INTEGERS: ELECTRONIC JOURNAL OF COMBINATORIAL NUMBER THEORY 5 (2005), #G03

VARIATIONS ON A THEME OF EUCLID

David Collins
Department of Mathematics, Occidental College, Los Angeles, CA 90041, United States
dcollins@oxy.edu

Received: 12/29/03, Revised: 11/5/04, Accepted: 4/12/05, Published: 5/10/05

Abstract

The game of Euclid is an impartial game played between two players. A position in the game is a pair of integers (a, b). A move consists of replacing the current position with one in which the larger of a and b has been reduced by any multiple of the smaller. The game ends when the two numbers are equal. The players alternate moves, and the winner is the last player to make a move.

Several variations take the form of restrictions on the moves available to the players. One important class of restrictions takes the form of a set Λ of positive integers from which the number of multiples a player removes on a turn must be chosen.

Of particular interest are versions with "dynamic" restrictions. In these variations of Euclid, the maximum multiple which can be removed on a turn is governed by some given function. In this way, the set of available moves changes as the game proceeds.

It is shown how all versions considered can be recast as sequential take-away games, and this transformation is frequently used to find winning strategies.

1. Introduction

The game of Euclid is an impartial game played between two players. A position in the game is a pair of integers (a, b). A move consists of replacing the current position with one in which the larger of a and b has been reduced by any multiple of the smaller, with the proviso that the result must remain positive. The game ends when the two numbers are equal. Without loss of generality, we can assume that a and b are relatively prime and that $a \leq b$ for any position (a, b). The players alternate moves, and the winner is the last player to make a move, i.e., to move to $(1, 1)$.

Euclid is a Nim-like game, and we can use the Sprague-Grundy theory of impartial games. We also employ the convention of referring to positions in which the next player to move has a win as N-positions and to those in which the previous player to move wins as P-positions (cf. [1], Ch. 4). An N-position has a positive Sprague-Grundy value, while a P-position has a Sprague-Grundy value of zero.

We begin our discussion of Euclid with a result which makes finding the winner easy.

Theorem 1 *The first player to have more than one available move has a winning strategy in Euclid.*

Proof. If the player to move has more than one choice, we have $b > 2a$. Choose n such that $a < (b - na) < 2a$. We need to consider only two of the possible moves from (a, b): to $(a, b - na)$ and to $(a, b - (n + 1)a)$. Call these moves A and B, respectively. Then, if B is a P-position, the first player can simply move there. Otherwise, moving to A will force the second player to move to B. Clearly, the first player will win in one of these cases. □

There are a few interesting things about this theorem. First, it is non-constructive: it does not actually tell us how to win. Later on, we shall see several different ways of constructing the winning strategy. The advantage of this form is that the principle of the first player with a choice winning applies to many variants of Euclid while the winning strategy itself varies. Second, it applies equally well to the *misère* form of the game where the last player to move loses. In many Nim-like games, solving the misère version is—relative to solving the standard version—misery (a miserable pun). Fortunately, this is not the case with Euclid, at least as far as knowing who has a winning strategy (as we shall see in Section 7).

The first and probably the simplest way to approach the task of finding explicit winning strategies in the game of Euclid is to use the golden ratio, ϕ. We claim that in Euclid, the first player has a win in exactly those positions in which $b/a > \phi$. The validity of this claim depends on the following two assertions:

1. From every position for which $b/a > \phi$, there is a move which leaves a position (a', b') with $1 \leq b'/a' < \phi$.
2. Any move from a position with $b/a < \phi$ leaves a position (a', b') with $b'/a' > \phi$.

These conditions are equivalent to the assertion that the positions in Euclid with a Sprague-Grundy value of 0 are precisely those with $1 \leq b/a < \phi$. The interested reader can consult [6] for a proof that these conditions do indeed hold.

