Larger Golomb Rulers

Tomas Rokicki and Gil Dogon
rokicki@gmail.com, gil.dogon@mobileye.com

Abstract

We present the construction of possibly-optimal Golomb rulers through a size of 40,000 marks, which provides support for the existence of subquadratic Golomb rulers for all sizes. Then we explore the subquadratic question at more length, finding a first gap at 492,116 through 492,118 at which sizes subquadratic Golomb rulers may not exist. Finally, we explore some even larger rulers, looking for those with a maximum quality (a measure of how much better than quadratic they are). Our results are strictly computational in nature, but derive from performance improvements to the known algorithms implementing construction methods based on finite fields.

1 Introduction

A *Golomb ruler* is a set of nonnegative integer values (marks) that includes zero and whose pairwise differences are all unique. That is, for every pair of distinct numbers (a, b) in the set, there is no other pair (c, d) such that $a - b = c - d$ (unless $a = c$ and $b = d$). The values in the set are called *marks*, like marks on a ruler. The *length* of the ruler is equal to the maximum value in the set.

Finding the shortest Golomb ruler with a given number of marks (n) may be a difficult task. Indeed, a massive distributed project is under way to prove that the known 28-mark ruler (with a length of 585) is the shortest possible[4]; it has been running for nearly two years and is about 12% complete.

Finding good (near-optimal) Golomb rulers using normal search techniques is difficult; we are not aware of any effective approach at this point in time, despite much interest and many investigations But a little bit of mathematics originated by James Springer in 1938[11] (with some follow-up help from R. C. Bose and S. Chowla[1][2]) gives construction techniques for rulers that appear to be optimal (and are guaranteed to be near-optimal). Indeed, with the exception of six small sizes, no better Golomb rulers have been found than those generated by these construction methods. All the computational effort expended over the past two decades to prove the optimality of rulers with 18 through 27 marks has supported the assertion that rulers generated by these 75-year-old ideas are the best we can do.

In addition to these two construction methods, there is another method by Ruzsa[9] that is also asymptotically optimal; we consider this as well, though for the range of mark counts we tried, it was never superior to the other two methods. A good overview of all three construction methods and much more background are given in [5].

James B. Shearer implemented the construction methods in a set of Fortran programs around 1986, and generated the best set of Golomb rulers through 160 marks possible with these methods; his information is published on IBM's research site[10]. He also implemented a brute-force search program that verified or found the optimal ruler lengths through 16 marks.

We have reimplemented Singer, Bose, and Chowla's ideas and run their constructions on modern machines, extending the set of possibly optimal rulers up through 40,000 marks. This has required some new ideas and improvements to the existing algorithms, which we describe below.

We believe these rulers are probably optimal, and are offering a $250 reward if anyone can beat any of these ruler lengths (for 36 through 40,000 marks) using any technique[8].

2 Modular Rulers

Every proper subset of a Golomb ruler is also a Golomb ruler (but one with fewer marks). (If the zero value is omitted you can subtract the smallest value from each value in the set to obtain a valid ruler.) Also, every Golomb ruler can be flipped—take the maximum value in the set, and subtract from this every value in the set to generate a new set that measures the same distances.

All of the constructions we use generate *modular* rulers. A modular ruler is one such that the marks can be placed on the circumference of a circle (of a specific size, the modulus), and all the distances measured are unique. While a standard Golomb ruler measures $n(n-1)/2$ distinct distances, a modular ruler measures $n(n-1)$ distinct distances. Every subset of a modular ruler is also a modular ruler (with the same modulus). Every modular ruler is also a Golomb ruler.

Thus, every ruler we claim to be optimal actually satisfies more constraints than strictly required; in addition to satisfying the normal Golomb requirements, it also is a modular ruler. While any Golomb ruler is also a modular ruler if you make the modulus large enough, it turns out the modulus for our rulers is typically only a little more than the total length of the ruler itself.

