Structure and Scheduling of the Sporadic (16, 2, 6) Fixed-Win Tournament

Jack van Ryswyck
math@javhar.net

Abstract

In a fixed-elimination tournament, a team is eliminated upon suffering exactly $\mathcal{L}$ losses. A tournament additionally has the fixed-win property if its champion is guaranteed to have won exactly $\mathcal{W}$ matches. Fayers [1] showed that, apart from single-elimination tournaments and two-team playoff series, only one such tournament exists: the sporadic $\mathcal{L} = 2, \mathcal{W} = 6$ case with exactly 16 participants.

This paper investigates the structure and scheduling of this exceptional tournament. The conditions established by Fayers force a decomposition into upper and lower brackets, an organization into rounds, and a three-team semifinal endgame.

The pairing of teams within these rounds presents a more intricate problem when minimizing repeat encounters between teams. Careful scheduling can avoid repeats during the bracket phase. This can be achieved using a fixed schedule through the first three rounds, but adaptive scheduling is required thereafter. More strongly, *every* possible three-round history admits an adaptive continuation that avoids repeats in the remaining bracket rounds, regardless of match outcomes. A constructive algorithm provides this guarantee.

Past the bracket phase of the tournament, up to three endgame matches remain. Adversarial outcomes can trivially force at least two of the endgame matches to repeat. Exhaustive minimax analysis shows that they can force three repeats, regardless of how the matches are scheduled.

Finally, a seeded scheduling scheme extends conventional single-elimination seeding to the bracket rounds while respecting repeat-avoidance constraints.

1. Introduction

1.1. Problem Statement

Following Fayers [1], an *$\mathcal{L}$-elimination tournament* is a competition in which each game has exactly one winner, a team is eliminated on its $\mathcal{L}$-th loss, and the tournament ends when only one team remains. The lone remaining team is crowned champion. Such a tournament has the *fixed-win property* with parameter $\mathcal{W}$ if the champion is guaranteed to have won exactly $\mathcal{W}$ games, regardless of how the individual games play out.

Denote a fixed-elimination fixed-win tournament for n teams as $(n, \mathcal{L}, \mathcal{W})$. Two families are obvious:

- $\mathcal{L} = 1, n = 2^{\mathcal{W}}$: ordinary single elimination.
- $\mathcal{L} = \mathcal{W}, n = 2$: a "best of $2\mathcal{W} - 1$" contest between two teams, common in playoff series.

The Fayers paper shows that outside of these families there exists only one other fixed-win tournament.

Theorem 1.1 (Fayers [1])**.** *The only possible $(n, \mathcal{L}, \mathcal{W})$ tournament with $\mathcal{L} \geq 2$ and $n > 2$ is $(16, 2, 6)$.*

The proof is by recursion on two numbers associated with each of the participants: the team's *win target* w and *elimination number* ℓ. These counts indicate how many more matches the team needs to win in order to be the tournament champion, and how many "lives" it has remaining, which is the number of additional losses it would take to eliminate the team from the tournament. This will be denoted as $\left[\begin{smallmatrix} w \\ \ell \end{smallmatrix}\right]$. The tournament starts with every team at $\left[\begin{smallmatrix} \mathcal{W} \\ \mathcal{L} \end{smallmatrix}\right]$. As Fayers shows, classifying all possible tournaments yields Theorem 1.1.

1.2. Motivation

Round-robin tournaments aside, competitive events generally need a fast, decisive elimination format. Single elimination is fast but harsh: one bad day, one bad bounce, and a top competitor is gone. Even barring such upsets, there is no guarantee that the two best teams will reach the final stages of the tournament. They may simply meet each other early.

Double elimination tournaments exist, but require some careful scheduling decisions. First, true double elimination ought to imply that a team is eliminated when *and only when* it has lost twice. A team that has lost only once should still be able to recover from the loss and win the tournament. This precludes repechage systems, where teams are typically relegated to competing for lower positions following a loss. It also precludes systems that can remove a team by transitive loss, when an opponent that defeated them subsequently loses to another competitor. That would eliminate a team despite having lost only one match, which is not strictly double elimination. The fate of a participant ought not to depend on the results of others.

When a tournament follows strict double elimination, a second issue appears. Competitors are typically divided into a "winners bracket", containing the undefeated teams, and a "lower bracket". The winners bracket is reduced by half every round, but the lower bracket is replenished with competitors dropping down from the winners bracket. The constant replenishment requires extra lower bracket rounds to bring the size of the lower bracket back down.

The motivation for imposing a fixed-win guarantee is that a match loss is already sufficiently punished by having spent the safety buffer. There is no need to have the number of wins required to be champion depend on the path taken through the tournament. The (16, 2, 6) tournament threads this needle. As we show in Section 4, it forces a genuine division into an upper bracket and a lower bracket. An optimal pairing mechanism described in Section 5 ensures that competitors do not need to play each other twice during the bracket phase of the tournament.

2. The (16, 2, 6) tournament

The unique existence of the (16, 2, 6) tournament is proven in Fayers [1]. But how can (16, 2, 6) be implemented? The diagram below shows a valid tournament structure.

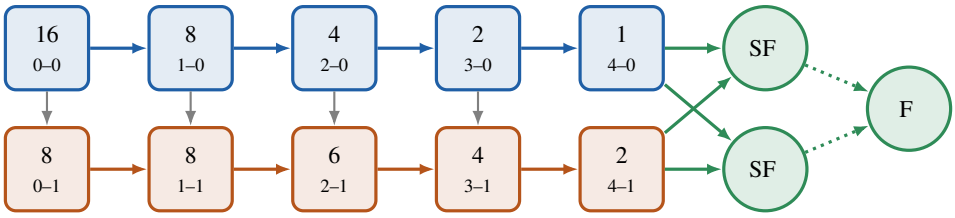


Figure 1: The (16, 2, 6) tournament structure.

The boxes in the diagram indicate the number of teams, along with their win-loss record. The *upper bracket* is formed by the blue boxes, containing only undefeated teams. The red boxes are the *lower bracket*, with teams that have suffered one loss. The tournament starts with all 16 teams in the 0-0 box.

Within each box, teams are played against each other in pairs. A match sends the winner horizontally to the next box, and drops the other team down vertically. Dropping down from a red box means elimination from the tournament.

The brackets produce one team, say U, at 4-0, and two teams, say L_1 and L_2, at 4-1. In the semifinals there is a twist: The undefeated team U appears in *both* semifinals. The winners of the semifinals play in the final, unless U wins both semifinals, in which case it does not need to play itself in the final.

From a scheduling perspective it seems somewhat objectionable that team U must play two semifinals, where L_1 and L_2 each only play one. On the other hand, L_1 and L_2 did already play an extra round prior to the semifinal: they each played 5 matches to get there, while U only played 4.

It may also seem odd that the final itself might not happen. However, this is no different from a playoff series, where "game 7" might not happen.

The total number of matches played is 30 or 31: There will be 15 teams eliminated with two losses each, while the champion incurred at most one defeat, for a total of 30 or 31 losses. Each match contributes exactly one loss to that number. In general, any $\mathcal{L}$-elimination tournament for n participants requires between $(n-1)\mathcal{L}$ and $n\mathcal{L}-1$ matches.

It can be verified directly that the fixed-win property indeed holds. Teams L_1 and L_2 can each feature in at most two of the remaining matches and must win both of them. Team U could appear in all three remaining matches and must win two of them. It all works out correctly – but is it the only way to schedule a (16, 2, 6) tournament?

We shall show that the structure of the tournament, modulo causality-preserving rescheduling, must indeed be organized this way. The actual pairing of the teams within each box is nontrivial, and will be addressed in Sections 5 and 6.

3. Preliminaries

3.1. Tournament States

Denote a *tournament state* as a collection of win targets and elimination numbers:

$$\left\{ \begin{array}{ccc} w_1 & \cdots & w_n \\ \ell_1 & \cdots & \ell_n \end{array} \right\}.$$

Definition 3.1. A tournament state is valid if a match scheduling policy exists such that, regardless of the match outcomes, team i wins the tournament iff it wins exactly w_i matches, and is eliminated iff it loses exactly ℓ_i matches. All ℓ_i must be positive. All w_i must be positive if $n > 1$, and zero if $n = 1$.

The validity of a state is defined recursively. A state is valid iff there is a pair (i, j) of participants such that both possible outcomes of a match between i and j lead to another valid state. The two successor states are reached by decrementing w_i and ℓ_j when i wins, and decrementing ℓ_i and w_j when j wins. A team whose elimination number reaches zero is removed. The recursion bottoms out in any state that contains only one team, which is valid iff its win target is zero and its elimination number is positive. Such a final state identifies the champion as the lone remaining team.

Whenever a state may contain different elimination numbers we will commonly group them and insert a separator, for instance: $\left\{ \begin{smallmatrix} w_1 & \cdots & w_r \\ 1 & \cdots & 1 \end{smallmatrix} \,\middle|\, \begin{smallmatrix} w_{r+1} & \cdots & w_n \\ 2 & \cdots & 2 \end{smallmatrix} \right\}$. The grouping and separator are only there for clarity; they have no functional meaning.

3.2. Results from Fayers

Several results from Fayers [1] will be essential here. Three main propositions used here involve the use of a dyadic sum.

Definition 3.2. The *dyadic sum* $\mathcal{D}_k$ of a state $S = \left\{ \begin{smallmatrix} w_1 & \cdots & w_n \\ \ell_1 & \cdots & \ell_n \end{smallmatrix} \right\}$ is $\mathcal{D}_k(S) := \sum_{\ell_i = k} 2^{-w_i}$.

Proposition 3.3 (Fayers, Prop. 2.1)**.** *For $n = 2$, the only valid states are $\left\{ \begin{smallmatrix} u \\ v \end{smallmatrix} \mid \begin{smallmatrix} v \\ u \end{smallmatrix} \right\}$.*

This corresponds to any state in a two-team playoff series, where each team's win target equals the other team's elimination number.