I shall refer the game considered above as the standard version of Euclid. This paper shall consider many variations of these standard rules. We shall consider first variations in which the number of multiples of a which can be removed from b in a position (a, b)

must be chosen from a fixed set Λ. We shall present a solution for the simplest of these restriction sets, $\Lambda_k = \{1, 2, \dots, k\}$, and extend this theory to many other restriction sets in Section 4. Theorem 2 gives a necessary and sufficient condition for which a larger restriction set is equivalent to some Λ_k. We next consider "dynamic" restrictions which are not fixed but rather change as the game progresses and solve Euclid for one class of such restrictions, the move size restriction, in Theorem 3 (Section 5). The Sprague-Grundy theory does not seem to be useful in analyzing these dynamic restrictions because the winner in a given position is determined not only by the position itself but also by the preceding move.

After a brief discussion of standard Euclid with three numbers rather than two (Section 6), we conclude with a discussion of *misère* forms. In Section 7, Theorem 4 provides the winning strategy for the misère forms of both standard Euclid and of versions with restriction sets equivalent to some Λ_k.

For a game $(a, b), a < b$, we will use the continued fraction representation (Section 2) of b/a throughout and will often view the game as played in the Stern-Brocot tree (Section 3).

2. Continued Fractions

A second method of describing the winning strategy (besides the ϕ-based approach seen above) is through continued fractions. This method has the advantage of being more amenable to the analysis of the various variations of Euclid which are the subject of this paper. Each fraction b/a can be represented as a continued fraction in two ways: as $[a_0, a_1, \dots, a_n]$ (called the short form) and as $[a_0, a_1, \dots, a_n - 1, 1]$. Later on, we shall see why it does not matter which form of the continued fraction expansion we use.

It is possible to reinterpret Theorem 1 in terms of continued fractions. Euclid was first analyzed in this way by Lengyel in [6]. In the position (a, b), $b/a = [a_0, a_1, \dots, a_n]$, a move affects only the leading continued fraction digit a_0. If $a_0 = 1$, the player to move will have only one option. We therefore assume that the game is played until some player has a real choice, and $a_i \geq 2$ with some i. (The existence of such i is guaranteed by writing b/a in the short continued fraction form.) Since this continued fraction is certainly greater than ϕ, that player has a winning strategy.

3. Sequential Take-away Games

The above approach works quite well for the standard version and some of its variations [6]. Nevertheless, we have found it more convenient in investigating a wide variety of generalizations to look at the game of Euclid as what we call a *sequential* take-away game.

A simple take-away game is an impartial game played between two players. A position consists of a single pile of counters, and the players take turns removing some positive number of counters from this pile, subject to the rules of the particular game. The position $(1, n)$ in Euclid corresponds to an utterly trivial take-away game with a pile size of $n - 1$ and no movement restriction whatsoever.

Now that we possess the concept of a take-away game, we can define a sequential take-away game as a game which is divided into a series of smaller "subgames" to be played in a given order, each of which is a take-away game having the same movement restrictions. The players alternate throughout the game; the player to move last in a given subgame moves second in the game following it (cf. [6]). Such a game is denoted by $[a_0, a_1, \ldots, a_n]$, where a_0 is the pile size in the first subgame, a_1 the pile size in the second, and so on.

The game of Euclid is a sequential take-away game, somewhat disguised. The continued fraction representation of the position (a, b) reveals the disguise. Let a/b have continued fraction expansion $[a_0, a_1, \ldots, a_n]$ (cf. [3]). The players' moves reduce the size of the first coefficient a_0 until it is zero (a_0 reflects the number of multiples of b which must be removed from a before the result is less than b). Then, the play continues from a position with continued fraction representation $[a_1, \ldots, a_n]$, and with the players reducing a_1. Thus, Euclid is a sequence of take-away games with successive pile sizes of a_0, a_1, a_2, and so on. The only exception is that one must be subtracted from the last partial quotient of the continued fraction expansion of a/b, because [1] (corresponding to the Euclid position $(1, 1)$) is a terminal position and does not reflect an option.