A perfect Golomb ruler with n marks measures exactly the distances 1 through $n(n-1)/2$. This is only possible through 4 marks (can you prove it is not possible for 5 or more marks?) A perfect modular ruler with n marks measures exactly the distances 1 through $n(n-1)$ and it is easy to show that its modulus must be $n(n-1)+1$. The 1938 Singer construction, which generates most of the rulers we suggest are optimal, generates perfect modular rulers, every time, but only for a number of marks that is one more than a prime power. Indeed, it generates many, many different perfect modular rulers. For instance, the Singer construction generates the following perfect modular ruler with 12 marks (and with modulus 133):

$$(0, 1, 12, 14, 22, 29, 54, 60, 63, 90, 110, 129)$$

Pick any number from 1 to 133, and you can find two values that differ by that number (modulo 133) from the above set. For instance, for a difference of 65, you can pick 90 and 22; $(22 - 90) \equiv 65 \bmod 133$. The best length 11 ruler that we can extract from directly this ruler is of length 94 (from the 129 wrapping around to the 90). But we can generate a new modular ruler from this one by multiplying each value by 20, which is relatively prime to 133, and sorting the result, giving us

$$(0, 3, 14, 16, 20, 41, 48, 53, 63, 71, 72, 107)$$

From this modular Golomb ruler we can extract a length-11 Golomb ruler of length only 72 (from the 0 on the left to the 72 towards the end); this is one of only two optimal length-11 Golomb rulers.

The other two constructions we consider, the Bose-Chowla construction and the Ruzsa construction, similarly generate near-optimal modular rulers.

3 The Methods

All the methods take a single input g, which is a prime power (or for Ruzsa's construction, just a prime) and require roughly the following steps:

PrimPoly Find a primitive polynomial for a specific field generated by that prime. This provides a numbering of the elements of the field.

Select From that field, select members that satisfy a specific criteria specific to the construction. This will generate the modular ruler.

Multiply From the original modular ruler, generate a lot of additional modular rulers by multiplying that ruler by numbers relatively prime to the modulus.

Sort Since the values of the multiplied modular ruler are not generated in numerical order, sort the modular ruler.

Scan For each of these modular rulers, try all contiguous sub-rulers as candidate Golomb rulers.

The Fortran programs written by James B. Shearer include all of these steps in a clear and straightforward manner. In order to extend the results as far as we did, the algorithms for each step needed improvement.

There have been a number of papers written on finding primitive polynomials in Galois fields, and implementations of these ideas are available on the web. The ideas in these papers are nontrivial and critical in order to generate primitive polynomials of Galois fields with millions or billions of elements. Both of us independently authored primitive polynomial generation code.

The selection criteria for each of the distinct constructions is fairly complex (especially for the Singer construction). One of us essentially copied Shearer's code into C for this; the other reimplemented everything from scratch. The performance of the Shearer code here was adequate for our purposes.

There were a number of improvements to the multiply step. Both the Singer and the Bose-Chowla construction generated $O(g^2)$ distinct modular rulers (where g is the prime power used for generation). Since we will be sorting and scanning for every multiplicand we consider, it is important to eliminate redundant multiplicands. All of the constructions exhibited symmetry based on a particular subgroup, so we were able to exclude 5/6 of the multiplicands for the Singer construction and 3/4 of the multiplicands for the Bose-Chowla construction. The Ruzsa construction required consideration of many fewer multicands $(O(g))$ so symmetry reduction was not needed here.

In addition, the multiply step itself can be slow, with a multiply and remainder calculated for each value. We improved this by maintaining a cache of small multiples of the input ruler modulo the modulus. When evaluating a new modulus m, if that was small compared to the previous modulus considered, we would just add-and-test the relevant multiple from the cache. The compiler turned this into efficient parallel SSE instructions.

Most of the time in our program was spent in the sort phase. We tried numerous variations of radix sort and finally adopted one that was able to sort more than 100 million elements per second.

In the scan phase, we want to consider every possible sub-ruler from the current modular ruler for every possible Golomb ruler length. A straightforward implementation (as Shearer used) would take time $O(n^2)$ for every modular ruler. We were able to improve this as follows. First, since optimal rulers with 27 marks and fewer are already known, we did not scan for shorter rulers; we started our set of scans at 27. (We only did this optimization for values of g of 1000 or more, so we could pick up the known values as well, all of which are generated with small values of g.) So let us say the shortest ruler found for 27 marks for the current modular ruler turned out to be 950. For each mark count, we decided we did not care about rulers of lengths greater than n^2. Further, from the known lower bound on the lengths of a Golomb ruler $(n^2 - 2n^{1.5} + \sqrt{n} - 2)$. When considering whether to scan at length $27 + x$, we would add our best at 27 (950) to the lower bound on the length of a ruler with x marks, and if that was greater than either $(27 + x)^2$ or our best found ruler at length $27 + x$, we would skip that scan since it could never improve or match our best. This simple idea let us skip almost all scans with little effort. For instance, for 14,533 marks, we considered 22,680,000 moduli. Without this optimization we would have had to do roughly 14,000 scans per moduli; with it, we ended up doing only an average of 14.7 scans per moduli, a reduction by a factor of nearly 1000.