Proposition 3.4 (Fayers, Prop. 2.2)**.** *The state $S = \left\{ \begin{smallmatrix} w_1 & \cdots & w_n \\ 1 & \cdots & 1 \end{smallmatrix} \right\}$ is valid iff $\mathcal{D}_1(S) = 1$.*

Proposition 3.5 (Fayers, Prop. 2.4)**.** *For $n \geq 3$, the state $S = \left\{ \begin{smallmatrix} w_1 & \cdots & w_{n-1} \\ 1 & \cdots & 1 \end{smallmatrix} \,\middle|\, \begin{smallmatrix} w_n \\ 2 \end{smallmatrix} \right\}$ is valid iff $w_n \in \{1, 2\}$ and $\mathcal{D}_1(S) = \frac{1}{4} w_n$.*

Proposition 3.6 (Fayers, Prop. 2.6)**.** *For $n \geq 3$ and $r \geq 2$, the state $S = \left\{ \begin{smallmatrix} w_1 & \cdots & w_{n-r} \\ 1 & \cdots & 1 \end{smallmatrix} \,\middle|\, \begin{smallmatrix} w_{n-r+1} & \cdots & w_n \\ 2 & \cdots & 2 \end{smallmatrix} \right\}$ is valid iff $\mathcal{D}_2(S) = \frac{1}{4}$ and $\mathcal{D}_1(S) = \frac{3}{4} - \sum_{\ell_i = 2} w_i \, 2^{-w_i - 1}$.*

For our purposes, the relevant part of the conditions on $\mathcal{D}_1(S)$ is that they depend only on the win targets of the $\ell = 2$ teams.

4. Structural Properties of the (16, 2, 6) Tournament

4.1. Admissible Matches

The conditions identified by Fayers impose considerable structure on the $(16, 2, 6)$ tournament, to be examined in this section. The structure is dominated by a division into two brackets.

Definition 4.1. Upper and Lower Bracket: An *upper bracket match* is a match played between two teams at $\ell = 2$. A *lower bracket match* is a match between two teams at $\ell = 1$. A *cross-bracket match* is a match between a team at $\ell = 2$ and a team at $\ell = 1$.

In the proposed tournament structure of Section 2, cross-bracket matches do happen in the semifinal. We shall show that they cannot happen before the semifinal, due to a very fundamental property that emerges.

Proposition 4.2. *While there are at least two undefeated teams left, no cross-bracket matches can be played.*

Proof. Suppose there are at least two teams at $\ell = 2$, and we play a match between a team at $\left[{}^{u}_{1}\right]$ and a team at $\left[{}^{v}_{2}\right]$. The presence of $\left[{}^{u}_{1}\right]$ means that $n \geq 3$, so Proposition 3.6 applies.

If $n = 3$, the state is $S = \left\{{}^{u}_{1} \mid {}^{v}_{2}\ {}^{w}_{2}\right\}$. Proposition 3.6 requires $\mathcal{D}_2(S) = \frac{1}{4}$, forcing $v = w = 3$, and it requires $\mathcal{D}_1(S) = \frac{3}{4} - 2(3 \cdot 2^{-4}) = \frac{3}{8}$. But $\mathcal{D}_1(S) = 2^{-u}$, which cannot equal $\frac{3}{8}$. So the state would not be valid.

Therefore we must have $n > 3$. If $\left[{}^{v}_{2}\right]$ beats $\left[{}^{u}_{1}\right]$, there will still be at least two $\ell = 2$ teams. Since we had at least four teams and eliminated one, there will still be at least three teams in total. Thus Proposition 3.6 still applies. The winning team is still at $\ell = 2$ but its win target has decreased by one, which changes $\mathcal{D}_2$. Proposition 3.6 required $\mathcal{D}_2$ to stay invariant at $\frac{1}{4}$, so the new state is invalid. □

Lemma 4.3. *Any match between two teams at $\ell = 1$ must feature two teams with the same win target.*

Proof. Let the participants of the match be X at $\left[{}^{x}_{1}\right]$, and Y at $\left[{}^{y}_{1}\right]$. If they are the only two remaining teams, then Proposition 3.3 already specifies $x = y = 1$. For the remainder, we need only consider the case with at least three remaining teams.

The joint contribution of X and Y to $\mathcal{D}_1$ is $2^{-x} + 2^{-y}$. If X wins, then X will move to $\left[{}^{x-1}_{1}\right]$, and Y will be eliminated. Their joint contribution to $\mathcal{D}_1$ then becomes $2^{-(x-1)} = 2 \cdot 2^{-x}$. If Y wins, it similarly becomes $2 \cdot 2^{-y}$.

Propositions 3.4, 3.5, 3.6 cover all possible states in a (16, 2, 6) tournament when at least three teams remain: every team at $\ell = 1$, one team at $\ell = 2$, and multiple teams at $\ell = 2$. They all require $\mathcal{D}_1$ to be fixed, or dependent only on the win targets of the $\ell = 2$ teams. If one of the propositions applies both before and after the match between X and Y, then $\mathcal{D}_1$ must remain unchanged by the match between X and Y. We must have $2^{-x} + 2^{-y} = 2 \cdot 2^{-x} = 2 \cdot 2^{-y}$ and therefore $x = y$.

The only case where one of the propositions applies before but not after the match is when the match reduces the tournament from three teams to two, so 3.5 and 3.6 stop applying, with an undefeated team still alive, so 3.4 does not apply. The match then started in state $\left\{{}^{x}_{1}\ {}^{y}_{1} \;\middle|\; {}^{z}_{2}\right\}$ and ends in state $\left\{{}^{x-1}_{1} \;\middle|\; {}^{z}_{2}\right\}$ or $\left\{{}^{y-1}_{1} \;\middle|\; {}^{z}_{2}\right\}$. From Proposition 3.3 we then have $x = y = 3$. □

Theorem 4.4. *While there are at least two undefeated teams, every match must be between two teams with the same win-loss record.*

Proof. Suppose there are at least two undefeated teams left. Proposition 4.2 demands that any match played will involve two teams at the same elimination number ℓ. Let X at $\left[{}^{x}_{\ell}\right]$ play Y at $\left[{}^{y}_{\ell}\right]$. Without loss of generality, X wins the match.

The case $\ell = 1$ is already covered by Lemma 4.3. Consider the remaining case, $\ell = 2$. Before the match, Proposition 3.6 requires $\mathcal{D}_2 = \frac{1}{4}$. If X and Y are the only two teams at $\ell = 2$, then $2^{-x} + 2^{-y} = \mathcal{D}_2 = \frac{1}{4}$, which, for integer x, y, implies $x = y = 3$. If on the other hand there are other teams at $\ell = 2$, then after the match there will still be at least two undefeated teams and Proposition 3.6 ensures that $\mathcal{D}_2$ is still $\frac{1}{4}$. The same reasoning as in Lemma 4.3 now applies to $\mathcal{D}_2$: The joint contribution of X and Y to $\mathcal{D}_2$ changes from $2^{-x} + 2^{-y}$ to $2 \cdot 2^{-x}$ or $2 \cdot 2^{-y}$, so again $x = y$.

In any case, X and Y are at the same $\left[{}^{w}_{\ell}\right]$ and therefore at the same win-loss record. □

4.2. The Semifinals

Proposition 4.2 prohibits cross-bracket matches while at least two undefeated teams remain. Once the upper bracket has produced a single survivor, however, cross-bracket matches become possible. We shall show that the upper bracket winner then has a specific sequence of two or three matches left, which correspond to the semifinals and final seen in Section 2.

First, a simple observation.

Lemma 4.5. *While at least three teams remain in the tournament, no team with $w = 1$ may play a match.*

Proof. A match can eliminate at most one team. If a team at $\left[{}^{1}_{\ell}\right]$ wins a match while at least three teams remain, it reaches $w = 0$ while another team is still alive. This contradicts the definition of a valid tournament state. □

We can now characterize cross-bracket matches.

Proposition 4.6. *In any cross-bracket match, the $\ell = 1$ team will enter the match at $\left[{}^{2}_{1}\right]$.*

Proof. By Proposition 4.2, a cross-bracket match is only possible when there is only one $\ell = 2$ team remaining. According to Proposition 3.5, this final undefeated team will be at $\left[{}^{1}_{2}\right]$ or $\left[{}^{2}_{2}\right]$.

1. Case $\left[{}^{1}_{2}\right]$: Since the undefeated team is at $w = 1$, by Lemma 4.5 it is only allowed to play when there is just one other team left. Proposition 3.3 then specifies that the other team must be at $\left[{}^{2}_{1}\right]$.
2. Case $\left[{}^{2}_{2}\right]$ with $n > 3$: Proposition 3.5 applies, so $\mathcal{D}_1 = \frac{2}{4}$. Let $\left[{}^{2}_{2}\right]$ play against $\left[{}^{u}_{1}\right]$. If $\left[{}^{2}_{2}\right]$ wins, it will move to $\left[{}^{1}_{2}\right]$. At least three teams will still be alive after the match, so Proposition 3.5 continues to apply. $\mathcal{D}_1$ is now $\frac{1}{4}$. Only 2^{-u} was removed from $\mathcal{D}_1$, so we had $u = 2$ and therefore the opponent was at $\left[{}^{2}_{1}\right]$.
3. Case $\left[{}^{2}_{2}\right]$ with $n = 3$: The state going into the match is $\left\{{}^{u}_{1}\ {}^{v}_{1}\ \middle|\ {}^{2}_{2}\right\}$ with $\left[{}^{2}_{2}\right]$ playing against $\left[{}^{u}_{1}\right]$. If $\left[{}^{2}_{2}\right]$ wins then the state becomes $\left\{{}^{v}_{1}\ \middle|\ {}^{1}_{2}\right\}$, so $v = 2$ by Proposition 3.3. If $\left[{}^{u}_{1}\right]$ wins then the state becomes $\left\{{}^{u-1}_{1}\ {}^{2}_{1}\ {}^{2}_{1}\right\}$, and by Proposition 3.4 we have $u - 1 = 1$. So both $\ell = 1$ teams were at $\left[{}^{2}_{1}\right]$.

In any case, the undefeated team plays against a $\left[{}^{2}_{1}\right]$ opponent. □

We can now characterize the first cross-bracket match, which is the next match that the upper bracket winner plays.

Lemma 4.7. *After the last upper bracket match, the upper bracket winner will be at $\left[{}^{2}_{2}\right]$. Its next match will be versus an opponent at $\left[{}^{2}_{1}\right]$.*

Proof. By Proposition 4.2, the first cross-bracket match cannot happen until there is only one undefeated team remaining. So the final upper bracket match must take place before the first cross-bracket match.