From this point on, we shall assume that one has already been removed from the last term whenever we consider a Euclid position (or any position) as $[a_0, a_1, \ldots, a_n]$. (Therefore, the "Euclid position representation" $[a_0, a_1, \ldots, a_n]$ for the game (a, b) differs slightly from the original continued fraction expansion of b/a. In fact, this is why it does not matter which of the two continued fraction forms was used in the first place.)

Looking at the game of Euclid as a sequential take-away game makes finding the Sprague-Grundy number of any position much easier; whereas without continued fractions they are rather obscured. To find the Sprague-Grundy values, we work from right to left: to find the Sprague-Grundy value of $[a_0, a_1, a_2, \ldots, a_n]$, we look first at the Sprague-Grundy value of $[a_n]$, then $[a_{n-1}, a_n]$, and so on. To find $g([a_i, \ldots, a_n])$, we need to know only $a_i, i < n$, and the Sprague-Grundy value of $[a_{i+1}, a_{i+2}, \ldots, a_n]$. In [7], Lengyel refines the use of this method to calculate the Sprague-Grundy values for Euclid, and shows how it is related to the Stern-Brocot tree [3]. Later on in this paper, we shall employ a similar recursive technique when looking at variations of Euclid.

4. Static Restrictions

Many different ways of extending the game of Euclid are covered in this paper. Probably the most natural way is to restrict the choice of multiples which can be removed to a given set Λ. This type of restriction remains constant throughout the game and is referred to as a static restriction. In the standard form of Euclid, the number of multiples of a which are removed from b is chosen from the infinite set $\Lambda = \{1, 2, \dots\}$. We can change this set Λ and produce any number of variations on the game. Lengyel [6] solves Euclid for the restriction sets $\Lambda_k = \{1, 2, \dots, k\}$. These games are the sequential extensions of Bachet's subtraction game, a single-pile take-away game where only $1, 2, \dots, k$ counters can be removed on a turn. (In general, a subtraction game consists of the players removing some number of counters from a single pile where the number removed must be chosen from a restriction set [1], Ch. 4.) The Sprague-Grundy function of Bachet's game has the period $(0, 1, \dots, k)$ of length $k + 1$.

Many subtraction-based games with finite or infinite restriction sets display a similar periodicity and are equivalent to Bachet's game for some k. For example, the single-pile take-away game with restriction set $\{1, 2, 3, 5, \dots, p_k, \dots\}$, with p_k being the kth prime number, has the same Sprague-Grundy function as Bachet's subtraction game with $\Lambda_3 = \{1, 2, 3\}$. Now, we show that games which are equivalent to one of Bachet's games in the one-pile version are equivalent in Euclid as well. More formally, let G be any static restriction game and let B_k be Bachet's subtraction game with the set $\{1, 2, \dots, k\}$ of allowed subtractions. We denote the Sprague-Grundy functions of G and B_k by g_G and g_{Bk}, respectively. We write $G \equiv_c B_k$ if games G and B_k have the same Sprague-Grundy function with c ($c \leq k$) being the Sprague-Grundy value of the terminal position, i.e., $g_G = g_{B_k}$ and $g_G(0) = g_{b_k}(0) = c$.

If G and B_k have the same Sprague-Grundy function for all $c \leq k$, we write $G \equiv B_k$. In the one-pile form, of course, the value of the terminal position is zero; but as we have seen above, a Euclid position $[a_0, a_1, \dots, a_n]$ is equivalent to the one-pile game with a_0 counters, the only difference being that the Sprague-Grundy value of the terminal position is $g([a_1, a_2, \dots, a_n])$ rather than zero. The following theorem was suggested by Lengyel.

Theorem 2 *$G \equiv_0 B_k$ if and only if $G \equiv B_k$.*

Proof. The "if" part of the statement is obvious. For the other part we need

Lemma. If $G \equiv_0 B_k$ then $\Lambda_G \supseteq \Lambda_k$ for the subtraction set Λ_G of game G.

