

Introduction

This document presents an empirical analysis concerning the integrity of TETR.IO solo modes (40 LINES and BLITZ). An investigation was sparked on April 14th after TETR.IO player FURRY achieved a [world record](#) in the 40 LINES game mode (colloquially referred to as “Sprint” in the stacker community) with a time of 14.183 seconds.

We hope to outline newly discovered methods which can be used to verify the legitimacy of any given solo mode TETR.IO replay. When applying these methods of verification to FURRY’s solo replays, we have obtained **evidence beyond a reasonable doubt** confirming that many of his top solo mode achievements, dating at least a few months back and including his previously mentioned **14.183 TETR.IO 40 LINES world record**, have been **illegitimately achieved**.

This investigation was led by [mat1jaczyyy](#) and [renge](#), and was reviewed by TETR.IO staff.

NOTE: We request that you remain respectful. Please do not harass any of the individuals mentioned in this post. Their case is between them and staff only; this document is shared for transparency reasons.

Objectivity

It should be noted that the authors of this document seek to help preserve the integrity of the TETR.IO leaderboards (and by extension, of the overall stacker community), and are solely motivated by the presence of apparent exceptional empirical data. TETR.IO replays are completely deterministic and contain exact inputs recorded from the live play: the exact same inputs will always yield the exact same results.

Although this investigation was initially sparked by a single player appearing suspicious on the surface, the main focus of the investigation is set on developing more exhaustive methods of analyzing TETR.IO solo replays that could be used to further assure whether a replay is legitimate or manipulated. These methods will be based on the nature of the inputs stored inside said replay as well as direct comparisons of those inputs to known legitimate ones, such as TETRA LEAGUE replay data.

The above-stated methods will be used to verify the legitimacy of FURRY’s top 40 LINES and BLITZ replays. The goal of this document is to present an unbiased analysis of all mentioned data, and we realize that we cannot analyze all replays hosted on TETR.IO, nor can we guarantee that our findings present the entire situation.

Motivation

In this section we will present numerous uncharacteristic details about FURRY's replays which had initially prompted us to start a more detailed investigation which has resulted in this document.

Odd 40 LINES replays

Here are some replays containing odd moments that are not alike most other top solo PBs:

[17.377 \(9547e8PcN.ttr\)](#)

- Very poor stacking in the opening
- Piece #34 was placed with 4 unnecessary inputs including two unnecessary double-taps
- Piece #47 puts the stack in a very unfavorable position that is recovered from with multiple softdrops
- A 17.xxx personal best using 6-3 stacking with 3.4KPP and including all of the above mistakes is just absurd, especially when the previous 18.xxx PB was 3.1KPP

[16.628 \(gDQ2zoUMJ.ttr\)](#)

- Attempts to perform PCO loops at a pace of 6 PPS
- Accidentally achieves 5 PC WR time with significantly suboptimal 55 pieces, missing one PC and getting one 2 line PC

[16.268 \(y4JkGpS2XK.ttr\)](#)

- 16.xxx PB with extremely high KPS (19.2)
- No hesitation shown after an early misdrop on piece #14 which is usually a reset
- Piece #37 sets up a TSD which was later misdropped on piece #41. Little hesitation shown in fixing said mistake
- Overall poor stacking, with some triple rotates (such as on piece #69)

[15.393 \(c3NXLYjAs.ttr\)](#)

- Extremely high 20.14 KPS run, virtually unseen in the genre

Suspicious BLITZ personal best

[911,403 \(68d1HCLn4.ttr, 911k_blitz_freestyle_sussy.ttr\)](#)

FURRY's personal best (BLITZ freestyle WR) was automatically marked as suspicious upon its completion and submission. A day later, TETR.IO staff uploaded the replay to the official in-game leaderboards.

The screenshot shows the 'RESULTS' screen for a BLITZ game. At the top, the 'FINAL SCORE' is 911,403. Below the score is a 'BACK' button and an 'AGAIN' button. A 'REPLAY' button is also visible. The 'STATS' section is divided into 'OVERVIEW' and 'FULL' tabs. The 'OVERVIEW' tab shows the following statistics:

STATS	VALUE
LEVEL	14
PIECES PLACED	498
PIECES PER SECOND	4.15
KEYS PRESSED	1714
KEYS PER PIECE	3442
KEYS PER SECOND	14283
SCORE	911,403
LINE	198
LINE PER MINUTE	99.00
T-SPINS	34
MAXIMUM COMBO	

At the bottom, it says 'BLITZ PLAYED BY FURRY ON 3/17/2022, 7:00:00 PM'. A warning message at the bottom right states: 'YOUR REPLAY LOOKS SUSPICIOUS OR OUT OF DATE (1 ISSUE). TRY RESTARTING TETR.IO.'




Fig 1. FURRY's BLITZ personal best replay was [marked as suspicious](#) upon submission.

The screenshot shows a Discord chat conversation. The user 'furry' (03/18/2022) asks '@osk !! Yesterday I got a 911k blitz which had an insane RNG but the client said the replay was suspicious so it was not submitted. I was wondering if you can verify it? 🙄' and provides a link to the replay. The user '@osk' (03/18/2022) responds with 'congrats on the huge pb! uploaded as R:68d1HCLn4'. A file named '911k_blitz_freestyle_sussy.ttr' (345.66 KB) is attached to the message.

Fig 2. The replay had to be [manually uploaded](#) by TETR.IO staff.

Fast improvement rate

Because different people improve at different rates, the period of time a player takes to reach a certain milestone is not inherently cause for suspicion. However, as players' efficiency becomes more optimized and they start reaching the limits of their physical speed, their achievements should start to slowly diminish and further personal bests should take more time to achieve.

In the following chart, we can see that this sort of exponential diminishing improvement is the case for many top sprint players. FURRY is the exception here, they display linear improvement instead, especially once they enter the sub-20 range:




Fig 3. Comparison of top sprint players' improvement rates across TETR.IO and Jstris

We observe FURRY's improvement graph curves a lot flatter than other top sprint players, having more of a linear trajectory instead. This curving amount is displayed with the dashed line representing a best fit based on the exponential function with an additional y-axis shift:

$$ae^{bx} + c$$

This indicates unnaturally fast improvement, especially past the sub-20 barrier.

High physical speed

Physical speed refers to how quickly a player is able to physically execute inputs while playing the game. Physical speed is usually measured in keys per second (KPS for short). A higher KPS value would indicate that the player is putting more effort and strain on their hands to do extra inputs when compared to a lower KPS value.

Glancing at FURRY's TETR.IO profile, we can easily observe a majority of his top 40 LINES replays have been played out with significantly higher KPS values than any other top sprint player on the 40 LINES leaderboard, deviating greatly from other players at his skill level.

Players with poor physical efficiency will reach their physical KPS limit at a much lower value, usually due to suboptimal key bindings.




Fig 4. Comparison of top sprint players' key bindings

Given that FURRY's bindings are identical to ICLY's and are similar to other top sprint players' bindings, we can assume they don't provide him a direct advantage. HIRYU is the only player shown here with significantly different bindings than other players; however, it seems that they aren't put into a disadvantage, despite the problematic Hard Drop key placement.

The following chart shows top sprint players' box plots of KPS values taken from their best 10 fastest completed sprints across both TETR.IO and Jstris:




Fig 5. Comparison of KPS values from top sprint players' best 10 sprints

Most of FURRY's best 10 sprints land in the high 17 - low 19 KPS range, with the lowest value at 17.2 KPS and the highest at 20.14 KPS. These values are larger than any of their peers, and are largely unseen in the stacker community with one exception: JUSTIN1L8. JUSTIN1L8 is an anomaly on Jstris, as they have demonstrated that it is possible for a real human to play at such high KPS values, with an outlier PB run finishing with a whopping 20.82 KPS. However, due to hand injuries caused by playing at such a high physical speed, JUSTIN1L8 has since moved away completely from regular input to 1-key finesse.

CABOOZLED_PIE is also included in the chart as a player who is considered to have a relatively high and inconsistent KPS value compared to other peers in their skill range, but comparable to other top sprint players. Although their time is not as fast as FURRY's due to high finesse leading to high keys per piece (KPP), they are able to achieve high KPS. While high KPS definitely provides an advantage, this illustrates that it is not the only requirement for acquiring world-class times. Good finesse is also required to make sure as little keypresses are required to finish a sprint run.

To conclude: though other players have been able to legitimately finish runs with a KPS as high as FURRY, any player nearing or passing 20 KPS would be a reason for concern.

Techniques of analysis

Input timestamp difference trends

The first method we developed to analyze a player's replays involves comparing the differences of the timestamps between each consecutive input event stored in the replay, regardless of the key and state the event describes. Timestamps are stored as frames, where one frame equals 16.66ms of real-time. Here's a sample visual representation of these differences:




Fig 6. An example visual of a player's differences between key events

Intuitively, these differences would only accurately represent the amount of time taken between each input event as the player executed it. However, because modern computers work by sampling data at constant intervals, they do not directly and instantaneously propagate a physical key event to TETR.IO. Instead, there are multiple different samplers keyboard events must go through until they reach the game client.

The first sampler that grabs information from real keypresses is the keyboard's scan rate. This is simply how often the microprocessor on the keyboard is able to scan the matrix of the keyboard for changes of state on the physical keys themselves.

Next, this microprocessor needs to transfer this key state change event to the computer itself. For USB keyboards, reporting these changes is executed on a set polling rate as defined by the USB standard. This is either 62.5Hz (every 16ms), 125Hz (every 8ms), 250Hz (every 4ms), 500Hz (every 2ms), or 1000Hz (every 1ms).

Most low-budget and membrane USB keyboards poll at 125Hz. High-end mechanical gaming USB keyboards almost always poll at 1000Hz, although ironically they often tend to have significantly slower scan rates than that, so they don't always end up taking full advantage of this. Sometimes that leads to data representing their scan rate rather than poll rate to show up.




Fig 7. An example replay showing use of a 125Hz USB keyboard

PS/2 keyboards use an interrupt-based approach instead, where the keyboard sends a packet containing changes immediately upon detecting a change. The CPU is interrupted and the OS has to process the key event immediately. Unfortunately, due to the age of the PS/2 standard, the bandwidth on the bus is incredibly low and as such not many key events can be transferred in a single second.

If the events come in close enough, this effectively puts a hard minimum time between any two consecutive events, which can sort of be observed like a polling rate if there are enough events happening close together. In practice, we have observed this gap to usually be either 120Hz or 180Hz. Keyboards built into laptops usually communicate over a PS/2 interface.




