset is 15, and the difference between the subset sums is zero. In addition to being optimal, this is also a perfect partition. If the sum of all the numbers is odd, a perfect partition will have a subset difference of one.

We first present previous work on this problem, including a polynomial-time greedy heuristic, and an exponential-time algorithm to find optimal solutions. We also present an elegant polynomial-time approximation algorithm due to Karmarkar and Karp[6], called set differencing or the KK heuristic, which dramatically outperforms the greedy heuristic. Our main contribution is to extend the KK heuristic to a complete algorithm, which we call Complete Karmarkar Karp (CKK). The first solution returned by CKK is the KK solution, and as the algorithm continues to run it finds better solutions, until it eventually finds and verifies an optimal solution.

We present experimental results comparing CKK to the best previous algorithm for finding optimal solutions. CKK appears to be asymptotically faster than the previous algorithm, and provides orders of magnitude improvement when perfect partitions exist. Due to the existence of perfect partitions, it is possible in practice to optimally partition arbitrarily large sets of numbers, if the number of significant digits in each number is limited. This limit is about twelve decimal digits for two-way partitions. This is not a limitation in practice, since no physical quantities are known with higher precision. We have performed the same experiments for three-way partitioning, with similar results, except that the practical precision limit for optimally partitioning large sets is only six decimal digits.

Optimal Number Partitioning

Richard E. Korf
Computer Science Department
University of California, Los Angeles
Los Angeles, Ca. 90024
korf@cs.ucla.edu

August 18, 1995

Abstract

Given a set of numbers, the two-way partitioning problem is to divide them into two subsets, so that the sum of the numbers in each subset are as nearly equal as possible. The problem is NP-complete, and is contained in many scheduling applications. Based on a polynomial-time heuristic due to Karmarkar and Karp, we present a new algorithm, called Complete Karmarkar Karp (CKK), that optimally solves the general number-partitioning problem, and significantly outperforms the best previously-known algorithms for this problem. By restricting the numbers to twelve significant digits, CKK can optimally solve two-way partitioning problems of arbitrary size in practice. Arbitrary-size three-way partitioning problems with up to six digits of precision are solvable. For numbers with greater precision, CKK first returns the Karmarkar-Karp solution, then continues to find better solutions as time allows. Over six orders of magnitude improvement in solution quality is obtained in less than an hour of running time. Rather than building a single solution one element at a time, CKK constructs subsolutions, and combines them in all possible ways. CKK is directly applicable to the 0/1 knapsack problem, since it can be reduced to number partitioning. This general approach may also be applicable to other NP-hard problems as well.