Proof of the lemma. We proceed by contradiction. Suppose that $\Lambda_G \not\supseteq \Lambda_k$. Let m be least element of the difference set $\Lambda_k \setminus \Lambda_G$. We have $1 \leq m \leq k$. Then clearly, $g_G(m) = 0$

because we can only move to $1, 2, \ldots, m-1$ in G, and they have non-zero values, for $g_G(i) = g_{B_k}(i) = i, 0 \le i \le k$, by $G \equiv_0 B_k$. On the other hand, $g_{B_k}(m) \ne 0$ because we can move to zero in B_k. Thus, $G \not\equiv_0 B_k$, a contradiction. □

Assume that $\Lambda_G \setminus \Lambda_k \ne \emptyset$. Now, let s be any element of $\Lambda_G \setminus \Lambda_k$, thus $s > k$ by the lemma. There are two cases.

1. For all $s \in \Lambda_G \setminus \Lambda_k$ there exists an $s' \in \Lambda_k$ such that $s - s' \equiv 0 \bmod (k+1)$. We now prove by contradiction that $G \equiv B_k$. (In this case, the option s really adds nothing to the options to remove $1, \ldots, k$, and does not affect the Sprague-Grundy function.) By the lemma, $\Lambda_G \supseteq \Lambda_k$, so that $g_G(i) = g_{B_k}(i)$ for all $i \le k$. Now let t be the least element for which $g_G(t) \ne g_{B_k}(t)$. Then $t > k$ and

$$g_G(t) = mex\left\{\left\{\bigcup_{i=1}^{k} g_G(t-i)\right\} \bigcup \left\{\bigcup_{\substack{s\in\Lambda_G\setminus\Lambda_k \\ s\le t}} g_G(t-s)\right\}\right\} =$$

$$= mex\left\{\bigcup_{i=1}^{k} g_{B_k}(t-i)\right\} = g_{B_k}(t).$$

   The second step in this equality is justified since $g_G(t-s) = g_{B_k}(t-s) = g_{B_k}(t-s')$ where $s' \in \Lambda_k$ and $s \equiv s' \bmod (k+1)$ by the hypothesis. We have a contradiction.

2. There exists an $s \in \Lambda_G \setminus \Lambda_k$ such that for all $s' \in \Lambda_k$, we have $s - s' \not\equiv 0 \bmod (k+1)$, i.e., $g_{B_k}(s - s') \ne 0$. In the games G and B_k, both with $g(0) = 0$, we consider the value of $g(s)$. We know that $g_{B_k}(s) = 0$, otherwise there would exist an option to move to $s - s'$ in B_k with $s - s' \equiv 0 \bmod (k+1)$, violating our assumption. On the other hand, $g_G(s) \ne 0$, because $g_G(s-s) = g_G(0) = 0$. This contradiction shows that $G \not\equiv_0 B_k$, so this case never occurs under the hypothesis of the theorem. □

Remark. If indeed G is equivalent to any B_k, the value of $k+1$ will be the least integer not in Λ_G. In fact, it is not too difficult to see that $G \equiv B_k$ if and only if Λ_G contains $\{1, \ldots, k\}$ but no multiples of $k+1$.

This theorem and the remark allow us to apply the theory of Euclid with restriction set $\Lambda_k = \{1, 2, \ldots, k\}$ to Euclid with other restriction sets. We summarize this approach here. In the position $[a_0, a_1, \ldots, a_n]$, each of the partial quotients can be reduced $\bmod(k+1)$. Then, the game plays exactly as in the unrestricted Euclid, except that it may take somewhat longer. In particular, the first player facing a position with $a_i \ge 2$ has a winning strategy. This extends our first theorem to the restriction sets Λ_k and restrictions equivalent to these by the above theorem. We now turn to some examples of infinite sets which are equivalent to Λ_k.

Example 1: $\Lambda = \{1$ and any number of odd numbers$\}$.

Since one is present but all multiples of two are excluded, this game is equivalent to B_1, a completely deterministic game. In this variation of Euclid, there is never a "real" choice (for the parity of the number of remaining moves changes by every move [6]), and the result is completely unaffected by skill.

Example 2: $\Lambda = \{1, 2, \ldots, 2^n, \ldots\}$.