Fig 8. An example replay showing use of a PS/2 keyboard, much more inconsistent 180Hz

When an application wants to listen to key events, it defines an event handler function (callback) to run upon receiving a keyboard event from the OS. If the application wants to get an accurate timestamp of when the key event happened, it should immediately query the OS for the current high-resolution timestamp. Because of this, an overworked CPU could cause inputs to register a later timestamp than in reality, as the routine that responds to key events might not run in an adequate time to record the timestamp accurately.

TETR.IO is a web-based game and as such runs inside a browser. At the moment a key event is received, TETR.IO calls `performance.now()` immediately to get a timestamp in milliseconds for that event from the browser. The nature of this could potentially introduce a lot of noise to the timestamp. Depending on the browser implementation, the precision could be reduced by rounding to shield the user against certain attacks. The time it takes from the browser's internal callback receiving the key event to executing relevant JavaScript code could also vary, potentially introducing noise. Another potential factor that introduces noise is the FPS at which the game is rendering at, however we were unable to conclude the exact effects of different FPS values to the player's inputs.

TETR.IO's internal engine runs at 60 full frames per second, with 10 hidden subframes inside a full frame. The timestamp in milliseconds assigned to the key event is mapped from real-time to an in-game subframe. It's at this point that finally the key event is registered in-game and can be processed to an in-game action. Given that the engine runs at 600 time points per second, the millisecond timestamp is rounded and some precision is lost here. One subframe is 1.66ms.




Fig 9. An example flat replay that's hard to tell anything about. Likely 500 or 1000Hz USB

With that context in mind, we can now calculate how often each difference in subframes between two consecutive key events is recorded inside a TETR.IO replay. The idea is that the same player running the same computer configuration and using the same keyboard hardware should always reveal the same trend of differences, assuming there are enough data points in the replay to conclusively determine the constant pattern from the many sources of noise outlined above.

We can use this comparison to question whether a player's solo mode differences match those of public TETRA LEAGUE data, which can be used as a trusted baseline of performance for a player. This is because solo mode attempts are played out locally in an untrusted environment, while online matches require a constant connection to a trusted server to communicate with the opponent. It is implausible to introduce foul play in an online match without desyncing and potentially triggering any anti-cheat mechanisms. Comparisons against trusted TETRA LEAGUE data will be a recurring theme throughout the investigation.

This form of comparison does not necessarily require an easily identifiable keyboard poll rate to verify solo replays as legitimate, as shown by this example:




Fig 10. The patterns in these replays are certainly odd, but they definitely match

If we want to test a replay for the most dominant sampling rate affecting the data (effectively filtering out noise), we can compute the discrete Fourier transform of the data and view it in the frequency domain instead of the usual time domain. You can learn more about this [here](#).




Fig 11. CHICHI_'s replay with a 125Hz USB keyboard from earlier in the frequency domain

Keypress duration comparison

The second method we developed to analyze a player's replays involves comparing the lengths of keypresses (inputs). Depending on the control scheme and finger layout of the player, they are expected to have a consistent range of these lengths per in-game key across all of their relevant gameplay, as this is exactly what defines their physical playstyle or footprint.

To help visualize this, we developed a box plot chart split up per key. Any player is expected to generally follow the same ratio pattern between keys across replays characteristic to them only. A player is also expected to play with slightly shorter or longer durations while keeping this pattern, depending on the context of the gameplay.

For example, in a multiplayer setting where misdropping is fatal and caution is necessary a player is expected to play leisurely, making their keypress durations longer than usual.

Inversely, in a singleplayer setting, where resetting is easily accessible, a player will usually attempt swifter keypresses to achieve a more optimized run, making their keypress durations shorter than usual.

This is especially true for the 40 LINES game mode, and is less the case for the BLITZ game mode, due to the length of runs in each game mode respectively (there is more at stake in case of a mistake late into the attempt).

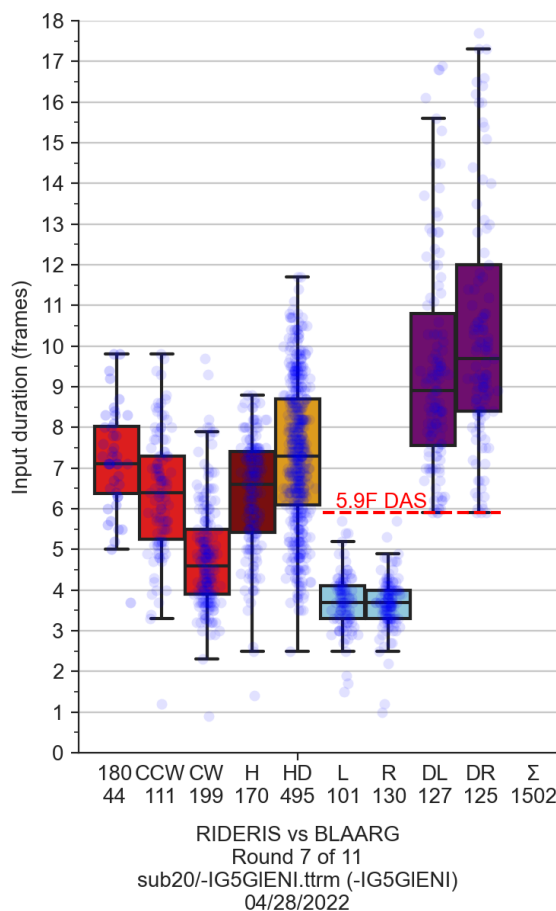


Fig 12. An example multiplayer round

The shorter the inputs, the harder they are to pull off consistently, meaning they occur less often in-game. Longer inputs don't occur often due to the high octane nature of the game when played at a high level. These two factors explain why a bell-curve-like distribution is observed for each key. When analyzing a replay, the shorter the duration of the inputs and the more consistent the distribution is, the more concern is raised.

It should be noted that during replay processing, DAS inputs are cut short at a Hard Drop event to prevent DAS preservations from overinflating against the main peak of a regular DAS press.

When questioning whether a player's solo mode keypress durations are realistic, public TETRA LEAGUE data for a player can be used as a trusted baseline of performance. This is for the same reasons outlined in the previous section regarding input timestamp trends. Here are some example comparisons of top sprint players' inputs from their 40 LINES personal bests with inputs from a randomly picked TETRA LEAGUE match:

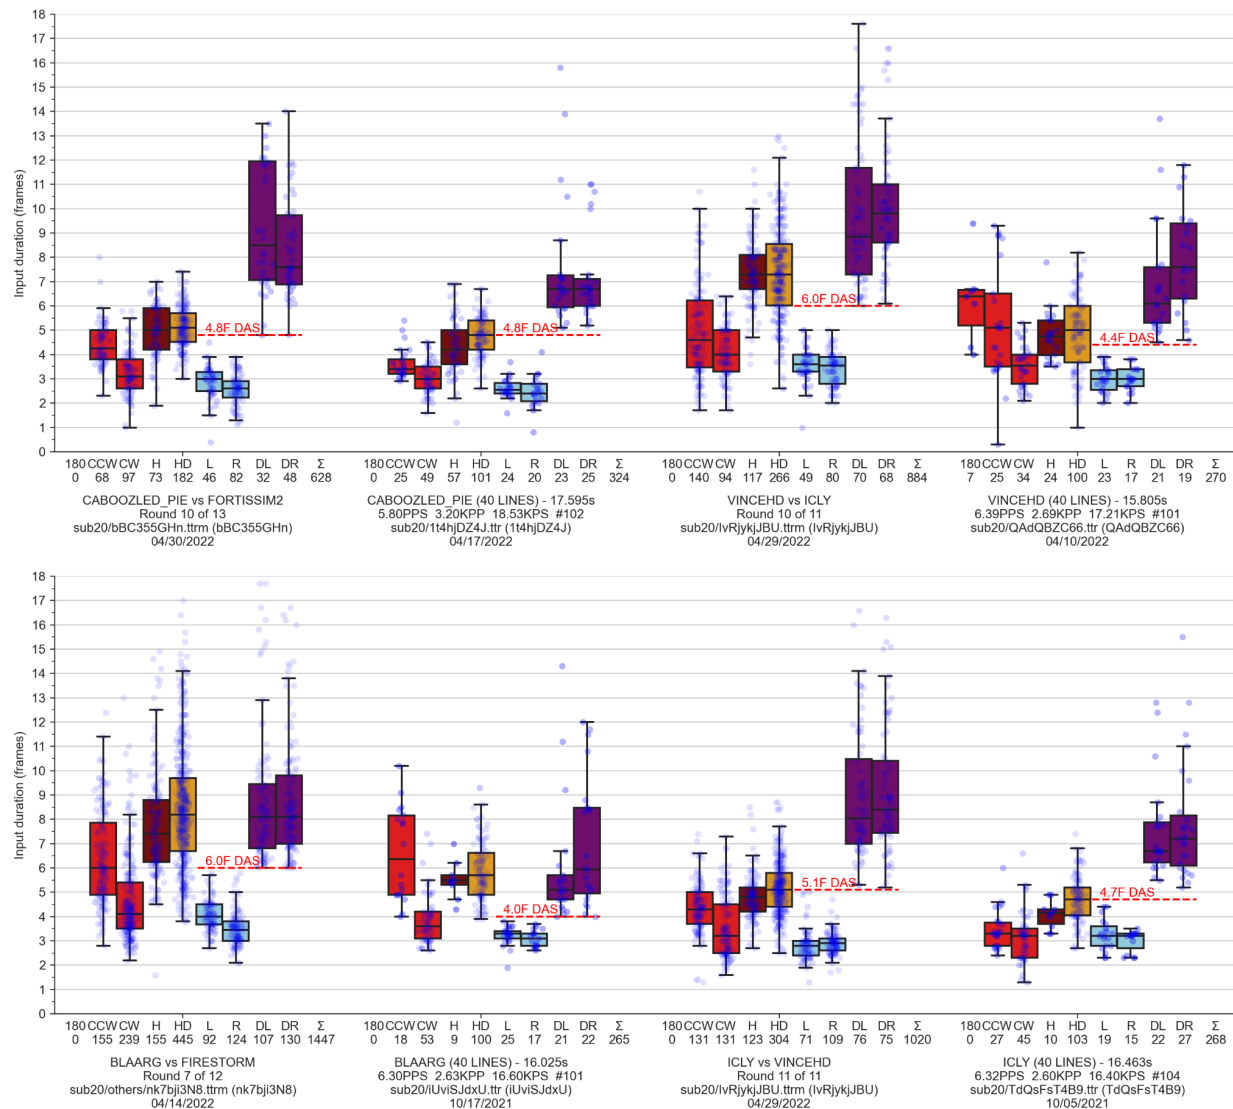


Fig 13. Example comparisons of 40 LINES inputs vs TETRA LEAGUE inputs

To get a solid indicator on just how much more swiftly a player is playing in a solo replay, we can calculate the p-value of the one-tailed Student's T-test using the inputs from the replay as a sample and a huge dataset of inputs from TETRA LEAGUE as a trusted population. A T-test is designed to compare multiple data sets to find out whether the data has been generated under the same conditions. The p-value calculated by this test shows the likelihood of a dataset being generated from a different source or under different conditions.