The very first match of the tournament created a lower bracket participant. Subsequently, each lower bracket match keeps one of its lower bracket teams alive, so at least one lower bracket participant must remain until the first cross-bracket match. Therefore at least three teams are alive when the final upper bracket match happens.

Theorem 4.4 requires that the participants in the final upper bracket match have equal win targets, and by Proposition 3.6 their contributions to $\mathcal{D}_2$ sum to $\frac{1}{4}$. They must therefore both be at $\left[{}^{3}_{2}\right]$, and their match leaves a single undefeated team at $\left[{}^{2}_{2}\right]$. Since it has no undefeated opponent remaining, it cannot play again without crossing brackets. Therefore the *first* cross-bracket match has the undefeated team entering at $\left[{}^{2}_{2}\right]$. By Proposition 4.6 its opponent will be at $\left[{}^{2}_{1}\right]$. □

This $\left[{}^{2}_{2}\right]$ versus $\left[{}^{2}_{1}\right]$ match corresponds to the first semifinal from Section 2. The $\left[{}^{2}_{2}\right]$ team is the upper bracket winner U, and the $\left[{}^{2}_{1}\right]$ team is the lower bracket team L_1.

For a match between teams A and B, we write $A \succ B$ to indicate that A won the match.

Theorem 4.8. *When only one undefeated team U remains, its next two matches are against two distinct teams L_1 and L_2, both at $\left[{}^{2}_{1}\right]$ when they meet U. Any further match is the final between the two teams that won those matches, both at $\left[{}^{1}_{1}\right]$.*

Proof. Let the first cross-bracket match feature the upper bracket team U and the lower bracket team L_1. Lemma 4.7 showed that U is at $[^2_2]$ and L_1 is at $[^2_1]$. Let $\mathcal{V}$ be the set of all other remaining teams, all of whom must be at $\ell = 1$. The set $\mathcal{V}$ is nonempty, as $\{^2_1 \mid ^2_2\}$ is not a valid state. When $\mathcal{V}$ contains only one team, that team is to be identified as L_2. However, it is possible that $\mathcal{V}$ contains multiple teams. Consider the two possible outcomes of the match U–L_1:

- Case $U \succ L_1$: L_1 is eliminated and U moves to $[^1_2]$. While $|\mathcal{V}| \geq 2$, U cannot play again due to Lemma 4.5. The teams in $\mathcal{V}$ must play among themselves until $\mathcal{V}$ is reduced to L_2. Then L_2 will be at $[^2_1]$ by Proposition 3.3. The next match is U–L_2, as no other teams are still alive. If $U \succ L_2$ then L_2 is eliminated and the tournament is concluded with U as the undefeated champion. If $L_2 \succ U$ then both teams will be at $[^1_1]$ and will play each other in the final.

- Case $L_1 \succ U$: L_1 moves to $[^1_1]$ and U moves to $[^2_1]$. All teams are then at $\ell = 1$, so Proposition 3.4 requires $\mathcal{D}_1(\mathcal{V}) = \frac{1}{4}$, and Lemma 4.3 requires that every remaining match feature two teams at the same win target. While $|\mathcal{V}| \geq 2$, all teams in $\mathcal{V}$ must have $w \geq 3$ since their 2^{-w} sum up to $\frac{1}{4}$, and therefore the only admissible match is within $\mathcal{V}$. As in the case above, the teams in $\mathcal{V}$ must play among themselves until L_2 emerges. L_2 will then be at $[^2_1]$ since $\mathcal{D}_1(\mathcal{V})$ still equals $\frac{1}{4}$. The next match is U–L_2, being the only two teams at the same win target. That match will eliminate one of its participants, leaving the tournament with only two teams, both at $[^1_1]$, to play the final.

In either case, the scenario outlined in the theorem is realized. □

The tournament therefore concludes with the following sequence of matches:

1. U–L_1 with L_1 at $[^2_1]$;
2. matches within $\mathcal{V}$, if any;
3. U–L_2 with L_2 at $[^2_1]$;
4. a possible final with both at $[^1_1]$.

The matches U–L_1 and U–L_2 are the semifinals indicated in Section 2. The final happens iff $L_1 \succ U$ or $L_2 \succ U$.

4.3. Rounds and Scheduling

The results in Sections 4.1 and 4.2 specify which teams may face each other in a match. But what of the order in which the matches are to be scheduled? This section will address the scheduling considerations.

By Theorem 4.4 and Lemma 4.3, every upper bracket match and every lower bracket match is between two teams with the same win-loss record. The only matches that feature two teams with a different win-loss record are the cross-bracket matches. Together with Proposition 4.2, Theorem 4.8 showed that the only cross-bracket matches are the first semifinal, U–L_1, and possibly the second semifinal, U–L_2. This naturally partitions the matches into *rounds*.

Definition 4.9. The **semifinal** consists of the next two matches played by the upper bracket winner, namely U–L_1 and U–L_2. The **final** is the possible match between the winners of the semifinal matches. All other matches are grouped as follows:

- **Upper bracket round:** All matches between teams at a given $[^w_2]$.
- **Lower bracket round:** All matches between teams at a given $[^w_1]$.
- **Bracket round:** An upper bracket round or a lower bracket round.
- **Round i:** All matches between teams that have played $i - 1$ previous matches. This comprises a pair of an upper bracket round and a lower bracket round, where these exist. A Round i match has both teams entering at $w + \ell = 9 - i$.
- **Bracket phase:** All matches in bracket rounds.

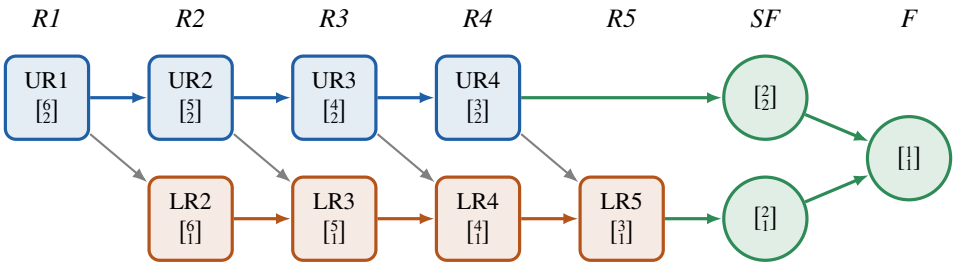


Figure 2: Names to be used to indicate the rounds and bracket rounds. The square boxes are bracket rounds, with the $\left[\begin{smallmatrix} w \\ \ell \end{smallmatrix}\right]$ numbers indicating the win target and elimination number of both teams entering a match in that bracket round. The three circles are not bracket rounds, but the semifinal and the final. The two SF circles indicate the three teams entering the two semifinal matches, with $\left[\begin{smallmatrix} 2 \\ 2 \end{smallmatrix}\right]$ referring to the upper bracket winner's status when entering the first of its two semifinals.

The naming of the rounds is depicted in Figure 2. Note that the diagram is skewed relative to the one shown in Figure 1: The bracket rounds are arranged in columns indicating the total number of matches played, rather than the win target. In the upper bracket there is no UR5, as UR4 already produced the upper bracket winner, which proceeds directly to the semifinal. In the lower bracket, the two other semifinalists emerge from LR5.

The arrows in Figure 2 indicate *causal links*.

Definition 4.10. Causally linked matches: Two matches M_i and M_j are *causally linked* when a participant in M_i can subsequently appear in M_j. The two matches need not be in consecutive rounds; there could be other matches in between.

The scheduling of matches need only respect the causal links. It is therefore entirely admissible, for instance, to play a UR2 match as soon as two participants have progressed from UR1, even when the other matches in UR1 have not yet been concluded. It is even possible to finish UR4 before starting LR2. However, since causal links necessarily point from earlier rounds to later rounds, it is always possible to schedule the matches such that each round only begins after the previous round is concluded. This may limit the information available to an adaptive pairing, but Section 5.5 will show that for repeat avoidance this costs nothing.

This comes into focus particularly when scheduling the two semifinal matches. The results in Section 4.2 allow for the U–L_1 semifinal to go ahead even when there are still multiple other lower bracket teams remaining. The proof of Theorem 4.8 only required that this collection $\mathcal{V}$ of lower bracket teams play among themselves until one of them emerges as the second semifinalist L_2. Since none of those $\mathcal{V}$ matches involve U or L_1, the causal links are only from the first semifinal to the second semifinal, and from the $\mathcal{V}$ matches to the second semifinal.

By Lemma 4.3, the matches within $\mathcal{V}$ preserve $\mathcal{D}_1(\mathcal{V})$. Moving them before the first semifinal therefore preserves the validity of the intermediate states, by Proposition 3.5. We can therefore, "without loss of causality", schedule all the $\mathcal{V}$ matches to take place before the first semifinal. Having done so, the semifinal matches and the possible final match are the last matches to be played in the tournament. The upper bracket and lower bracket rounds are all concluded before the first semifinal. That first semifinal starts in the state $\{\begin{smallmatrix} 2 & 2 \\ 1 & 1 \end{smallmatrix} \mid \begin{smallmatrix} 2 \\ 2 \end{smallmatrix}\}$. This is the **canonical semifinal**.

4.4. Tournament Structure

The structural properties of the tournament are now sufficiently described that we can understand *why* $\mathcal{L} = 2, \mathcal{W} = 6$ is the unique starting point of any $(n, \mathcal{L}, \mathcal{W})$ tournament with $\mathcal{L} \geq 2$ and $n > 2$. The case $\mathcal{L} \geq 3$ with $n > 2$ was already proved impossible in [1].

From the canonical semifinal we can trace the tournament back to its starting state. The entrants to the semifinal and its preceding lower bracket round LR5 are shown on the right in Figure 3. On the left is the recursive relationship between cohort sizes in the bracket rounds. Each upper bracket round receives half the cohort of the preceding one; each lower bracket round receives half the cohort from both bracket rounds in the preceding round.

The resulting backtrack chain is shown in Figure 4. Since there are no teams in the $\left[\begin{smallmatrix} 7 \\ 1 \end{smallmatrix}\right]$ lower bracket round, the tournament starts with every team in the $\left[\begin{smallmatrix} 6 \\ 2 \end{smallmatrix}\right]$ upper bracket round.