Since the multiples of three are not powers of two, it is now clear that the version of Euclid with $\Lambda = \{1, 2, \ldots, 2^n, \ldots\}$ is completely equivalent to B_2. At any point, we need only use the options $\Lambda_2 = \{1, 2\}$ to reduce any non-zero position to zero. The additional options serve only to shorten the game as in example 1.

Example 3: $\Lambda = \{1, 2, 3, 5, \ldots, p_k, \ldots\}$.

The set of primes works in an analogous way. This variation of Euclid reduces to B_3. Again, since multiples of four are not prime, the multiples of four are the positions with Sprague-Grundy value of zero. In any position, we need only to use the removal options $\Lambda_3 = \{1, 2, 3\}$.

Example 4: $\Lambda = \{p^k$ for all primes $p, k = 0, 1, \ldots\}$.

We leave it to the reader to show that the removal set of prime powers is equivalent to Λ_5.

5. Dynamic Euclid

We next consider "dynamic" versions of Euclid in which the maximum multiple which can be removed on a given turn is governed by the game function, f. In this way, the set of available moves dynamically changes as the game proceeds. Typically, the reduction technique of Section 4 cannot be applied when dynamic restrictions are imposed on the available moves. In one class of dynamic games, the maximum number of counters $f(n)$ which can be removed is a function of the pile size n. In [4], Holshouser, Reiter, and Rudzinski show how to calculate the Sprague-Grundy values $g(n)$ given a game function $f(n)$. We note that Theorem 1 of [4] guarantees that many of these games cannot be equivalent to Λ_k for any k.

In its Euclid form the pile size restriction makes the maximum number of multiples of

a which can be removed from b in (a, b) a function of the maximum number which could be removed were there no restriction. The Sprague-Grundy numbers of each position in a pile size restriction can be found recursively using the representation of Euclid as a sequential take-away game introduced earlier and working from right to left. Suppose that a Euclid position (a, b) has the representation $[a_0, a_1, \ldots, a_{n-1}, a_n]$ Then, the last segment $[a_n]$ is a single-pile take-away game with pile size a_n, so $g([a_n])$ can be found using the theory of [4]. Then, the Sprague-Grundy function $g([a_{n-1}, a_n])$ is clearly just $g([a_{n-1}])$ permuted so that the output 0 is replaced by $g([a_n])$ and the values less than or equal to $g([a_n])$ are shifted down by one, or, more formally,

$$g([a_{n-1}, a_n]) = \begin{cases} g([a_{n-1}]) - 1, & \text{if } g([a_{n-1}]) \leq g([a_n]) \text{ and } g([a_{n-1}]) \neq 0, \\ g([a_{n-1}]), & \text{if } g([a_{n-1}]) > g([a_n]), \\ g([a_n]), & \text{if } g([a_{n-1}]) = 0. \end{cases}$$

Similarly, we can now find $g([a_{n-2}, a_{n-1}, a_n])$ as a function of $g([a_{n-2}])$ and $g([a_{n-1}, a_n])$, and continue recursively in this fashion until the Sprague-Grundy value of the entire game is known. Thus, we see how results of [4] extend directly to Euclid.

Another class of dynamic take-away games makes the maximum number of counters which can be removed on a given turn a function of the previous move. The theory of one-pile games with this *move size* restriction has been solved by Schwenk in [8]. Here, we offer a complete extension of his theory to the game of Euclid. It will be noted that in Euclid, the restriction is on the maximum number of multiples of a which can be removed from b in a position (a, b), $b \geq a$; and this restriction is a function of the number of multiples removed on the previous turn. Games with a move size restriction are especially difficult to handle because the characterization of the Sprague-Grundy numbers is problematic and does not immediately give the winner and winning strategy. This is why Schwenk chooses a different approach.

Space prevents us from giving all the details of Schwenk's theory here; the interested reader should consult [8]. What follows is a brief summary. In a take-away game, let the maximum number of counters to be removed on the kth $(k \geq 2)$ move be $f(T(k-1))$, where $T(k-1)$ is the number removed on the previous move and f is a non-decreasing function. (Also, we exclude the possibility of removing all counters on the first move in the one-pile version.)