This p-value will not be used to accept or reject a null hypothesis as is the usual way to go when doing statistics in this way, but rather solely as an indicator as explained above. You can learn more about general applications of Student's T-test [here](#).

This T-test p-value can be calculated for every top player on the TETR.IO leaderboards. Once these values are calculated, it is possible to compare a player's p-values with a dataset of expected ranges of p-values from other players. If they are consistently an outlier on the unlikely side on multiple keys at once, we can argue that their run is perhaps too unrealistic compared to everyone else's relative performance.

To simplify these comparisons, we'll be comparing the natural log of p-values instead. Here's a chart displaying p-values for every player who is sub-20 seconds on the TETR.IO leaderboards as of the time of writing this document:

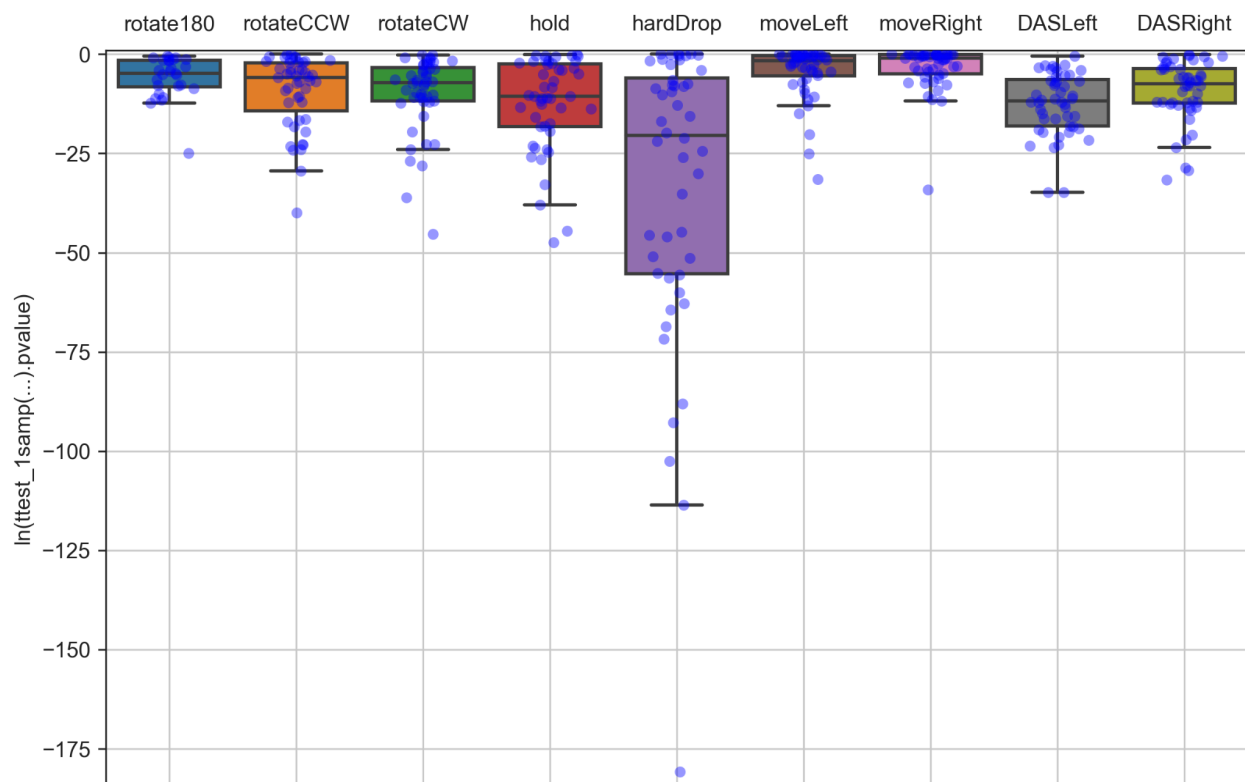


Fig 14. Dataset of p-values for every sub-20 40 LINES player on TETR.IO

As is expected, the largest difference is usually on the Hard Drop key. This is natural as it will be the key that speeds up the most, given the slow and safe style of play in TETRA LEAGUE vs the riskier, speedy style of solo modes. The DAS Left and DAS Right keys are also quite swifter than the rest, as most players lower their DAS significantly in solo modes to aid with speed. The other keys usually follow suit, as it is more convenient for a player to play more swiftly on all keys rather than just certain ones.

Analyzing FURRY's replays

Initial analysis

Upon learning of the previously described techniques, we applied them to FURRY's publicly available replay files from the TETR.IO leaderboards. In this section, we'll present the data from the replays and the results of our analysis. It's worth noting that at this point in the investigation more advanced comparison methods using Student's T-test and discrete Fourier transforms were not developed yet. Instead, we simply relied on visually understanding charted data. For this reason those will not be discussed here, but at a later point in the document.

We are using an initial assumption that the keyboard FURRY uses at all times is a USB 125Hz keyboard. All TETRA LEAGUE replays from FURRY that we've looked at have a 125Hz pattern that matches closely those of other players, namely JYIN and CHICHI_ whom we've reached out to and verified that they are indeed using 125Hz keyboards. RECOUNT also seems to be using a 125Hz keyboard, although we weren't able to reach them to confirm.

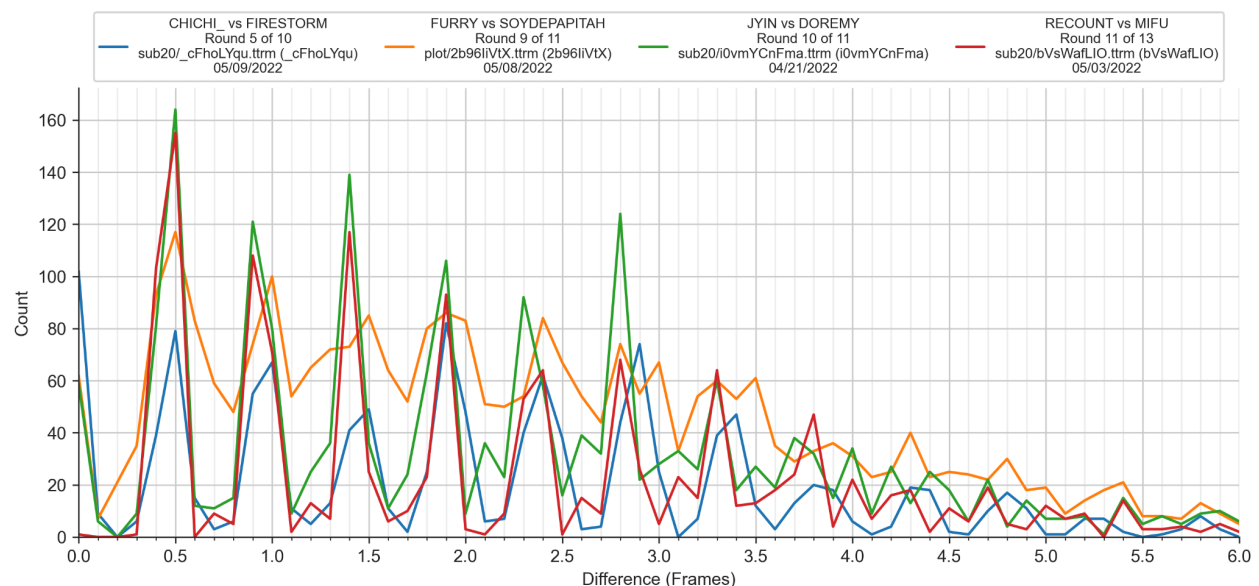


Fig 15. FURRY's TETRA LEAGUE data matches other players' known 125Hz keyboard data

Additionally, we've reached out to a player who wishes to stay anonymous that claims to own the same keyboard that FURRY uses to play. We've been able to determine that the keyboard they have is definitely 125Hz. FURRY has also posted a picture of their keyboard taken right before the 40 LINES world record to [Twitter](#) (tweet is now deleted), implying this is the keyboard used in the run. We were able to confirm the keyboard our anonymous party owns shares physical characteristics with FURRY's keyboard, although it seems to have been rebranded.