These techniques permitted us to fully explore all three finite-field construction methods through a mark count of 40,000. In no case did the Ruzsa construction method yield the best result; approximately 2/3 of the time the projective plane method gave the best ruler, and the affine plane method gave the best results in the remaining cases. The final length of the best rulers found were in general very close and slightly less than the number of marks squared. Rulers with n marks and of length less than n^2 are called subquadratic rulers. For each mark count we were able to find a subquadratic ruler.

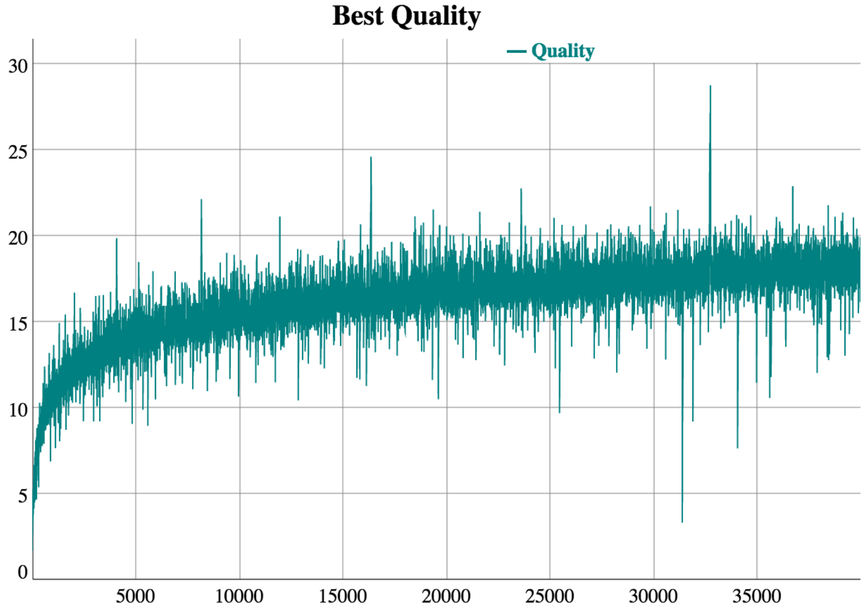

Figure 1: The best quality obtained for mark counts of 7 through 40,000 using the finite field construction methods for Golomb rulers.

The *quality* $b(R)$ of a Golomb ruler R is the difference between the number of marks in the ruler and the square root of its length. A positive quality indicates a subquadratic ruler. For a given number of marks, the quality of the best ruler found by our computations varied from about 1 to about 29, as shown in Figure 1. For a given small range of mark counts, the best quality generally lies within a small range, with occasional spikes up and down. In addition, the quality trends up with increasing mark count and clearly in a sublinear fashion.

The full data, including all modular rulers and the necessary multipliers and moduli to generate the best rulers, is available at [8].

4 Golomb Rulers: Pushing the Limits

The previous section extended the calculation of near-optimal rulers based on affine and projective planes through 40,000 marks. These constructions generate the best known rulers for all sizes in excess of 16 marks. To our knowledge, despite significant interest, no one has been able to find any Golomb ruler with more than 16 marks with length less than the best of that generated by the affine or projective plane construction. This leads to hypothesis A:

Hypothesis A: For all $n > 16$, the optimal Golomb ruler can be generated by an affine plane or projective plane construction.

No counterexample to hypothesis A is known despite a great deal of CPU expended over the past several decades, not the least of which was the distributed.net's OGR-24 through OGR-28 efforts which have consumed thousands of CPU years over the past 16 years. Also, the search for Golomb rulers has been a central problem for constraint-based search for some time, with no counterexamples found.

Our results, and previous work by other authors [3] have found that subquadratic rulers exist at least through 65,000 marks. This leads to hypothesis B:

Hypothesis B: subquadratic rulers exist for all n.

If OGR(n) is the optimal Golomb ruler of length n, then assuming hypothesis B is true, $b(\mathrm{OGR}(n))$ is positive for all n. Another question of interest is, how large can the quality $b(\mathrm{OGR}(n))$ be?

Hypothesis C: $b(\mathrm{OGR}(n))$ is unbounded.

This hypothesis was conjectured by Erdös and Turán in 1941[6].