For example, we might have a game in which one can remove only up to twice as many counters as were removed on the previous move; in this case $f(n) = 2n$ (cf. [8] and [5], Section 1.2.8, #37). The idea of Schwenk's theory is to define a sequence H_i such that $H_1 = 1$ and $H_{k+1} = H_k + H_j$, where j is the smallest index such that $f(H_j) \geq H_k$. In our example, $H_1 = 1$, $H_2 = H_1 + H_1 = 2$, $H_3 = H_2 + H_1 = 3$, and it is not too difficult to see that in general $H_{k+1} = H_k + H_{k-1}$. Thus, the sequence H_i is the Fibonacci sequence in our case.

Schwenk demonstrates that each number N is represented uniquely by a sum of H_is, where $N = H_{j_1} + H_{j_2} + \cdots + H_{j_m}, j_1 < j_2 \cdots < j_m$, and $f(H_{j_i}) < H_{j_{i+1}}$. In our example, this statement becomes *Zeckendorf's Theorem*, which states that each number is represented uniquely as a sum of non-consecutive Fibonacci numbers. The number of elements in this sum is the *norm* $|N|$ of N. Clearly, $|N| = 0$ if and only if $N = 0$. In [8], Schwenk proves

Theorem A The first player to be able to reduce the norm has a winning strategy. The only winning strategy is to reduce the norm.

Schwenk's idea is to make up N as a sum of "losing" positions, H_js. In fact, each H_j as a starting position is a P-position of norm one which cannot be decreased on the first move. If a player can reduce $N = H_{j_1} + \cdots + H_{j_m}$ with $j_1 < j_2 \cdots < j_m$ and $|N| = m \geq 2$, by removing H_{j_1} then the other player cannot immediately remove H_{j_2}, for $f(H_{j_1}) < H_{j_2}$. Therefore, any legal removal by the other player will not decrease the norm. Also note that all winning removals can be characterized as partial sums of $H_{j_1} + H_{j_2} + \cdots + H_{j_m}$ (cf. [5]).

To extend the Schwenk theory to Euclid, we need to employ a new function. Define $h(N), N = H_{j_1} + \cdots + H_{j_m}$, as the least H_i such that $f(H_i) \geq H_{j_1} = H_{j_1}(N)$. The definition implies that $|h(a)| = 1, f(h(a)) \geq H_{j_1}(a)$, and if $H_i < h(a)$ then $f(H_i) < H_{j_1}(a)$.

We now consider the Euclid position $[a_0, a_1, \ldots, a_n]$. If $n = 0$ then we are back to the one-pile version covered by Theorem A [8]. We now have the following theorem:

Theorem 3 *Assume that* $n \geq 1$.

(1) The position $[a_0, a_1, a_2, \ldots, a_{n-1}, a_n]$ *has the same winner as the first position in the series* $[a_0, \ldots, a_{n-1} - h(a_n)]$, $[a_0, \ldots, a_{n-2} - h(a_{n-1})]$, $\ldots$, $[a_0, \ldots, a_i - h(a_{i+1})]$, $\ldots$, $[a_0 - h(a_1)]$ *in which the last partial quotient* $a_i - h(a_{i+1})$ *is non-negative.*

(2) If $a_i - h(a_{i+1}) < 0$ *for* $i = 0, 1, \ldots, n-1$, *the first player is always the winner.*

Any position can be reduced to a shorter position (i.e., one with fewer terms in the Euclid position representation) with the same winner in one of these two ways. By recursively applying these reductions, we eventually reach a one-pile position whose winner is known by Schwenk's theory.