Fig 16. FURRY's keyboard from the tweet on the left, opposite our anonymous party's keyboard

FURRY's top 10 records in both 40 LINES and BLITZ, however, do not match the expected 125Hz pattern at all. Instead, it seems that the wave-like pattern oscillates at a higher rate than that. A few analyzed replays are shown here alongside a TETRA LEAGUE replay for reference:


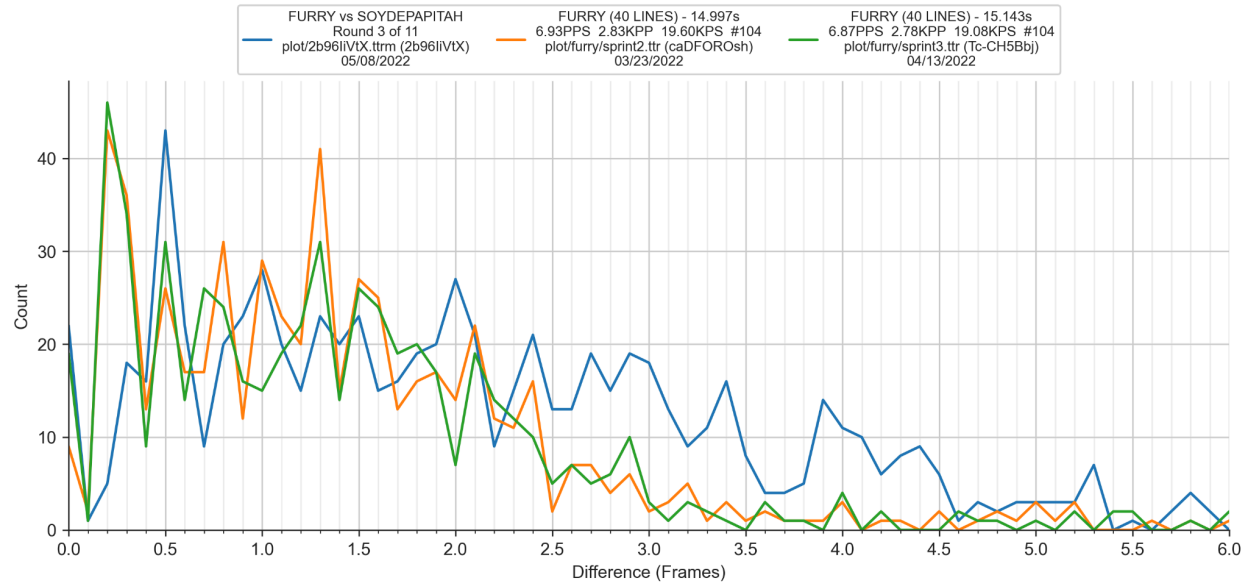


Fig 17. The current 40 LINES world record run does not match TETRA LEAGUE data

It's as if the 40 LINES world record's data has been stretched and fit towards lower difference values. We think this indicates the two were not created under the same conditions, and the only plausible explanation for this would be if FURRY had used different keyboards for these 40 LINES finishes.

	Date	Finesse	KPP	KPS	Pieces	PPS	Time
#2	3/23/2022, 8:20:27 AM	2F (98.08%)	2.83	19.60	104	6.93	0:14.997
#3	4/13/2022, 9:52:23 AM	2F (98.08%)	2.78	19.08	104	6.87	0:15.143



#4	3/6/2022, 5:44:43 PM	7F (94.06%)	2.94	19.43	101	6.61	0:15.283
#5	2/23/2022, 4:08:44 PM	6F (95.1%)	3.04	20.14	102	6.63	0:15.393




Fig 18. FURRY's other top 40 LINES runs also don't match TETRA LEAGUE data

FURRY's other 40 LINES replays also don't seem to follow the pattern observed in TETRA LEAGUE data, but instead closely match the world record.

	Date	Finesse	SPP	Level	Pieces	PPS	Score
#1	3/18/2022, 3:27:41 PM	57F (91.57%)	1830.13	14	498	4.15	911,403




Fig 19. FURRY's BLITZ personal best does not match TETRA LEAGUE data

Much like the 40 LINES replays, FURRY's BLITZ personal best also doesn't fit against TETRA LEAGUE data in a similar way. This particular BLITZ replay was initially automatically flagged as [suspicious](#) upon submission, and had to be manually submitted by TETR.IO staff.

#2	4/11/2022, 10:36:24 AM	88F (89.74%)	1771.01	13	458	3.82	811,124
#3	1/28/2022, 6:50:50 PM	66F (91.97%)	1656.60	13	461	3.84	763,693




Fig 20. FURRY's other top BLITZ runs also don't match TETRA LEAGUE data

We've actively checked multiple weeks' worth of FURRY's TETRA LEAGUE data for a pattern that would match his solo mode replays. So far, we've only been able to observe data that matches a regular USB 125Hz keyboard. Given our assumption that FURRY is using the same keyboard at all times, we've been unable to find a plausible explanation for such differences in our comparisons.

Aside from these trends, we've also looked at the difference between FURRY's keypress durations in TETRA LEAGUE and in solo mode replays:

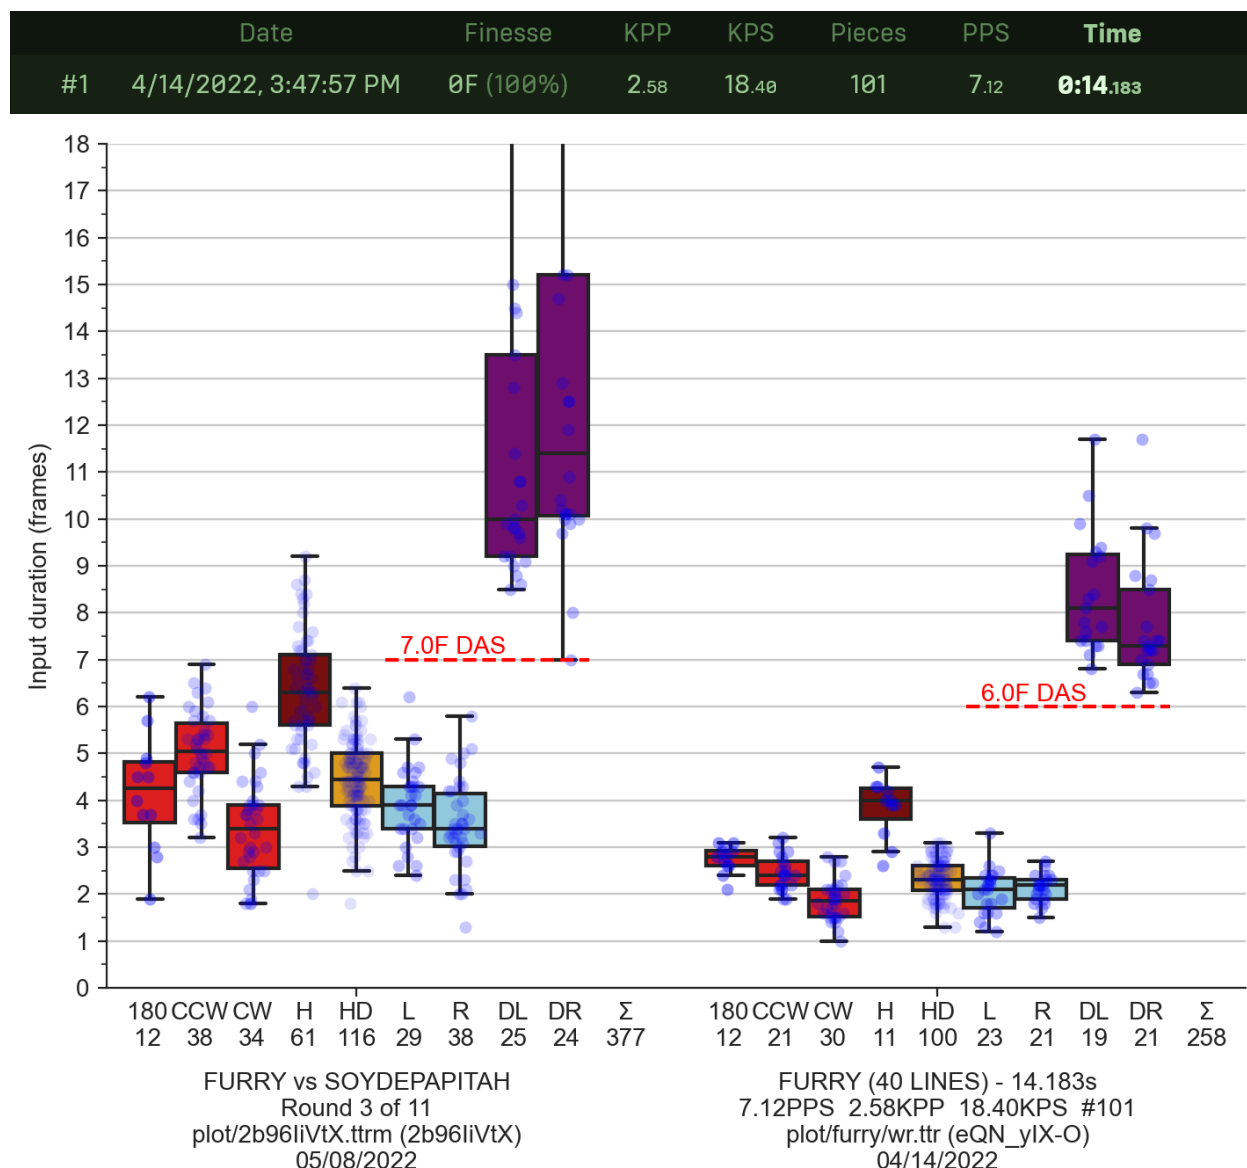


Fig 21. FURRY's keypresses in TETRA LEAGUE compared to his 40 LINES world record

We observe FURRY's keypresses to be significantly shorter and more consistent in 40 LINES compared to his TETRA LEAGUE data, more so than [other top players](#). It should also be noted that there is not a single move keypress between 4F and 6F immediately under his DAS.

	Date	Finesse	SPP	Level	Pieces	PPS	Score
#1	3/18/2022, 3:27:41 PM	57F (91.57%)	1830.13	14	498	4.15	911,403