5 Finding Subquadratic Rulers

Using the programs created for our previous work, we set out to explore these hypotheses for larger values of n. First, we tried to extend the work of Dimitromanolakis past 65,000. There are three known asymptotically-optimal algebraic constructions: the Ruzsa construction, the affine plane construction, and the projective plane construction. The first two construction methods have very fast (linear time) modular ruler construction methods. The Ruzsa is trivially linear-time; the affine construction is linear from ideas in [7], but for the projective plane ruler construction only a quadratic-time algorithm is known. The affine plane construction usually generates better rulers than the Ruzsa construction, so we use that as our main workhorse. We did not consider all possible derived Golomb rulers but only scanned a modular Golomb ruler until we showed the existence of a subquadratic ruler for all sizes in the range between the generating prime power and the previous prime power. This very quickly covered the entire range from 40,000 through to 500,000 except for eight gaps between prime powers.

For the larger gaps that caused more problems, we started with the projective plane construction rather than the affine plane construction; even though generating the modular ruler takes longer for the projective plane construction, the actual scanning of the derived rulers, which dominates the run time ($O(n^3)$), can be done more rapidly.

6 Multiplicative Groups for Faster Scanning

When generating derived modular rulers from a given modular ruler, you multiply the original ruler by a value m that is relatively prime to the modulus of the rulers (and the multiplication is done modulo the modulus) and then sort the resulting values. For instance, for the projective plane construction, with a prime of 13, the modulus is $13^2 + 13 + 1 = 183$, which is factored into 3*61. You would consider rulers generated by multiplying the original rulers by the values 2, 4, 5, 7, 8, 10, 11, and so on. The multiplication, remainder, and the sort operations are the most costly operations in the scan when you are focused on a small range of final Golomb rule sizes; we were focused on only the small range for which we had not already found subquadratic rulers.

Rather than considering the distinct multipliers sequentially, you can instead consider them in a sequence that emphasizes multiplication by small integers, especially 2 or 3. Given an already sorted set, you can perform multiplication, modulo, and sorting by such a small integer much more rapidly than the naive implementation. For the multiplication and modulo steps, you can instead use addition and subtraction, where the value to be subtracted changes only a small number of times throughout the range. The sort step can be done by merging a small number of already-sorted subranges. For instance, given the projective-plane-derived modular ruler

$$(0, 1, 8, 24, 37, 41, 59, 107, 119, 128, 134, 139, 153, 181)$$

with the modulus 183, multiplication by two gives

$$(0, 2, 16, 48, 74, 82, 118, 214, 238, 256, 268, 278, 306, 362).$$

The value to subtract to keep it in range is 0 for the first half (through 118) and then 183 for the rest, giving us

$$(0, 2, 16, 48, 74, 82, 118, 31, 55, 73, 85, 95, 123, 179).$$

The sort step is just a simple merge of the two sorted subsequences, giving

$$(0, 2, 16, 31, 48, 55, 73, 74, 82, 85, 95, 118, 123, 179).$$

The multipliers used must be relatively prime to the modulus. For projective-plane-derived rulers, the modulus for a prime power g is $g^2 + g + 1$, which is never divisible by 2, so most of the time the generation of the next derived ruler can use the multiplier 2. For affine-plane-derived rulers, the modulus for a prime power g is $g^2 - 1$, which is usually divisible by 2 (except for when g is a power of 2) and usually divisible by 3 (since $g^2 - 1$ is $(g-1)(g+1)$ and since g is usually not divisible by 3, one of $g-1$ or $g+1$ usually is) so the smallest multiplier we can frequently use is 5. The time required for the merging of the sorted sublists increases with the logarithm of the multiplier; this may not seem like much, but for a multiplier of 5 we will spend more than twice as much time merging as we would for a multiplier of 2.

For the example above, with $g = 13$ and a modulus of 183, we only ever need use a multiplier of 2, as the first 20 powers of 2 mod 183 generate all necessary multipliers:

$$(1, 2, 4, 8, 16, 32, 64, 128, 73, 146, 109, 35, 70, 140, 97, 11, 22, 44, 88, 176)$$

(The other 100 multipliers that are relatively prime to 183 generate equivalent rulers by symmetry; we omit the details.)

Therefore, even though generating the original modular Golomb ruler is slower with the projective plane construction, if we plan to scan a significant fraction of the possible generated rulers, the projective plane construction ends up being faster. Another benefit was that the projective plane construction typically (about 2/3rds of the time) generated slightly shorter rulers.