Proof. We first prove the second assertion of the theorem. Suppose that in $[a_0, a_1, \ldots, a_n]$, $a_i - h(a_{i+1}), i = 0, 1, \ldots, n-1$, is always negative. Then the first player can win by adopting the strategy of reducing the norm of a_0 by *one* on each move. In this way, that player will eventually reduce a_0 to zero by removing some H_i satisfying $H_i \leq a_0 < h(a_1)$. We have $f(H_i) < H_{j_1}(a_1)$ by definition, so that even the smallest norm-reducing removal is precluded and the next player will be unable to reduce the norm of a_1. Therefore, the first player will be able to apply the same strategy of reducing the norm by one. It is clear that the first player will win each subgame in this fashion and thus the entire game.

We can now prove the first part of the theorem. Assume that a_i is the greatest index such that $a_i - h(a_{i+1})$ is non-negative. Then, the winner of $[a_0, a_1, \ldots, a_i - h(a_{i+1})]$, say Player A, will be able to reach the position $[h(a_{i+1}), a_{i+1}, \ldots, a_n]$. We must show that this is a P-position. Clearly, $h(a_{i+1})$ has norm one, so the only way for Player B to reduce it is to eliminate it entirely. This, however, allows Player A to reduce the norm of a_{i+1} by removing $H_{j_1}(a_{i+1})$, by the definition of $h(a_{i+1})$ (for $f(h(a_{i+1})) \geq H_{j_1}(a_{i+1})$). In this case, (2) shows that Player A wins, since each $a_l - h(a_{l+1}), l \geq i+1$, is now negative. If, on the other hand, Player B does not reduce the norm of $h(a_{i+1})$ but instead moves to $[k, a_{i+1}, \ldots, a_n]$, then Player A can win by adopting the strategy of reducing the norm of k by one on each turn (cf. Theorem A guarantees that Player B, by missing the opportunity to decrease the norm, allows Player A to do so). In this way, Player A eventually will be the one to move to $[a_{i+1}, \ldots, a_n]$. However, Player A will reach this position by removing some $H_j < h(a_{i+1})$ (recall that Player B faced $[h(a_{i+1}), a_{i+1}, \ldots, a_n]$). Thus, again by the definition of $h(a_{i+1})$ (for $f(H_j) < H_{j_1}(a_{i+1}) \leq a_{i+1}$), Player B will be unable to reduce the norm of a_{i+1} and loses by (2). □

Example 1. In all of the following examples assume that the move function is $f(n) = 2n$. Consider the position $(25, 87)$, i.e., $[3, 2, 11]$. This reduces to $[3, 2 - h(11)] = [3]$ by (1). Thus, the first player can win by moving to the P-position $[2, 11]$, i.e., $(12, 25)$, first. How the first player wins from here by playing second is not directly given by Theorem 3 but is outlined in its proof. If the second player moves to $(1, 12) = [11]$, the first player can reduce the norm of $11 = 3 + 8$ by moving to $[8]$. If the second player instead moves to $[1, 11] = (12, 13)$, the first player moves to $[11] = (1, 12)$. Then, the second player can remove one at most twice and is unable to reduce the norm. Thus, the first player wins in all cases, as Theorem 3 claims.

Example 2. Now consider the position $[1, 5, 11]$. This reduces to $[1, 5 - h(11)] = [1, 3]$ by (1). Now, since $1 - h(3) = -1$, we use (2) and reduce to $[1]$. Thus, the first player wins by removing 1 and moving to $[5, 11]$. Now, any response allows the first player to reach the P-position $[2, 11]$ and win as above.

Example 3. Finally, consider the position $[2, 3, 13]$. We first try $[2, 3 - h(13)]$, but because $3 - h(13) = -5$ is negative, we instead reduce to $[2 - h(3)] = [0]$. Thus, this is a P-position and a loss for the first player. The reader can check all the variations.

6. The Euclid of Three Numbers

Consider an extension of Euclid in which a position is (a, b, c), where a, b, c are integers. There are many ways to generalize the rules of standard Euclid. Suppose we decree that a legal move is to remove the same multiple of the smallest integer from each of the larger two. Then it is not hard to see that the idea of Theorem 1 applies.