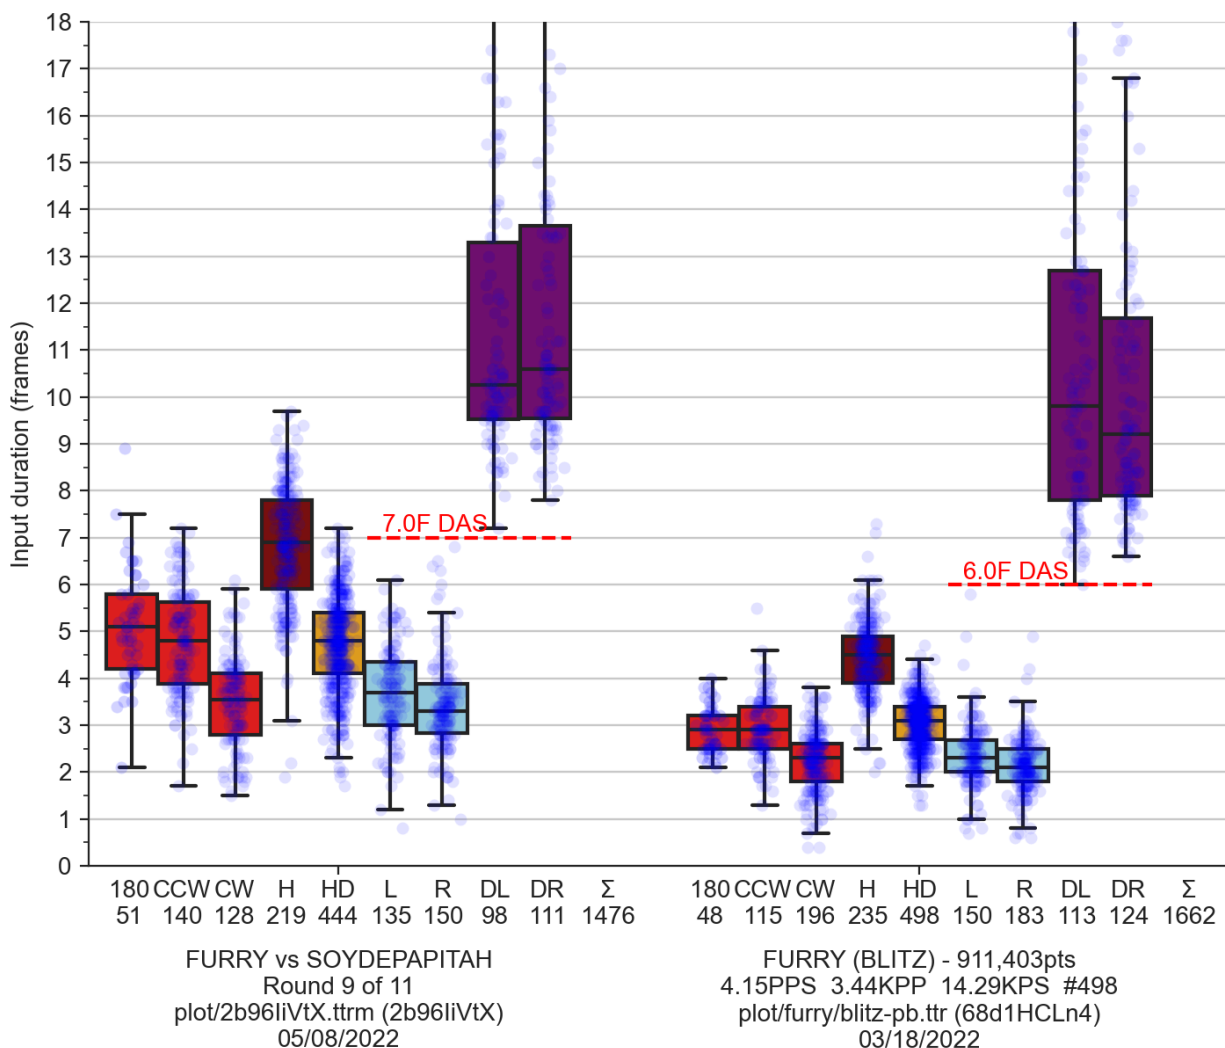



Fig 22. FURRY's keypresses in TETRA LEAGUE compared to his BLITZ personal best

A similar, swifter pattern emerges when analyzing his BLITZ personal best, albeit to a lesser degree than the 40 LINES replays. Again, his swiftness changes from TETRA LEAGUE to BLITZ much more than his peers.

Here's some more data showing the same phenomenon occurring with other 40 LINES and BLITZ replays:

	Date	Finesse	KPP	KPS	Pieces	PPS	Time
#2	3/23/2022, 8:20:27 AM	2F (98.08%)	2.83	19.60	104	6.93	0:14.997
#3	4/13/2022, 9:52:23 AM	2F (98.08%)	2.78	19.08	104	6.87	0:15.143
#4	3/6/2022, 5:44:43 PM	7F (94.06%)	2.94	19.43	101	6.61	0:15.283



	Date	Finesse	SPP	Level	Pieces	PPS	Score
#2	4/11/2022, 10:36:24 AM	88F (89.74%)	1771.01	13	458	3.82	811,124
#3	1/28/2022, 6:50:50 PM	66F (91.97%)	1656.60	13	461	3.84	763,693
#4	12/27/2021, 12:45:01 PM	89F (88.63%)	1532.20	13	475	3.96	727,793

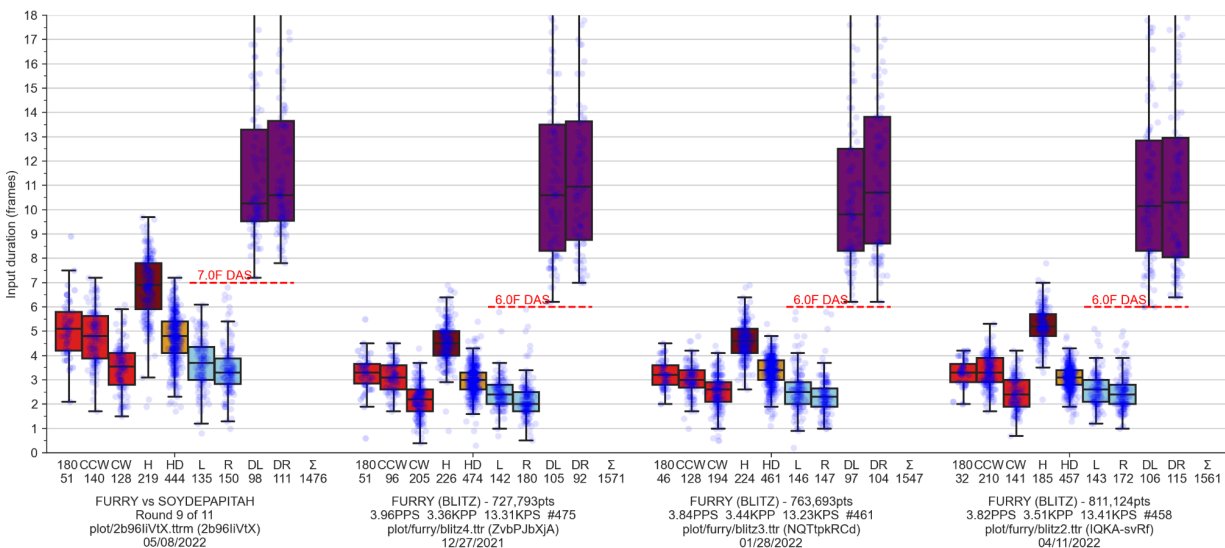


Fig 23. FURRY's other 40 LINES and BLITZ replays also don't match TETRA LEAGUE data

While there's nothing conclusive that can be deduced from these, it definitely seems FURRY "swifts up" a lot more than their peers, which is additional motivation for concern.

Voice call

After failing to find plausible explanations for contradictions our initial analysis uncovered, we presented our findings to FURRY in a voice call on 05/03/2022 to gather further insight and understanding of the data we obtained.

We started with introductions and an explanation of the phenomena discussed earlier. First we discussed keypress durations. FURRY was shown VINCEHD and BLAARG's keypress duration comparisons between TETRA LEAGUE and 40 LINES as examples of other players' data. Worth noting is that although VINCEHD and BLAARG tune their DAS lower for solo modes than multiplayer similarly to FURRY and most other high-level 40 LINES players, their input duration charts do not show nearly the same consistency and swiftness change as FURRY's.

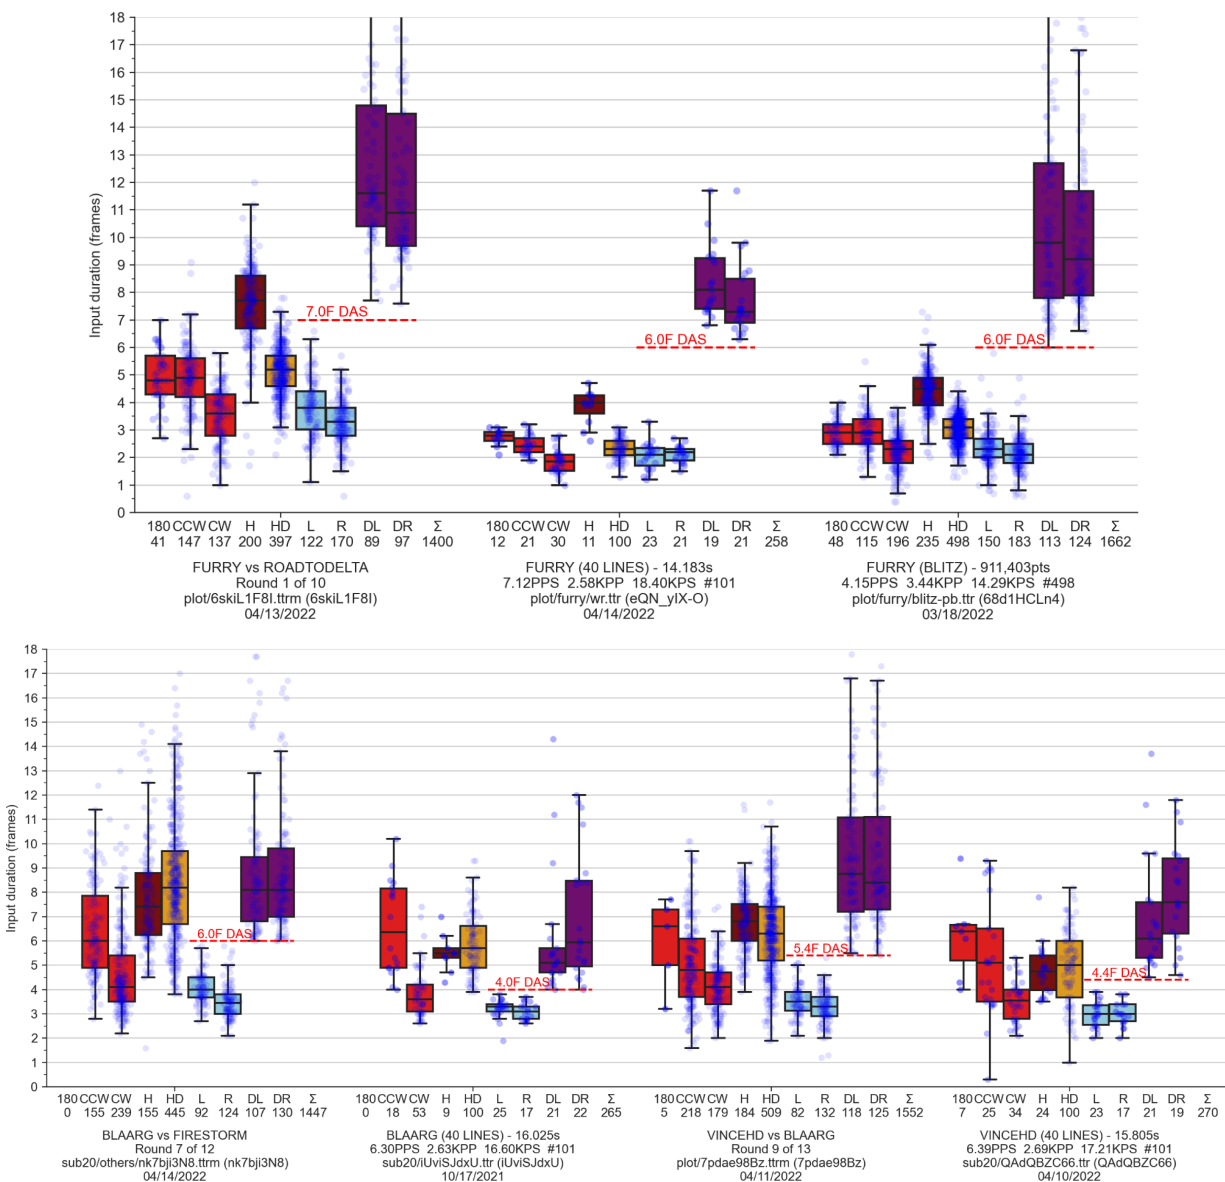


Fig 24. Keypress duration data presented to FURRY during the voice call

FURRY is shocked about the disparity of his data compared to VINCEHD and BLAARG. He explains that his 40 LINES playstyle involves his hands lightly "vibrating" on the keys, whereas his playstyle in TETRA LEAGUE involves longer-lasting and more stable keypresses, and that this is the most likely reason why his inputs are extremely short in 40 LINES.