With the cohort sizes $|\mathbf{UR}_i|$ and $|\mathbf{LR}_i|$ we can track the indicator $q_i := |\mathbf{LR}_i|/|\mathbf{UR}_i|$. From the recursive relationships we obtain that $q_{i-1} = q_i - 1$. The indicator ends up at $q = 4$ in the column with $\left[\begin{smallmatrix} 2 \\ 2 \end{smallmatrix}\right]$ and $\left[\begin{smallmatrix} 3 \\ 1 \end{smallmatrix}\right]$ in Figure 4, and decreases

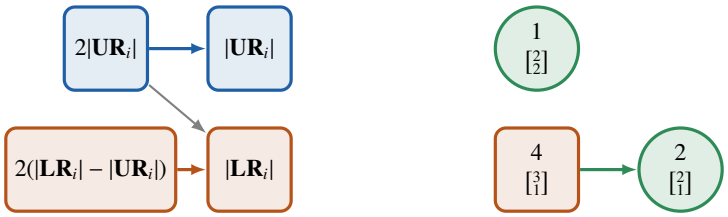


Figure 3: The recursive relationship between bracket round cohort sizes, and its base case, the entrants of the semifinal and the preceding lower bracket round LR5.

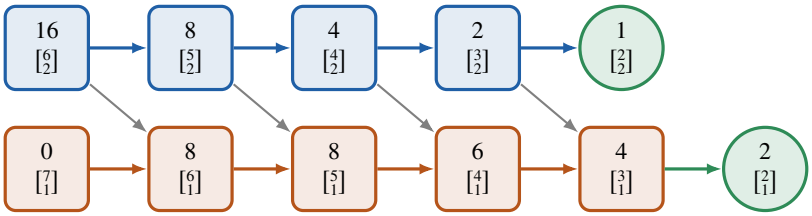


Figure 4: Tracing the cohort sizes back from the canonical semifinal to the start of the tournament.

by one for each step to the left, until it reaches $q = 0$ and therefore an empty lower bracket. This confirms (16, 2, 6) as not merely a numerical coincidence, but the inevitable unique starting point for the tournament.

5. Repeat Avoidance and Adaptive Scheduling

5.1. What Can and Cannot be Avoided

Suppose the tournament contains two teams that defeat everyone else. These teams must eventually meet. Their meeting must be in the upper bracket, as they both remain undefeated until then. When they meet, one of the teams drops into the lower bracket and the other stays in the upper bracket. As they both defeat every other opponent in those brackets, they will both emerge as bracket winners and enter the semifinal. There they will meet again, since the upper bracket winner is in both semifinals. If the result of that semifinal is the reverse of their first match, then they will both advance to the final where they will meet again for a third time.

By defeating all other opponents in their paths, they are essentially playing a best-of-3 playoff series among themselves. Thus, no matter how matches are scheduled, results can conspire to produce repeat matches in the semifinal and final. The possibility of repeat matches is an inherent property of double elimination.

Nevertheless, it would be attractive to minimize and delay repeat matches as much as possible. Can we avoid repeat matches before the semifinal? It turns out that it is indeed possible, but not trivial, and not by any fixed schedule.

5.2. The Limits of Fixed Scheduling

In order to prove that no fixed schedule can avoid repeat matches before the semifinal, first we define what a fixed schedule is.

Definition 5.1. Adaptive versus Fixed Schedules: A schedule is *adaptive* if the pairing of teams in some round is allowed to depend on the results of earlier rounds. A *fixed* schedule is non-adaptive: For each match in the tournament it is known ahead of time which two previous matches feed into it.

To show that fixed schedules can always produce repeat matches before the semifinal, introduce the following colouring mechanism that can be applied to any schedule that is fixed from R1 to R4.

Definition 5.2. LR4 Colouring: LR4 comprises three matches. Assign different colours to the LR4 matches. Propagate the three colours back to the matches in R1 through R3:

- To a lower bracket match, assign the colour at which its winning team would end up if it reached LR4.
- To an upper bracket match, assign the colour at which its losing team would end up if it reached LR4.

Lemma 5.3. *In any fixed schedule for R1 through R4, causally linked matches in R1, UR2, UR3 must have different LR4 colours in order to prevent a repeat match happening in or before LR4.*

Proof. Consider causally linked upper bracket matches M_i and M_j during the first three rounds of the tournament, as shown in Figure 5. Teams X and Y play each other in M_i. Nothing prevents the following scenario from playing out: X beats Y in M_i, keeps on winning until it arrives at M_j, then loses. In that scenario Y drops into a lower bracket match with M_i's colour, followed by X dropping into a lower bracket match with M_j's colour. If those colours are the same, then X and Y are both heading for the same match in LR4. Therefore they cannot both survive LR4. If they keep beating everybody else in the lower bracket matches, the elimination of one of the two teams means they must have met again in LR4 or earlier. □

With that lemma we can now see that it is impossible for any fixed R1-R4 schedule to avoid repeat matches before the semifinal.

Theorem 5.4. *In any fixed schedule for R1 through R4, there exists a scenario where a repeat match happens in or before R5.*

Proof. Consider a UR4 match M_4, in which team U beats team D. Identify the first three matches (M_1, M_2, M_3) that D played, in which it defeated opponents (A, B, C). The progression is shown in Figure 6. If the three matches (M_1, M_2, M_3) do not have three different LR4 colours, then by Lemma 5.3 there can be a repeat match in or before LR4. If they do have three different colours, then teams (A, B, C) cannot meet in LR4 and could therefore all reach LR5, where they are joined by team D. Since LR5 only contains four teams, and D in that scenario has already played each of the other three, team D's LR5 match must repeat a match it played previously. □

5.3. Adaptive Scheduling Can Avoid Repeats

This section will show that adaptivity suffices in LR4 and LR5 without imposing any particular scheduling rule on the preceding rounds. In fact, every state reachable at the start of R4 admits a continuation that avoids repeats in both lower bracket rounds. Consequently, any fixed schedule that guarantees repeat-free play through R3 can be extended adaptively to avoid repeats before the semifinals.

Definition 5.5. The **meeting graph** of a set of teams has those teams as vertices, with an edge between two teams if they have previously played each other. The **upper bracket meeting graph** is the same graph, containing only the edges corresponding to upper bracket matches.

Lemma 5.6. *The upper bracket meeting graph is a forest.*

Proof. Assume there is a cycle in the upper bracket meeting graph. Each node in a cycle is incident to at least two edges of the cycle. The first match played in the cycle removes one of its two nodes from the upper bracket, preventing the removed node from being incident with another cycle edge. □

The full forest is shown in Figure 7. This graph contains all sixteen teams and all upper bracket matches played in the entire tournament. It is unique up to relabelling of nodes. The graph is a single tree, namely the binomial tree B_4.

The upper bracket forest at the start of R4 is shown in Figure 8. It is only missing the one purple edge, corresponding to the UR4 match that has not been played yet at that stage. The graph decomposes into two identical components, partitioning the teams into two groups of eight. Teams U_1 and U_2 are in UR4. Teams C_1 and C_2 have lost in UR3 and are in LR4. There are four more teams in LR4, namely the four winners of LR3. They are among the twelve unlabelled nodes. It is not known ahead of time which four unlabelled nodes survive into LR4.

Lemma 5.7. *At any stage in the tournament before the semifinals, the meeting graph of matches already played between the surviving teams is a subgraph of the upper bracket forest, and is therefore itself a forest.*

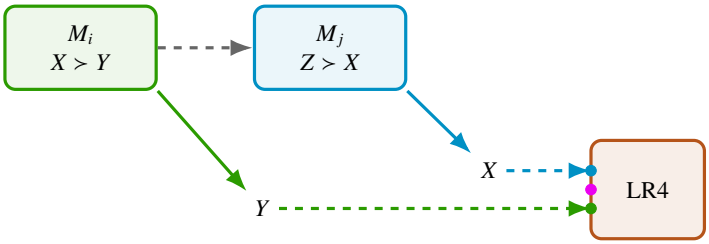


Figure 5: Causally linked upper bracket matches M_i and M_j must have different colours in order to prevent the possibility of a repeat meeting between X and Y in or before LR4.

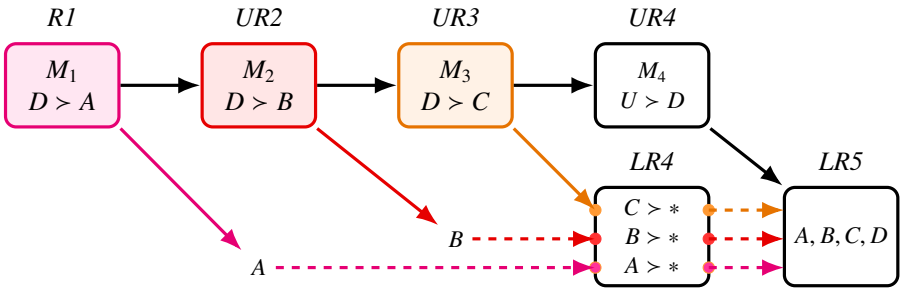


Figure 6: Causally linked path of four matches in R1-UR4. Colours shown are the LR4 colours of the matches. If the three colours are distinct then (A, B, C, D) can all reach LR5, where D has already played against all three other teams.

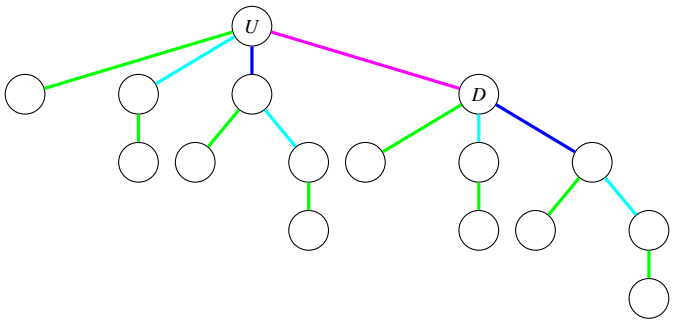


Figure 7: The upper bracket meeting graph after R4, when the upper bracket has been completed. The colour of the edges indicates the round in which the match was played: green for R1, cyan for UR2, blue for UR3, purple for UR4. The round can also be identified by the longest path from the edge downward to a leaf node. The team at the top of an edge is the winner of that match. Only five teams are still alive after R4, including team U as the lone remaining undefeated team, and team D who lost in UR4 and enters LR5. The other three teams in LR5 are among the unlabelled nodes.