7 The Challenge: Large Prime Gaps

For the gaps that turned out to be difficult, we used the projective plane construction. With this, we were able to cover all but one gap in the range through 500,000. The remaining gap was 492,116 through 492,118. In order to fully explore the projective plane construction, we had to construct all 39 billion possible different modular rulers and scan each of them for potential solutions. Next, we did the same for the affine plane construction, constructing and scanning all 10 billion possible different modular rulers. None of them generated subquadratic rulers for the three target mark counts given above.

This result is not completely surprising; the gap between the consecutive prime powers 492,113 and 492,227 is 114, which is a maximal gap (the first gap of that size or greater in the sequence of prime powers). The typical prime gap for primes around this size is ln(p) or about 14. (The projective plane construction for $g = 492,113$ found subquadratic Golomb rulers for 492,113 and 492,114 marks.) Empirically, we see that the projective plane and affine plane constructions typically create subquadratic rulers only for values close to the prime power used to generate them (and for some trivially small values); for this maximal prime gap, the size of the prime gap was too large to span.

We did not evaluate the next prime power up (492,251) for this range because of the computational time required and because we believe it is extremely unlikely to generate subquadratic rulers that small, based on the relationship between the size of the range of subquadratic rulers generated and the prime power used to generate them.

So we have found that for all sizes through 492,115 marks a subquadratic Golomb ruler exists; for all sizes through 500,000 marks, only three values

(492,116 through 492,118) do not have a known subquadratic ruler. In addition, one of the following must be true:

Hypothesis A is false; there is some subquadratic ruler for 492,116 that is not generated by the projective plane or affine plane construction;

Hypothesis B is false; there is no subquadratic ruler for 492,116;

A subquadratic ruler for 492,116 can be constructed by the projective plane or affine plane construction for a prime power of 492,251 or greater.

Some bug exists in our programs. To avoid this we each independently implemented the ideas in this paper and compared the results thoroughly.

We believe one of hypothesis A or B is false.

There has been much attention in the literature to using general optimization and search techniques to find Golomb rulers. So far these techniques have been shown to be effective only for a very small number of marks (through 16 to 20 marks), far short of the 492,116 marks we are interested in. Alternative algebraic or algorithmic construction methods are significantly inferior to the projective plane or affine plane construction (with the exception of one method, the Ruzsa method, which is only somewhat inferior and which we have also explored for the range in question without finding a subquadratic ruler).

Despite great interest, no construction techniques of comparable effectiveness have been found (the projective plane and affine plane construction techniques date back to 1938 and 1941, respectively).

8 Rulers of the Highest Quality

We now turn our attention to hypothesis C. Empirically we have found that powers of two greater than 2^{10} generate rulers with the largest $b(R)$. We do not yet have an explanation on why powers of two generate better Golomb rulers. The symmetry of the set of derived rulers for $g = p^n$ has order $4n$ (for the affine plane construction) and $6n$ (for the projective plane construction), so the higher n for powers of two at roughly equivalent sizes would yield fewer derived rulers to explore, and thus we would expect slightly worse rulers to be found. Yet, consistently, powers of two yield Golomb rulers of higher quality than other prime powers; every record quality (by increasing n) between 4,000 and 40,000 was found by the affine or projective plane constructions starting from a modular ruler derived from a power of two.

We assume this continues for values of n greater than 40,000. We focused on these rulers to try to extend the known best $b(R)$. We explored modular rulers of size $2^{16} = 65,536$ through $2^{21} = 2,097,152$. Since our purpose was to opportunistically look for rulers of high quality, rather than to prove an optimal quality, we only considered derived rulers (sub-rulers) whose length is close to the power of two (within 110). Further, for rulers larger than 32,768, we only considered the projective plane and affine plane constructions (though we have not yet run the affine plane construction for sizes 2^{20} or 2^{21}). The number of rulers searched grows exponentially with n by approximately a factor of four. For 2^n, we have searched the full range of inequivalent possibilities for a modular ruler, which is $\phi(2^{2n} + 2^n + 1)/6n$ in the projective case, and $\phi(2^{2n} - 1)/4n$ in the affine case (ϕ is Euler's totient function).