Next, we presented the input timestamp difference trends to FURRY, comparing FURRY's BLITZ personal best to a TETRA LEAGUE match. We explained that, with a 125Hz keyboard, the expected data should always keep the pattern from the TETRA LEAGUE match, with a peak every 0.48 frames (8 ms) the charts that result from FURRY's BLITZ personal best should be impossible to produce with the given hardware.

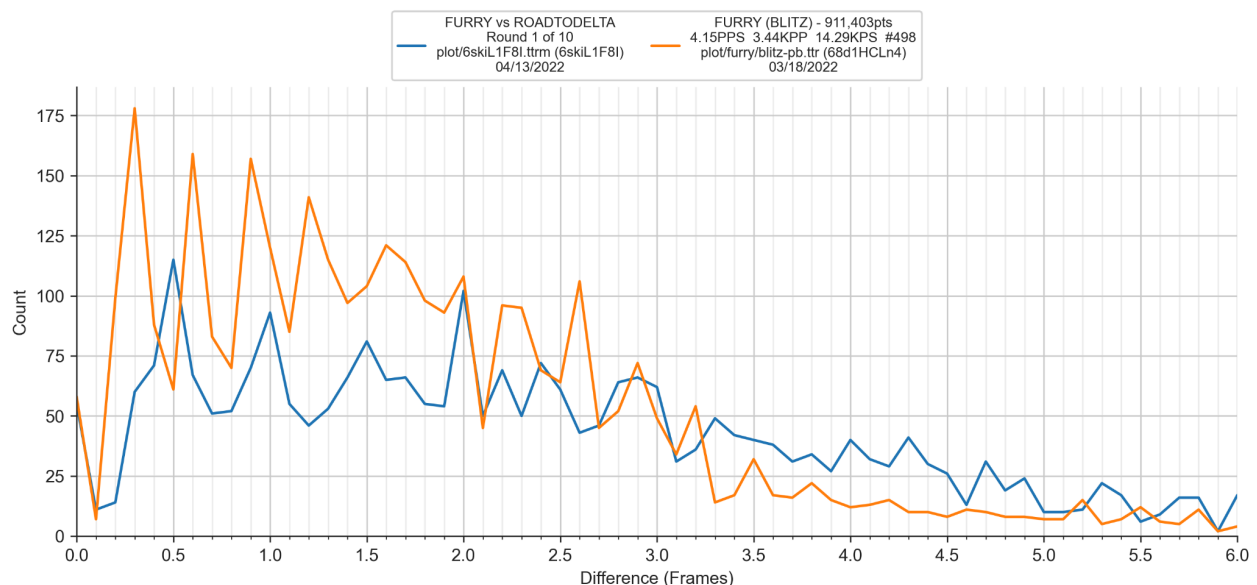


Fig 25. Input timestamp difference trend data presented to FURRY during the voice call

FURRY responds saying that his keyboard could be defective, as it is cheap (\$4 keyboard), very used, and has a broken spacebar. He also said that he does not know why the data comes out looking strange, and asked about the effects of game FPS and lag spikes on the data.

As of the time of writing this document, we lack the resources to tell precisely how FPS and lag spikes affect the recorded timestamps. However, we can make an educated guess that only frequent, consistent and severe lag spikes would cause the data to behave in an erratic way opposed to 125Hz as we see here. We believe the game itself would become much harder to play under these conditions, to the point where achieving a new personal best, let alone a world record, would be unreasonable.

We then asked FURRY to play two runs of BLITZ and two runs of 40 LINES to gather trusted solo mode replay data for further analysis. We asked FURRY to do this while screen-sharing his gameplay to help us confirm the live runs' legitimacy. Before beginning these runs, FURRY notes that for almost a month now (since the 40 LINES world record was set) he has changed his playstyle completely due to learning multiplayer and can no longer speed up and play fast.

In the two BLITZ runs, we instructed FURRY to take things easy and not necessarily go for a high score, since we're not looking for a correlation between his average performance and peak performance. FURRY ended the first BLITZ with a score of 398,105 (2SKAM6pYG.ttr) and the second with a score of 462,667 (o2Wdf21eX.ttr). The two 40 LINES runs that he completed for us were a 32.222s finish with 3.35pps (lqXEGldlf.ttr), and a 27.173s finish with 3.72pps (JpgQntZX5.ttr).

After the first run of 40 LINES, FURRY told us that he was playing using his multiplayer style and that his hands were freezing from the cold, further inhibiting his performance. After the second run, we asked if he could finish a 4.5pps run (~22-23s sprint time), to which FURRY answered that he was unsure if he could even manage 4 PPS given the conditions.

Afterwards, we played a few multiplayer matches against FURRY. We asked FURRY to send us all of the replays played out during the voice call, which amounted to the two aforementioned BLITZ runs, two 40 LINES runs, and two multiplayer matches (a FT3 and FT1).

We concluded the call by asking FURRY to run a keyboard polling rate test program called [Keyboard Inspector](#), which verified our assumption that FURRY is using a 125Hz keyboard.




Fig 26. The polling rate of FURRY's keyboard was determined to be 125Hz

FURRY mentioned during this conversation that he has a particularly poor Internet connection. As a result, FURRY claims he experiences frequent disconnections with the TETR.IO servers, even in the middle of playing solo modes, and that this has caused many of his replays to be marked as suspicious. As of the time of writing this document, we cannot confirm if this is standard TETR.IO behavior. During the voice call itself, FURRY has experienced multiple connection losses which caused him to temporarily disconnect from the voice call, particularly during moments when his network was stressed due to file uploads and downloads.

We have consistently assured FURRY throughout the call that it is okay to ask questions and tried our best to relieve as much pressure as possible. FURRY has shown that he understands how the data visualizations were generated and what they mean.

Post-call analysis

We used the same techniques to analyze the replays collected during the voice call:




Fig 27. Comparison of FURRY's TETRA LEAGUE data, personal bests, and call replays

The replays FURRY played out for us during the voice call are much more similar to his TETRA LEAGUE replays, and we can make out a pattern resembling a 125Hz keyboard as is expected. His personal bests do not match neither TETRA LEAGUE data nor the data from the call.

After the call, we developed more advanced comparison methods using Student's T-test and discrete Fourier transforms. We can now view this data in the frequency domain, instead of the time domain as we have so far:



Fig 28. The same comparison viewed in the frequency domain instead of the time domain

The frequency domain displays how much of the data is present among each given sampling frequency. This lets us clearly see which sampling frequency is the most dominant, and as such was most likely an actual sampling rate present during the replay's creation.

FURRY's TETRA LEAGUE data shows a very clear peak almost exactly at 125Hz. The replays collected from the call peak very close to this value, almost exactly at 120Hz. Meanwhile, his BLITZ personal best shows two peaks: there is one around 185Hz, and a slightly smaller one around 210Hz. His 40 LINES world record does not contain a lot of inputs, so the peak isn't quite as obvious, but it seems to reach the highest point at 235Hz. This suggests that these two

replays have not been created under the same conditions as TETRA LEAGUE data and our voice call with FURRY, especially because they lack a peak on the expected 125Hz frequency. Looking at FURRY's other top 40 LINES and BLITZ replays, we observe they also lack a peak on the expected 125Hz frequency, and are rather similar to each other in their respective game mode:

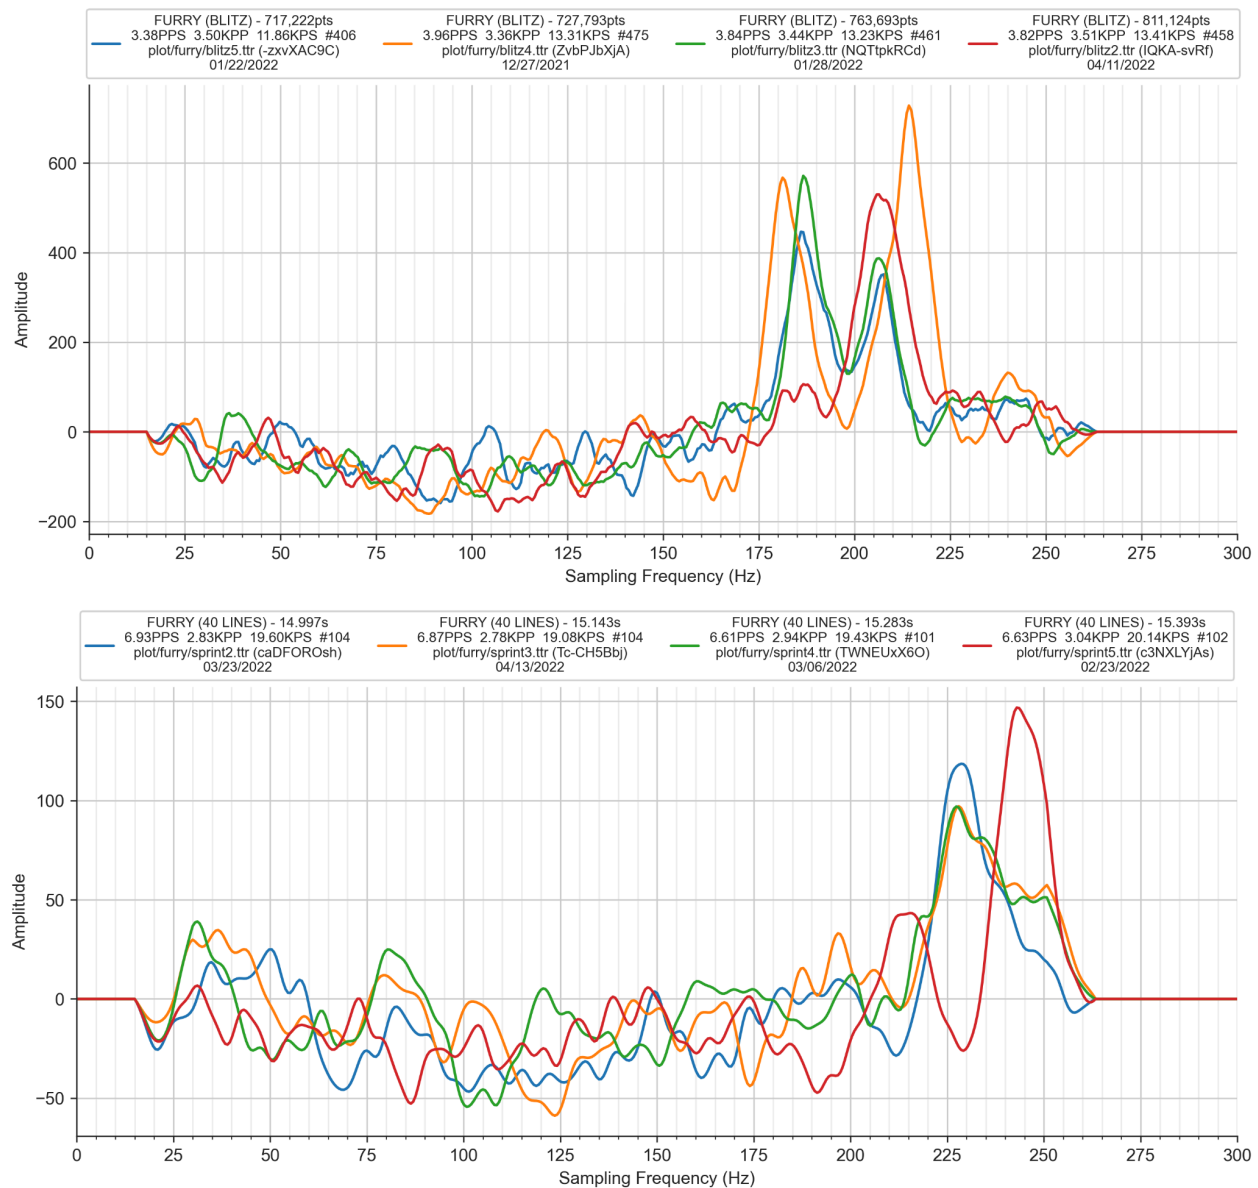


Fig 29. FURRY's other top 40 LINES and BLITZ replays viewed in the frequency domain

Next, we compared keypress durations for the same replays:

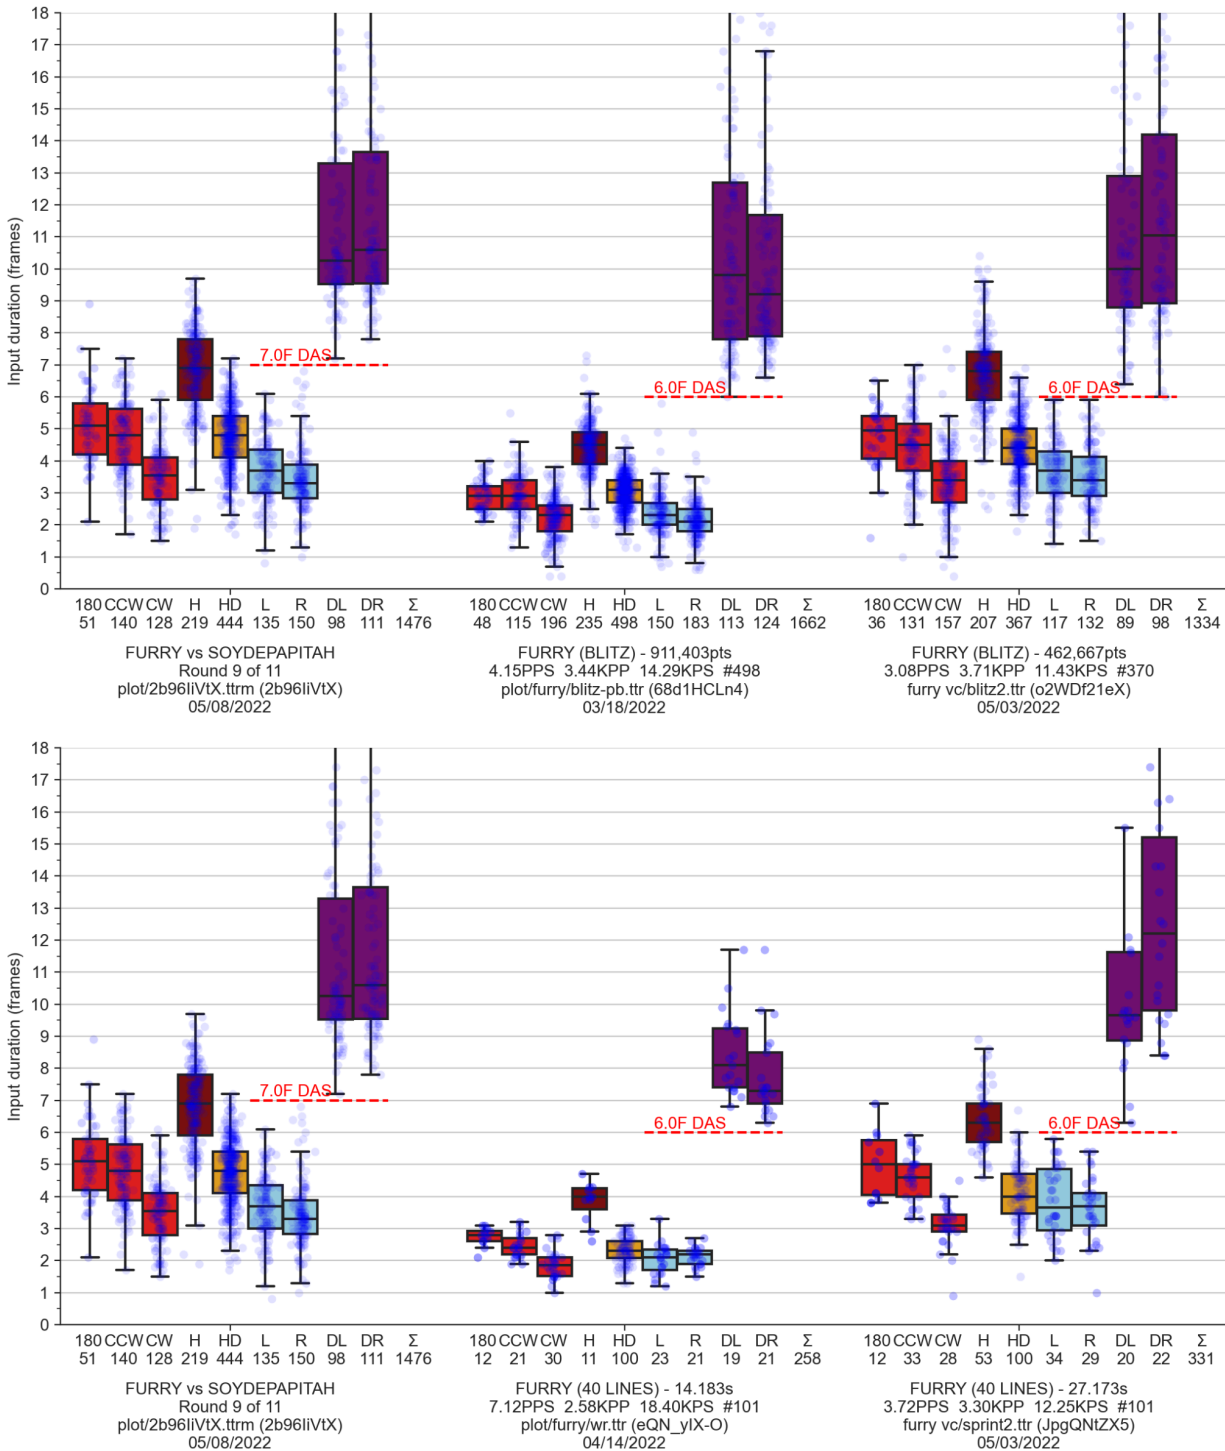


Fig 30. Comparison of FURRY's TETRA LEAGUE data, personal bests, and call replays

In BLITZ, FURRY's keypresses don't change between TETRA LEAGUE and the call replay. When it comes to 40 LINES, FURRY's keypresses do get swifter in the call replay compared to TETRA LEAGUE data, but not nearly as much as they do in his personal best. This is expected because the call replays were relaxed runs and not FURRY's peak performance.

With the newly developed Student's T-test, we decided to compare FURRY's swiftness change between his TETRA LEAGUE performance and 40 LINES personal best against every other sub-20 second player on the TETR.IO 40 LINES leaderboards:




Fig 31. Results of the T-test for FURRY against every other sub-20 second player

FURRY changes his swiftness between TETRA LEAGUE and his 40 LINES world record the most out of any other sub-20 second player for all keys except Hold, DAS Left, and DAS Right. This is expected for the DAS keys given that he plays at a significantly higher DAS than most of his peers, and his DAS change between TETRA LEAGUE and solo modes is not as high. FURRY lowers his DAS from 7F to 6F for solo modes, while most others go from 5.5F to 4F.

We argue the results of the T-test presented here are **extremely unrealistic** and point towards the game being **externally modified** in some way when playing out their world record.

With this newfound data, we attempted to contact FURRY in the days following our initial voice call, asking if we could call again to try to live-recreate runs that would be able to disconfirm our suspicions on his top records in 40 LINES and BLITZ. Unfortunately, we were unable to organize the meeting after multiple attempts.