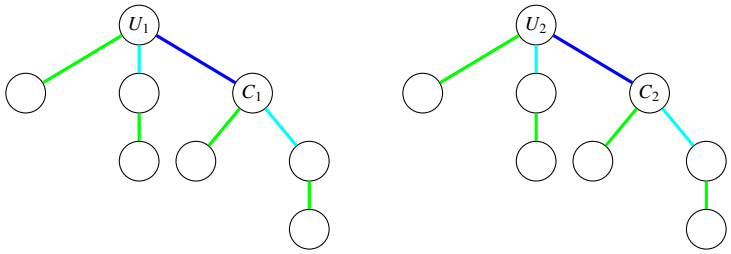


Figure 8: The upper bracket meeting graph at the start of R4. It is obtained from the Figure 7 graph by removing only the purple UR4 edge, whose match has not been played yet, leaving two identical graph components. Eight teams are still alive entering R4. The two U-teams, as winners of the UR3 matches, will play each other in UR4. The winner of that match will become the U-team in Figure 7, the other one will become the D-team. LR4 will feature the two C-teams, who lost in UR3, along with four teams who won in LR3. The latter four are among the unlabelled nodes.

Proof. Every lower bracket match eliminates its losing team. Therefore, if both participants of a previously played match are still alive, that match must have been an upper bracket match, which is contained in the upper bracket forest. □

LR5 will feature four teams. One of them is the D-team, namely the team that lost the U_1–U_2 match in UR4. The other three will come from LR4. The schedule will need to be able to find a matching between these four LR5 teams.

Lemma 5.8. *A forest on $k \geq 3$ nodes with maximum degree 2 contains $\lfloor \frac{k}{2} \rfloor$ vertex-disjoint pairs, each consisting of two nonadjacent nodes.*

Proof. A forest with maximum degree 2 is a disjoint union of paths. Its vertices can therefore be ordered $v_1, \ldots, v_k$ with edges only between consecutive vertices. Put $h = \lceil \frac{k}{2} \rceil$, then $h \geq 2$. Each pair (v_i, v_{i+h}) is nonadjacent. There are $\lfloor \frac{k}{2} \rfloor$ such pairs: if k is even then $h = \frac{k}{2}$ and the pairs are $(v_1, v_{1+k/2}), \ldots, (v_{k/2}, v_k)$; if k is odd then $h = \frac{k+1}{2}$ and the pairs are $(v_1, v_{1+(k+1)/2}), \ldots, (v_{(k-1)/2}, v_k)$. □

Lemma 5.9. *The complement of a forest on four vertices has a perfect matching if and only if the forest contains no $K_{1,3}$ star.*

Proof. The centre of a $K_{1,3}$ star is connected to all three other nodes, so there is no available pairing for it. Conversely, if a forest contains no $K_{1,3}$ star, its maximum degree is at most two. Apply Lemma 5.8 with $k = 4$ to obtain two pairs of nonadjacent nodes. □

By Lemma 5.7, the LR5 meeting graph is a forest. Thus the only obstacle to a repeat-free pairing in LR5 is a four-node $K_{1,3}$ star. The scheduling of LR4, aside from needing to be repeat-free itself, will therefore also need to avoid a $K_{1,3}$ star reaching LR5.

Any team present in the meeting graph of LR5 was also present in the meeting graph of R4. Any edge in the LR5 meeting graph was also already present in the meeting graph of R4: the only new edge in the entire upper bracket forest is the purple edge, but one of the two teams on that edge is still in the upper bracket in R5. So the LR5 meeting graph is an induced subgraph of the R4 meeting graph. Note that it is not an induced subgraph of the LR4 meeting graph, as the LR5 meeting graph will contain team D, which was not in LR4.

Any $K_{1,3}$ star present in the *LR*5 meeting graph must therefore have been present in the *R*4 meeting graph.

Lemma 5.10. *An R4 star can be eliminated in LR4 by pairing up two of its lower bracket teams.*

Proof. Refer to Figure 9, showing one component of the R4 upper bracket meeting graph. The only nodes with degree 3 are U_i and C_i. Both of them are alive in R4. The U-star is present iff both $*$-teams are alive. Those two teams are nonadjacent and can therefore be paired, ensuring that one of them will be eliminated and the star will not be present in R5. Similarly, the C-star is present iff both ‡-teams are alive. They can be paired, eliminating the star. □

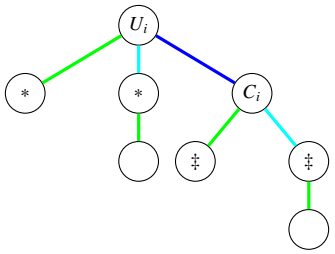


Figure 9: Star blockers in one component of the R4 upper bracket forest. If both $*$-teams are still alive, pairing them blocks the U-star. If both ‡-teams are still alive, pairing them blocks the C-star.

Theorem 5.11. *For any state reached after R3 there exists a repeat-free pairing for LR4 such that, regardless of the results of LR4, there will be a repeat-free pairing available for LR5.*

Proof. There are six teams in LR4. Employ the star-blocking moves of Lemma 5.10. Each star-blocking move is repeat-free and uses two LR4 nodes. Distinguish two cases based on how many LR4 nodes remain after blocking the stars.

- **Case 1: At least four LR4 nodes remain.** The LR4 meeting graph is an induced subgraph of the R4 upper bracket forest. Since the two U-nodes are not present in LR4, the LR4 meeting graph has maximum degree 2. Lemma 5.8 ensures that the remaining LR4 nodes can be paired without repeats.
- **Case 2: At most two LR4 nodes remain.** This means at least four nodes were used to block stars, so there were at least two stars.
  - **Case 2a: There were two stars in one group.** The two stars collectively contain five LR4 nodes: a C-node, two $*$-nodes, and two ‡-nodes. The sixth LR4 node is the C-node from the other group. Blocking the stars used up the $*$-nodes and the ‡-nodes. The two C-nodes remain, one from each group.
  - **Case 2b: There was one star in each group.** Each star already contains three LR4 teams: A U-star contains a C and two $*$-nodes, a C-star contains C itself and two ‡-nodes. Two stars therefore collectively contain all six LR4 nodes. Each star used two LR4 nodes for its blocker. One node remains from each group.

  In both cases 2a and 2b, there were two nodes that were not used to block stars, and those nodes are from different groups. They are therefore nonadjacent and can be paired up.

In every case, all LR4 matches are repeat-free, and all stars are blocked. No stars survive into the LR5 meeting graph. By Lemma 5.9, LR5 admits a repeat-free pairing. □

The algorithm for scheduling LR4 and LR5 can be summarized as follows.

Algorithm 1. Adaptive scheduling of LR4 and LR5:

- *LR4*
  1. *Consider the four teams that played in UR3. For any UR3 team whose R1 and UR2 opponents are still alive, pair up those two former opponents.*
  2. *Arrange the remaining LR4 teams in a sequence* $(v_1, \ldots, v_k)$ *such that any former meetings were between two consecutive nodes in the sequence. Pair teams that are* $k/2$ *positions apart.*
- *LR5*
  1. *Arrange the four LR5 teams in a sequence* $(v_1, \ldots, v_4)$ *such that any former meetings were between two consecutive nodes in the sequence. Pair* (v_1, v_3) *and* (v_2, v_4)*.*

Application of this algorithm will find that LR4 step 1 will encounter at most two such teams and that the pairing of their former opponents is repeat-free, and that LR4 step 2 and LR5 step 1 will produce repeat-free pairings as well.

Algorithm 1 is a sufficient construction, but it is not exhaustive. Other LR4 pairings can also block the stars: A U-star can also be blocked by pairing up its C-node and one of its $*$-nodes. A U-node corresponds to a UR3 team that reached UR4, and its C-node is its UR3 opponent, which is alive in LR4.

5.4. Where Adaptivity is Actually Required

Theorem 5.11 provides a repeat-free continuation of any history of the first three rounds. It did not even require that the history be itself free of repeats. In order to avoid repeats before the semifinal, we therefore merely need to avoid repeats in the first three rounds. This is not difficult to accomplish.

Lemma 5.12. *Repeat matches can be avoided in the first three rounds, even with a fixed schedule.*

Proof. Repeat matches are impossible in the upper bracket, since a previous opponent will no longer be in the upper bracket. Repeat matches are impossible in LR2 as well, since both participants of an LR2 match have played only one match before, which they both lost, so they did not meet each other. Only LR3 requires attention. Its eight participants comprise four teams that won in LR2 and four teams that lost in UR2. One simple way to avoid repeats is to pair the LR2 survivors among themselves, and the UR2 drop-downs among themselves. Every LR3 match then features two teams that lost in the same round, and therefore cannot have met. □

We can then state the general theorem:

Theorem 5.13. *In the (16, 2, 6) tournament it is possible to avoid repeat matches before the semifinal with a fixed schedule in the first three rounds, and adaptive scheduling thereafter. It is not possible with a fixed schedule for the first four rounds.*

Proof. Lemma 5.12 proved that it is possible to have a fixed schedule for the first three rounds that avoids repeats. Theorem 5.11 provides the method for repeat-free continuation in R4 and R5. Theorem 5.4 proved that no fixed scheduling of the first four rounds can ensure repeat-free play before the semifinals. □

It is not always possible to adaptively schedule R4 and R5 at the same time. An example of a case where R4 results need to be known before R5 can be scheduled is shown in Figure 10. At the start of R4, the eight labelled teams are still alive. In order to avoid a repeat in LR4, C_1 and X cannot be paired up, as they have already met. In order to avoid a possible repeat in LR5, the C_2 star must be blocked. The only available pairing in LR4 to do so is Y–Z, since U_2 is not in LR4. This in turn means that B_2 and C_2 cannot be paired up, as that would leave C_1–X as the remaining match. With the pairings C_1–X and B_2–C_2 ruled out, it is possible that both C_1 and X survive into LR5, and it is possible that both B_2 and C_2 survive.