For the largest size, we were able to find a Golomb ruler at $n = 2,097,125$ marks of length 4,397,762,317,463, for a $b(R)$ of 40.758. Based on this result, and how much larger this is than the $b(R)$ values for rulers smaller than 40,000, we suspect that $b(\mathrm{OGR}(n))$ is indeed unbounded. Table 1 summarizes our results.

g	n	d	$b(R)$	type	# rulers
2,048	2,046	4,124,805	15.04	aff	60,016
2,048	2,046	4,118,063	16.70	proj	54,498
4,096	4,084	16,517,156	19.87	aff	138,240
4,096	4,091	16,583,409	18.73	proj	139,968
8,192	8,182	66,601,079	21.05	aff	859,950
8,192	8,177	66,501,869	22.13	proj	728,208
16,384	16,371	267,316,415	21.19	aff	2,370,816
16,384	16,367	267,074,038	24.60	proj	1,820,448
32,768	32,750	1,070,915,788	25.15	aff	8,910,000
32,768	32,756	1,071,073,269	28.74	proj	11,748,240
65,536	65,515	4,288,124,720	31.23	aff	33,554,432
65,536	65,523	4,289,307,076	30.20	proj	23,224,320
131,072	131,042	17,162,400,535	36.65	aff	168,424,950
131,072	131,033	17,163,471,622	33.57	proj	142,888,536
262,144	262,125	68,691,708,709	33.97	aff	362,797,056
262,144	262,122	68,685,856,466	32.13	proj	424,189,440
524,288	524,252	274,802,524,432	35.90	aff	2,411,191,314
524,288	524,242	274,803,720,226	34.75	proj	2,066,689,584
1,048,576				aff	5,921,280,000
1,048,576	1,048,529	1,099,332,551,684	38.39	proj	4,704,480,000
2,097,152				aff	28,901,432,448
2,097,152	2,097,125	4,397,762,317,463	40.76	proj	34,426,570,752

Table 1: Record qualities by powers of two. The two missing rows are values we did not fully explore.

9 Discussion

We have extended the known results on Golomb rulers in three significant ways. First, we have extended the computation of the likely-optimal rulers from 160[10] to 40,000. Second, we have extended the search for subquadratic rulers from 65,000[3] to 500,000, except for a gap of three values 492,116 through 492,118 for which we've proved the known constructions do not suffice. Third, we have found Golomb rulers with the best known qualities, including one with a quality of 40.76.

Decades of search support the assumption that the known finite-field methods (which are over 70 years old at this point) generate the best rulers larger than 16. In hopes of spurring additional investigation, we are offering a prize of $250 to anyone who can find any Golomb ruler in the range 30 through 40,000 that is shorter than the ones we have computed.

At the same time, the existence of subquadratic rulers through nearly half a million support the assumption that subquadratic rulers always exist. Our computational results indicate that one of these two suppositions is likely false, because the finite field construction methods only work well near prime powers, but the prime powers exhibit increasing gaps as they get larger. If there is a subquadratic ruler of size 492,116, it is well beyond our current techniques to find it. At the same time, we have little hope of proving its nonexistence.

10 Acknowledgements

We thank Martin Gardner for introducing us to Golomb rulers decades ago, and Al Zimmermann for reawakening our interest in them through a recent programming contest[12].

References

[1] R. C. Bose, An affine analogue of Singers theorem, Journal of the Indian Mathematical Society, 6 (1942), 115.

[2] R. C. Bose and S. Chowla, Theorems in the additive theory of numbers, Commentarii Mathematici Helvetici, 37 (1962-63), 141147.

[3] Apostolos Dimitromanolakis, Analysis of the Golomb Ruler and the Sidon Set Problems, and Determination of Large, Near-Optimal Golomb Rulers

[4] distributed.net, OGR-28 Project Overview
`http://stats.distributed.net/projects.php?project_id=28`

[5] Konstantinos Drakakis, A Review of the Available Construction Methods for Golomb Rulers, Advances in Mathematics of Communication, vol. 3 no. 3, 2009.

[6] P. Erds and P. Turn, On a Problem of Sidon in Additive Number Theory and on Some Related Problems, Journal of the London Mathematical Society, 16 (1941), 212215.

[7] B. Lindstrm, Finding finite B2-sequences faster, Mathematics of Computation, 67 (1998), 1173-1178.

[8] Tomas Rokicki and Gil Dogon, Larger Golomb Rulers, `http://cube20.org/golomb/`.

[9] I. Z. Ruzsa, Solving a linear equation in a set of integers I, Acta Arithmetica, LXV (1993), 259-282.

[10] James Shearer, Golomb Rulers, `http://www.research.ibm.com/people/s/shearer/grule.html`

[11] James Singer, A Theorem in Finite Projective Geometry and Some Applications to Number Theory, Transactions of the American Mathematical Society, **43** (1938), 377385.

[12] Al Zimmermann, Al Zimmermann's Programming Contests, `http://www.azspcs.net/`.