Assume without loss of generality that in (a, b, c) we have $a \leq b \leq c$. Then, assuming that $b > 2a$, we choose n such that $a < b - na < 2a$. Then one of the moves $(a, b - na, c - na)$ or $(a, b - (n+1)a, c - (n+1)a)$ will win, for if the latter does not, the first player can take the former and force his opponent to move there.

The reader may object that this is hardly the most natural way of generalizing Euclid to three numbers, and the author sympathizes. To me, the most natural extension would be to make a legal move the decreasing of any of the three integers by a multiple of any other, provided the result remains positive. Unfortunately, I have not been able to find any winning strategies in this variation of Euclid, and doing this remains the most intriguing unsolved problem in the domain. The general theme that a player with many options has a winning strategy seems to hold, but there are some surprising P-positions; for example, $(4, 9, 16)$.

7. Misère forms

As we mentioned in Section 1, the misère form of Euclid is a win for the first player with a choice of moves. Our next theorem gives a winning strategy.

Theorem 4 *The first player to have a choice can win misère Euclid by adopting the following strategy: when faced with the position $[a_i, a_{i+1}, \ldots, a_n]$, with $a_i \geq 2$, make the same move as in the unrestricted version if at least one of $a_{i+1}, \ldots, a_n \geq 2$. Otherwise, play so as to leave an odd number of ones (whereas in unrestricted version one would leave an even number). This strategy works for Euclid with no restriction, restriction sets Λ_k, and other equivalent restriction sets.*

Proof. If at least one of $a_{i+1}, \ldots, a_n \geq 2$, playing as in the unrestricted version will ensure that the first player will make the next choice. Thus, the position will be reduced to a smaller one in which the first player still has control. If $a_{i+1} = a_{i+2} = \cdots = a_n = 1$, there are no more choices. In this case, playing to leave an odd number of ones will ensure that the second player's moves always leave an even number, eventually making the last move by leaving zero, and giving the first player victory. Note that it is always possible for the first player to leave an odd number, because $[1, a_{i+1}, \ldots, a_n]$ and $[a_{i+1}, \ldots, a_n]$ are two options. This completes the proof for the unrestricted version. This strategy also applies to the misère games with restriction set Λ_k and other equivalent restriction sets (in the sense described in Section 4). All we need to do is to reduce all the partial quotients $\bmod(k+1)$ and apply the above strategy. □

Acknowledgments

The author wishes to thank Occidental College and, in particular, the Undergraduate Research Center for approving this research and providing funding. He especially wishes to thank his faculty mentor Tamás Lengyel for many suggestions and careful reading of this paper as well as for suggesting this research topic. He also thanks the referee for helpful suggestions which improved the presentation of the paper.

References

[1] Elwyn R. Berlekamp, John H. Conway, and Richard K. Guy. *Winning Ways for Your Mathematical Plays*. London, New York: Academic Press, 1982.

[2] John H. Conway. *On Numbers and Games*. London, New York, San Francisco: Academic Press, 1976.

[3] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics*. Second edition. Reading, Massachusetts: Addison-Wesley Publishing Company, 1994.

[4] Arthur Holshouser, Harold Reiter, and James Rudzinski. *Pilesize Dynamic One-Pile Nim and Beatty's Theorem*. Integers, Electronic Journal of Combinatorial Number Theory **4**(2004), #G03, 1–13.

[5] Donald E. Knuth. *The Art of Computer Programming*, vol. 1: Fundamental Algorithms, Third Edition, Reading, Massachusetts: Addison-Wesley Publishing Company, 1997.

[6] Tamás Lengyel. *A Nim-Type Game and Continued Fractions*. Fibonacci Quarterly **41**(2003), 310–320.

[7] Tamás Lengyel. *On calculating the Sprague-Grundy function for the game Euclid*. To appear in the 11th International Conference on Fibonacci Numbers and Its Applications, 2004.

[8] Allen J. Schwenk. *Take-Away Games*. Fibonacci Quarterly **8**(1970), 225–234.

[9] Michael Zieve. *Take-Away Games*. Games of No Chance, MSRI Publications, vol. 29, 1996.