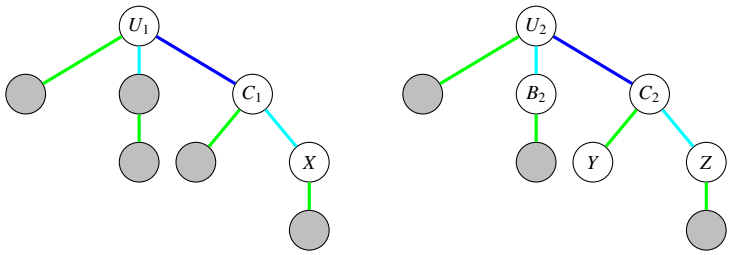


Figure 10: An R4 meeting graph where the eight surviving teams are marked. Grey nodes represent eliminated teams. LR4 can be scheduled, but LR5 cannot yet be scheduled. Even if it were revealed that U_1 wins the UR4 match, it would still be impossible to safely schedule LR5 until the LR4 results are known.

Now suppose more information becomes known: the UR4 match is played and results in a win for U_1. Even with UR4 completed, it is still not yet possible to commit to a schedule for LR5. Team U_2 will be in LR5, where it would need to be paired with the winner of Y–Z, as the other two LR4 matches could produce B_2 and C_2, whom U_2 has already met. That leaves the winners of the other two LR4 matches as the other LR5 pair. Those winners could be C_1 and X, who have already met. Therefore, the LR4 results still need to be awaited before LR5 can be scheduled.

5.5. The Minimax Number of Repeat Matches

As seen in Section 5.3, repeat matches can be avoided before the semifinals. Section 5.1 asserted that repeat matches cannot be avoided in the semifinals and the possible final. This means that adaptive scheduling can limit the number of repeat matches to at most three, but match outcomes can conspire to force at least two repeats. Can clever scheduling limit it to two, or could outcomes conspire to force three repeats?

This amounts to a game played between two entities, the **Scheduler** and the **Collider**. The Scheduler arranges the matches, the Collider picks the outcome of each match. The Scheduler's goal is to minimize the number of repeat matches. The Collider's goal is to maximize it. With optimal play, we know the Scheduler can force the outcome to be at most 3, and the Collider can force it to be at least 2.

Some care needs to be taken specifying the rules of the game. The Collider simply picks match results, and is free to pick any of the two possible outcomes for a match. The Scheduler needs to ensure the (16, 2, 6) structure is maintained. From the results in Section 4, the only admissible matches are:

- Teams with the same win-loss record may play each other, except that two teams with a 4-1 record may meet only when no undefeated team remains.
- If there is exactly one undefeated team left, it may play against a team with a 4-1 record.

These rules cover the upper and lower bracket matches, as well as the semifinals and the possible final.

The rules do not specify in what order the Scheduler needs to arrange the matches beyond respecting the causal links between them, nor when exactly the Collider responds. Some variants of playing the game:

1. **Fixed:** Schedule all matches ahead of the tournament, then pick all the results. Section 5.2 already showed that in this variant the Scheduler cannot even avoid repeat matches before the semifinal.
2. **Round by round:** Schedule one round of play, including all upper and lower bracket matches; then pick all the results of that round. Here and below we consider the two semifinals to be one round, and the possible final to be one round as well.
3. **Round by round, match by match:** Schedule one single match, then immediately pick the result. Continue until all matches of the upper and lower bracket in that round are completed, then move on to the next round.
4. **Unrestricted:** Schedule one single match *in any round*, then immediately pick the result. Rounds do not need to be completed before next-round matches start. The Scheduler may, for instance, opt to play out the entire upper bracket before scheduling any lower bracket matches.

These variants give progressively more freedom to the Scheduler. Each variant permits the Scheduler to follow any strategy available in the previous variant, while providing additional options. Consequently, the minimax number of repeat matches cannot increase from one variant to the next.

The unrestricted game can be solved by retrograde minimax analysis. A state S consists of the win-loss records of all surviving teams, together with the graph of their previous encounters. This graph contains an edge between two surviving teams if they have played each other before. Matches involving eliminated teams need not be retained, since they can no longer cause repeats. In any non-terminal state, the Scheduler is to move next, picking a pair of teams. The Collider then picks a result of the match, leading into the next state.

Let $\mathcal{P}(S)$ denote the set of admissible pairings in state S. For a pairing $p = \{p_\blacktriangleleft, p_\blacktriangleright\} \in \mathcal{P}(S)$, the immediate cost $c(S, p)$ equals 1 if $p_\blacktriangleleft$ and $p_\blacktriangleright$ met previously, and 0 otherwise. Let $S_{p,\blacktriangleleft}$ and $S_{p,\blacktriangleright}$ be the successor states when $p_\blacktriangleleft$ and $p_\blacktriangleright$ win, respectively.

The game-theoretical value $\mathrm{GTV}(S)$ of state S is defined as the minimum worst-case number of *additional* repeat matches that the Scheduler can achieve, regardless of the Collider's choices. The Bellman recurrence is

$$\mathrm{GTV}(S) = \begin{cases} 0, & \text{if } S \text{ is terminal,} \\ \min\limits_{p=\{p_\blacktriangleleft, p_\blacktriangleright\} \in \mathcal{P}(S)} \Big[c(S,p) + \max(\mathrm{GTV}(S_{p,\blacktriangleleft}), \mathrm{GTV}(S_{p,\blacktriangleright})) \Big], & \text{otherwise.} \end{cases}$$

The outer minimum represents the Scheduler's choice of pairing, while the inner maximum represents the Collider's choice of outcome. Each match consumes one remaining life, so the state graph is acyclic and the recurrence can be evaluated backward from the terminal states.

The game is invariant under relabelling of teams. Two states are therefore identified if there exists a bijection between their surviving teams that preserves both their win-loss records and their previous-encounter graphs.

For the unrestricted variant, an implementation of the state-space generator, retrograde solver, and separate verification is provided in the accompanying repository [2]. The implementation uses canonical forms of vertex-coloured graphs to identify such isomorphic states. Furthermore, admissible pairings are partitioned into orbits under the automorphism group of each state, so only one representative of each orbit needs to be considered. The database containing the state space is available in the accompanying artifact [3].

Theorem 5.14. *In variants 2, 3, and 4, the minimax value of the Scheduler–Collider game is 3.*

Proof. Exhaustive search of the full state space of variant 4 returns a minimax value of 3. The minimax values for variants 2 and 3 must be at least equal to the value for variant 4. In all three variants the Scheduler can avoid repeat matches before the semifinal, so the minimax value is at most 3. □

Figure 11 shows the number of distinct states, as well as the number of distinct *interim* states, at each depth. The depth reflects the number of matches played. An interim state arises after the Scheduler has scheduled a match, when it is the Collider's turn to move. Table 1 lists the total counts of the number of distinct states, modulo the isomorphism symmetries.

The figure and table also include the number of *equilibrium states*, which are states that can arise when both the Scheduler and the Collider play optimally. Also shown are *critical states*. These are interim equilibrium states where the two possible match outcomes lead to states with different game-theoretical values. Critical states are where the Collider is at risk of making a mistake.

The game-theoretical value shows that the Collider can always force three repeat matches, even when the Scheduler is given maximum freedom. However, where the Collider has a trivial strategy to force two repeat matches, the strategy to force three repeat matches appears to be significantly more complex. Figure 12 shows the percentage of interim equilibrium states that are critical at each depth. The game appears to be fairly easy for the Collider in the early phases. Indeed there are no critical interim states before depth 9. The complexity then rises sharply, peaking at match 20 and then rising again towards the end of the game, when the size of the state space collapses.

The fraction of critical states provides an informal indication of how frequently the Collider faces a consequential decision, rather than a formal measure of strategic complexity. For instance, the peak at depth 29 involves states featuring at most three remaining teams. At that point, the decision which team to keep alive in order to force another repeat in game 30 is trivial.

The unrestricted variant would not be very desirable in an actual sports tournament. It benefits neither participants nor spectators if, for instance, all the lower bracket teams must wait for the entire upper bracket to finish, and subsequently the upper bracket winner must wait for the entire lower bracket to be played out. Since the Scheduler does not benefit from the unrestricted rules anyway, the remainder of the paper adopts the conventions:

- Round-by-round scheduling is used.
- Repeat matches need to be avoided before the semifinals. The number of repeat matches in the semifinals is of no concern.

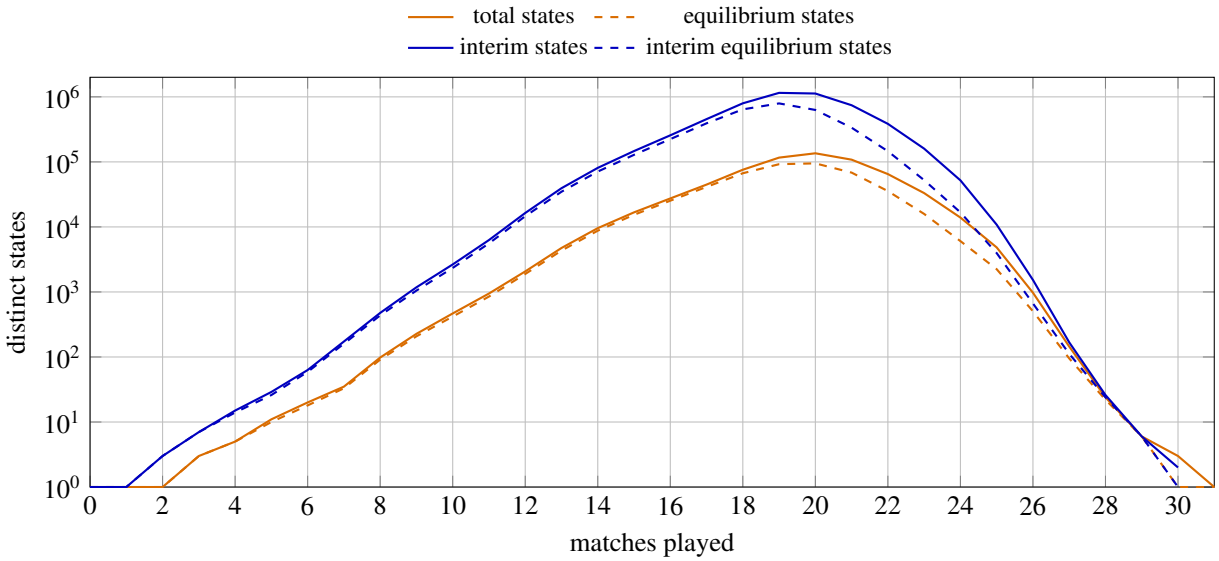


Figure 11: Distinct state counts (orange, bottom two curves) and interim state counts (blue, top two curves).

	states		**interim states**	
total	661,114		5,431,178	
equilibrium	481,272	72.8%	3,494,664	64.3%
critical			442,706	12.7%

Table 1: Number of distinct states in the Scheduler–Collider game. Each percentage is relative to the preceding row.