Another interesting replay that we haven't discussed so far is the 16.172 40 LINES liveplay available on FURRY's YouTube channel ([POV: you're anthony07](#)). Only after the call, we realized this replay is still available on TETR.IO through the "recent scores" menu:

	RECENT	40 LINES	BLITZ			
Game Mode	Date	Finesse	Pieces	PPS	Result	
40 LINES	5/3/2022, 5:29:40 PM	27F (88.12%)	101	3.72	0:27.173	➔ VIEW
40 LINES	5/3/2022, 5:26:49 PM	22F (90.74%)	108	3.35	0:32.222	➔ VIEW
BLITZ	5/3/2022, 5:22:38 PM	85F (87.3%)	370	3.08	462,667	➔ VIEW
BLITZ	5/3/2022, 5:18:21 PM	67F (88.62%)	334	2.78	387,105	➔ VIEW
40 LINES	4/17/2022, 6:14:59 AM	10F (94.12%)	102	1.41	1:12.275	➔ VIEW
40 LINES	4/16/2022, 6:56:53 PM	7F (93.46%)	107	3.08	0:34.765	➔ VIEW
40 LINES	4/16/2022, 6:52:15 PM	27F (86.24%)	109	2.31	0:47.187	➔ VIEW
40 LINES	4/16/2022, 6:31:21 PM	6F (95.19%)	104	5.93	0:17.527	➔ VIEW
40 LINES	4/16/2022, 6:13:54 PM	5F (96.08%)	102	6.31	0:16.172	➔ VIEW
40 LINES	4/16/2022, 6:11:10 PM	4F (97.3%)	111	6.09	0:18.222	➔ VIEW

Fig 32. FURRY's most recent 10 scores on 05/15/2022

The first four available replays are the ones we asked FURRY to play out in front of us during the voice call. The next three replays seem to be random runs where FURRY is just messing around, practicing setups and T-spin builds. The last five replays are from the session during which the POV run was played out. The second to last replay is the one from the POV video.

The three random 40 LINES runs have a dominant sampling frequency of 125Hz, just like the replays from the call:

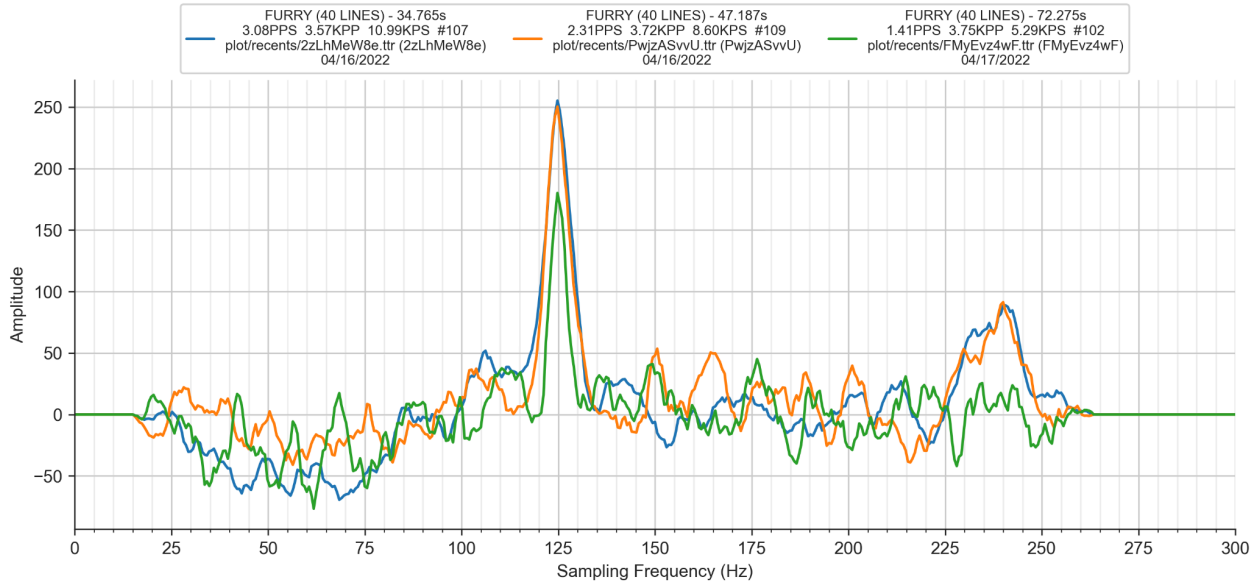


Fig 33. FURRY's three random replays viewed in the frequency domain

However, looking at the last three replays from the list (including the POV replay), they lack a 125Hz sampling frequency peak, but rather have a very obvious peak around 210Hz much like FURRY's other top replays:

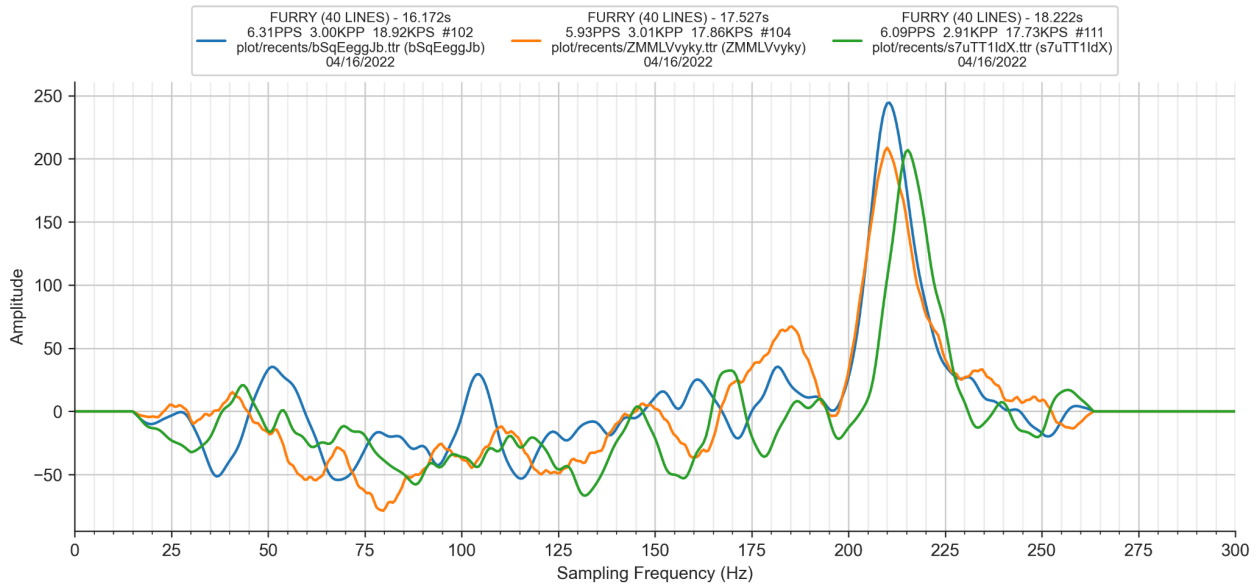


Fig 34. FURRY's last three recent replays viewed in the frequency domain. POV run is blue.

Given that the keyboard FURRY used to play out all these replays is one that has a USB report rate of 125Hz and we could not find another plausible explanation that would cause such a discrepancy between the expected data and actual data from the replays, **we argue that FURRY externally modified their game in some way when playing out these replays.**

Conclusion

After looking over our findings multiple times over and composing this document, we strongly believe that **FURRY used external assistance**, most likely by **slowing down their game**, to **illegitimately achieve their top records**. This unfortunately includes his **14.183 40 LINES world record**.

With the evidence in this document, we reached out to osk and the TETR.IO moderation team to present them the data we had collected and the conclusions we reached from it, in hopes they would analyze the provided evidence in full and take whatever actions they feel are the most appropriate.

On May 17th 2022, **FURRY publicly admitted to having cheated his top records**, and has since been permanently banned from TETR.IO without opportunity to appeal the ban. FURRY expresses remorse and apologizes for what he has done. For more information on TETR.IO staff's end of the investigation, please read [osk's blogpost](#) on the subject.

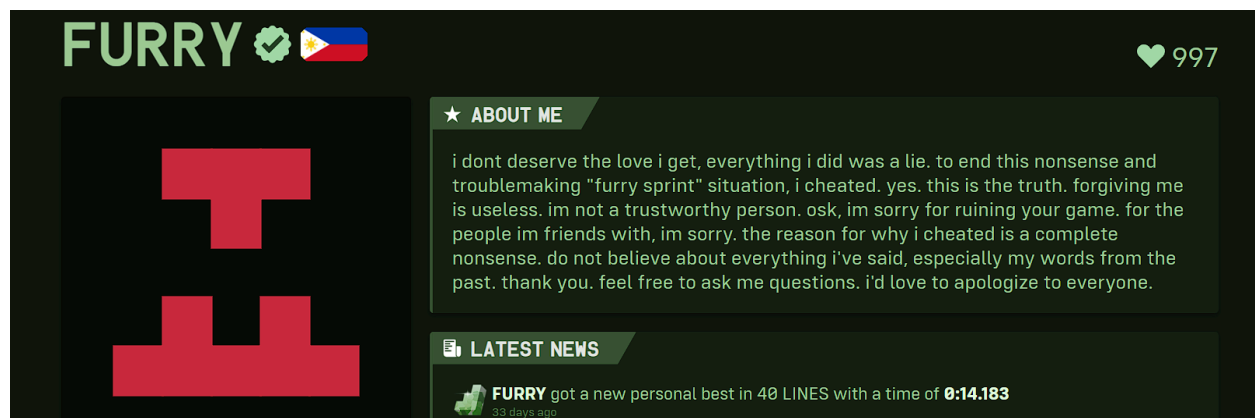


Fig 35. FURRY's public apology, on his TETR.IO profile prior to account ban

Going forward, we hope that the community will become stricter with regard to standards for top-level achievements, including frequent handcam videos and live-streamed gameplay. We think that it is crucial for stacker games' moderators to study the methodology presented in this document to prevent players from abusing leaderboards in the future.

Extra Credits

Arisalchigaya – for proofreading the document
frey – for their individual efforts and rechecking our findings
Fug Ebin Meme, hiryu, Syperion and others – for helping with data collection and organization
Ghostforce, kmzzz – for helping with and rechecking our math
Justin1L8 – for helping with understanding KPS data
mystery, z2sam – for helping organize early in the investigation