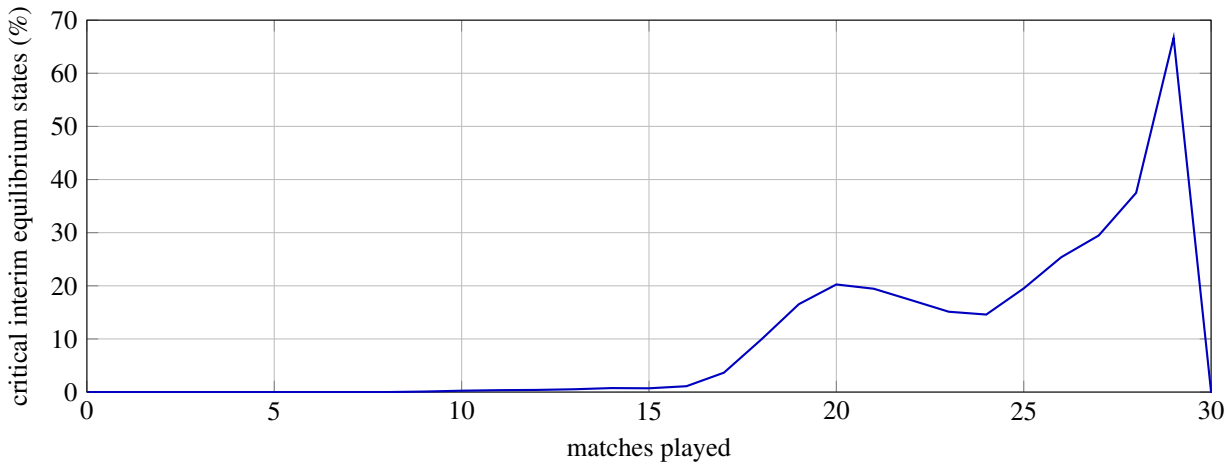


Figure 12: Percentage of interim equilibrium states that are critical.

6. Seeded Draws

Having resolved to avoid repeats before the semifinals, a more important issue than the exact number of repeat matches in the semifinals now presents itself. How can the tournament enhance the likelihood that the best teams survive to the semifinals? After all, one major motivation for the multiple-elimination format was to limit the risk of a strong team being eliminated early due to either an upset result or an early encounter between two strong teams.

Repeat avoidance during bracket rounds still leaves considerable freedom in the choice of pairings. That freedom can be used to preserve the preference that the strongest participants feature in the final stages of the tournament. In practice, elimination tournaments commonly employ **seeding** to reduce the probability that the strongest competitors eliminate one another in the early rounds. Relatively simple seeding rules can have a substantial effect on whether the strongest competitors survive to the later rounds, sometimes rivaling the effect of more elaborate changes to the tournament structure [4].

For the (16, 2, 6) tournament, assume the teams are appropriately seeded 1 through 16. As per common convention, higher seeds have lower numbers: Seed 1 is the highest seeded team.

6.1. Inherited Seeds

The canonical 16-team seeded schedule for the upper bracket is shown in Figure 13. The initial round is scheduled with seed i meeting seed $17 - i$. The rationale is that a higher-seeded team should not be penalized by being given a more difficult route than a lower seed. This is intended to prevent the perverse incentive for any team to artificially lower its seeding in scenarios where seeding depends on prior results.[1]

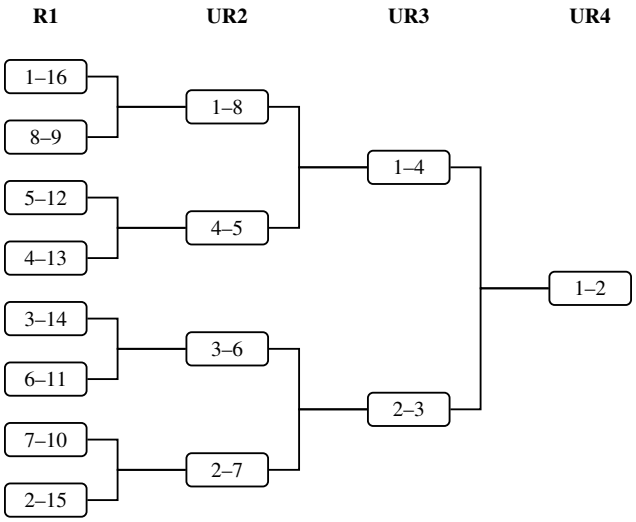


Figure 13: The standard 16-team seeding in the upper bracket, if no upset results occur.

The progression of the seeds through the upper bracket rounds in Figure 13 shows the scenario where the higher seed beats the lower seed in every match. An **upset** occurs when the lower seed wins instead. If an upset does happen, there are two possible scheduling responses:

- **Re-seeding:** At the start of each round, teams are re-seeded. The highest remaining seed becomes seed 1, and so on. This dooms the lowest seeded team to play the highest remaining seed in every round it reaches.
- **Inherited seeds:** A team that scored an upset win inherits the seed of the team it defeated. The victory is "rewarded" with the protected status of the higher seed.

[1]This protection holds only in expectation: Match results that defy the seeding's expectations can place a stronger opponent in seed 1's path than in seed 2's, and Schwenk [5] showed that this can lead to seed 2 having a better overall chance of winning the tournament.

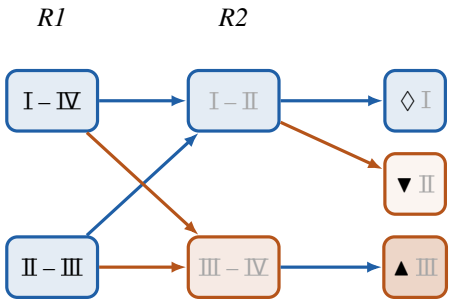
group	Ⅰ	Ⅱ	Ⅲ	Ⅳ
G_A	1	8	9	16
G_B	2	7	10	15
G_C	3	6	11	14
G_D	4	5	12	13

Table 2: Allocating the sixteen seeded teams into four groups.

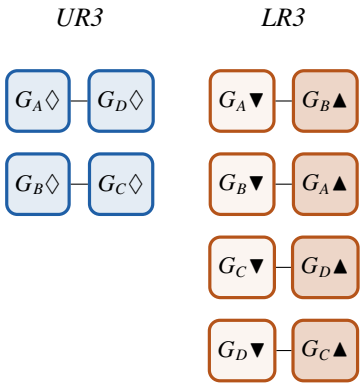
When looking for a fixed tournament schedule for the first three rounds, re-seeding is not an option as it is explicitly adaptive: The pairings depend on the results of the previous round. For the (16, 2, 6) tournament, assume therefore that inherited seeds are used. Consequently, each slot in the tournament schedule can be identified by the original seed that was expected to be there, had there been no upsets. This extends to lower bracket slots as well.

6.2. A Fixed Opening

An effective example of a fixed schedule for the first three rounds is illustrated in Figure 14. Teams are divided into four groups, $\{G_A, G_B, G_C, G_D\}$. Within each group, the R1 match pairs the highest seed with the lowest seed, and pairs the two middle seeds. The UR2 match then pairs the winners within each group, and the LR2 match pairs the two teams that lost their first match. The team that won both its matches advances to UR3, while the two teams with one win and one loss advance to LR3. The remaining team lost both its group matches and is eliminated.



(a) A group of four participants, with Roman numerals indicating the ordering of seeds within the group. In R1, the highest seed plays the lowest seed, and the two middle seeds play each other. The winners of the two R1 matches are paired up in UR2, the other two are paired up in LR2. Greyed-out numerals indicate inherited seeds.



(b) Scheduling of R3 following the two group matches in four groups $\{G_A, G_B, G_C, G_D\}$.

Figure 14: Group-based scheduling of the first three rounds. From each group emerge one unbeaten team marked ◊, and two teams with one win and one loss each, marked ▼ and ▲.

This schedule does not repeat any matches, regardless of the outcomes. An R2 match cannot be a repeat since its two participants either both won their R1 match, or they both lost, so they cannot have played each other. An R3 match cannot be a repeat since it pairs up two teams from different groups, where all R1 and R2 matches were within one group.

Standard seeding assigns seeds to groups according to Table 2. With this assignment, the R1 matches exactly follow the reverse-order pairing: seed i plays seed $17 - i$ in its opening match.

Tracing the inherited seeds through the first three rounds, as shown in Table 3, shows that R2 and UR3 also follow reverse-order pairing. Only LR3 deviates slightly. Reverse-order would have {5–12, 6–11, 7–10, 8–9}, but these would repeat seed pairings from R1 and could therefore repeat matches.

By the Continuation Theorem 5.11 of Section 5.3, this three-round opening always admits an adaptive repeat-free pairing in LR4, followed by an adaptive repeat-free pairing in LR5.

R1		R2		R3	
G_A	1 – 16	G_A	1 – 8	$G_A\Diamond$–$G_D\Diamond$	1 – 4
G_B	2 – 15	G_B	2 – 7	$G_B\Diamond$–$G_C\Diamond$	2 – 3
G_C	3 – 14	G_C	3 – 6		
G_D	4 – 13	G_D	4 – 5		
G_D	5 – 12				
G_C	6 – 11	G_A	9 – 16	$G_D\blacktriangledown$–$G_C\blacktriangle$	5 – 11
G_B	7 – 10	G_B	10 – 15	$G_C\blacktriangledown$–$G_D\blacktriangle$	6 – 12
G_A	8 – 9	G_C	11 – 14	$G_B\blacktriangledown$–$G_A\blacktriangle$	7 – 9
		G_D	12 – 13	$G_A\blacktriangledown$–$G_B\blacktriangle$	8 – 10

Table 3: Tracing the inherited seeds through the first three rounds.

6.3. Choosing Among Feasible Pairings

The adaptive scheduling method of Section 5.3 typically leaves multiple feasible pairings. A pairing is feasible if it avoids all previous encounters and, in LR4, guarantees that a repeat-free pairing in LR5 will remain possible regardless of the outcomes of R4. Theorem 5.11 ensures that such pairings always exist. However, feasible pairings need not be unique.

To express a preference among feasible pairings, the reverse-ordering principle can be employed: If the seeds are listed in sorted order, then their opponents should appear in reverse-sorted order. In particular, the highest seed should play against the lowest seed.

This preference can be quantified by the Kendall-τ rank correlation between the seed of each team and the seed of its opponent. We use the inherited seeds of the positions, not the original seeds of the teams. Preferably, higher seeds have lower seeded opponents, so perfect anti-correlation with $\tau = -1$ is the ideal. Between possible schedules for a round, τ values closest to -1 should be preferred.

This principle was already employed in the fixed opening schedule of Section 6.2 for LR3. The schedule with $\tau = -1$ was not feasible, as it could produce repeat matches: the four reverse-ordering pairs were already used in R1. The proposed schedule achieves $\tau = -\frac{5}{7}$:

$$\text{LR3:} \quad 5\text{–}11, \quad 6\text{–}12, \quad 7\text{–}9, \quad 8\text{–}10.$$

6.4. The No-Upset Progression

Suppose that in every match the better seeded team wins. Original and inherited seeds then coincide along the proposed opening of the schedule in Section 6.2. LR4 in that case contains seeds $\{3, 4, 5, 6, 7, 8\}$. Applying the rules of Algorithm 1 finds that no UR3 team has both its R1 and UR2 opponents still alive, as their R1 opponents were all eliminated in LR2. It suffices therefore to find a repeat-free pairing. The only previously played matches are 3–6 and 4–5, both from UR2. Of the $5!! = 15$ possible schedules for six teams, 10 schedules avoid these repeats. One of them is the perfect anti-correlation schedule:

$$\text{LR4:} \quad 3\text{–}8, \quad 4\text{–}7, \quad 5\text{–}6.$$

If R4 also does not produce any upsets, the LR5 teams are $\{2, 3, 4, 5\}$. The only previous meetings are 2–3 in UR3, and 4–5 in UR2. The $\tau = -1$ schedule is feasible:

$$\text{LR5:} \quad 2\text{–}5, \quad 3\text{–}4.$$

The resulting semifinalists are teams $\{1, 2, 3\}$. In this upset-free scenario, the upper bracket proceeded as in Figure 13, the lower bracket proceeded as in Figure 15.

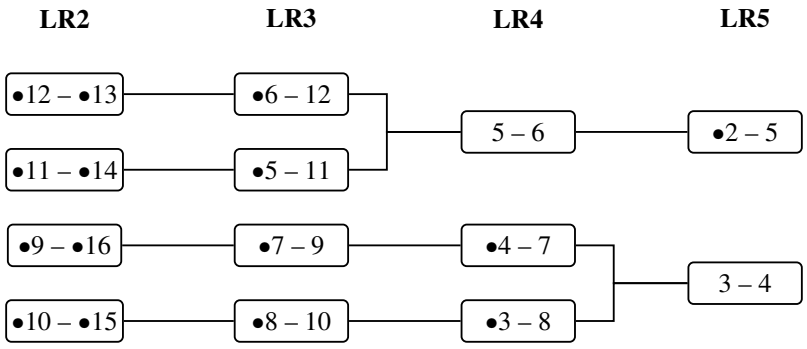


Figure 15: The standard 16-team seeding in the lower bracket, if no upset results occur. Seeds labelled with • indicate teams that dropped into the lower bracket, having lost in the previous upper bracket round. Seeds without • survived the previous lower bracket round.

7. Conclusion

The sporadic (16, 2, 6) fixed-elimination fixed-win tournament has considerably less structural freedom than its definition might suggest. Fayers's validity conditions force its upper and lower bracket round structure and its three-team semifinal. The remaining freedom lies primarily in the choice of pairings within these rounds.

This freedom is sufficient to avoid all repeat encounters during the bracket phase. However, while the first three rounds can be scheduled in advance, the remaining two require adaptive pairings. The constructive continuation algorithm guarantees repeat-free play in these rounds regardless of the preceding match results. An exhaustive minimax analysis further establishes that, even with unrestricted scheduling, match results can conspire to force three repeat matches overall.

Finally, the repeat-avoidance constraints leave sufficient flexibility to incorporate conventional seeding principles. A fixed, seeded opening combined with adaptive scheduling provides a complete tournament schedule that preserves the fixed-win property while avoiding repeats before the semifinals.

Appendix A. Extension: A Relaxed Win-band Requirement

To enable multiple-elimination tournaments with a different number of participants, the strict fixed-win requirement that the champion must have won exactly $\mathcal{W}$ matches can be relaxed into a win-band requirement: the champion must have won between $\mathcal{W}$ and $\mathcal{W} + k$ matches for $k \geq 1$. This can be achieved by inserting extra rounds that are played in only one bracket.

Appendix A.1. Extra Lower Bracket Rounds

From a sporting perspective it could be argued that inserting an extra round in the lower bracket is defensible: The teams in the upper bracket are rewarded for winning their matches by getting to skip one match. The "scenic route" to winning the tournament, through the lower bracket, then takes longer than the undefeated path through the upper bracket.

The recursive relation between the cohort sizes in the upper and lower bracket, as displayed in Figure 3, can be augmented trivially with a recursive step that applies to an extra lower bracket round.

$$|\mathbf{UR}_{i-1}| = |\mathbf{UR}_i|, \qquad |\mathbf{LR}_{i-1}| = 2|\mathbf{LR}_i|.$$

The chain is again guaranteed to lead back to a starting state with all teams in the upper bracket, as can be seen by tracking the q-indicator as introduced in Section 4.4. This indicator, traced backwards in time, decreases by one for a regular round. It *doubles* with an extra lower bracket round. The length of the tournament depends on which q gets doubled, as can be seen in Figures A.16 and A.17.

Inserting an extra lower bracket round into the canonical (16, 2, 6) tournament increases n to any power of 2 between 32 and 256. Inserting two extra lower bracket rounds increases n to any power of 2 between 2^6 and 2^{16}. It

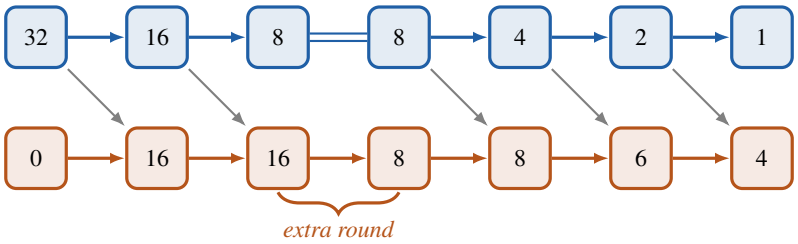


Figure A.16: Inserting an extra lower bracket round early in the tournament. The double line between the upper bracket rounds indicates that the teams progressed without playing that round.

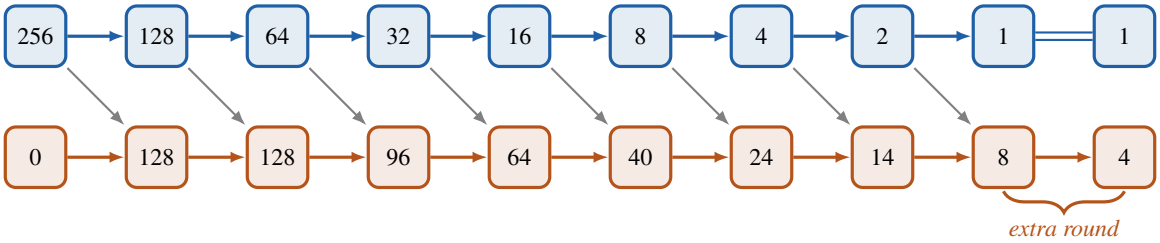


Figure A.17: Inserting an extra lower bracket round late in the tournament increases tournament length.

will necessarily be a power of 2, due to the simple doubling relationship in the upper bracket. With $n = 2^m$ participants and k inserted lower bracket rounds, the champion wins between $m + 2$ and $m + 2 + k$ matches. Each inserted round contributes an additional required victory precisely when the eventual champion has already entered the lower bracket.

Appendix A.2. Extra Upper *Bracket Rounds*

The construction of Appendix A.1 can formally be dualized: insert an extra round in the *upper* bracket.

Tracking the q-indicator again, this insertion sends q to $(q-1)/2$, so it shortens the tournament rather than lengthening it. Absent any extra lower bracket rounds, the only possible nontrivial moment it could be inserted is shown in Figure A.18. We now have $n = 8$, with the champion requiring 4 or 5 wins.

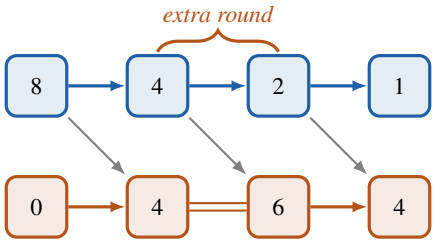


Figure A.18: Inserting an extra *upper* bracket round.

As a sporting format the extra upper bracket round would not be defensible, as it would punish a team for winning.

References

[1] M. Fayers, Multiple-elimination knockout tournaments with the fixed-win property, Discrete Mathematics 290 (2005) 89–97.

[2] J. Van Ryswyck, Booker: State-space construction and verification for fixed-elimination tournaments, 2026. URL: `https://doi.org/10.5281/zenodo.24325904`. doi:`10.5281/zenodo.24325905`.

[3] J. Van Ryswyck, Booker datasets, 2026. URL: https://doi.org/10.5281/zenodo.24329329. doi:10.5281/zenodo.24329329.

[4] P. A. Scarf, M. M. Yusof, A numerical study of tournament structure and seeding policy for the soccer World Cup Finals, Statistica Neerlandica 65 (2011) 43–57. doi:10.1111/j.1467-9574.2010.00471.x.

[5] A. J. Schwenk, What is the correct way to seed a knockout tournament?, The American Mathematical Monthly 107 (2000) 140–150. doi:10.1080/00029890.2000.12005171.