

Breaking the Ankh

Deterministic Propagation Netcode in 'Spelunky 2'

Guillermo NWDD

Race condition survivor at BlitWorks

Guillermo NWDD

BlitWorker @BlitWorks
Undergrad thesis tutor @UOCuniversitat
Game console prog. @ENTI_BCN/@UniBarcelona

Also, random Master's degree student

Blog: <https://rhet.dev>

Linkedin: <in/xnwdd>

Twitter: [@xnwdd](https://twitter.com/xnwdd)

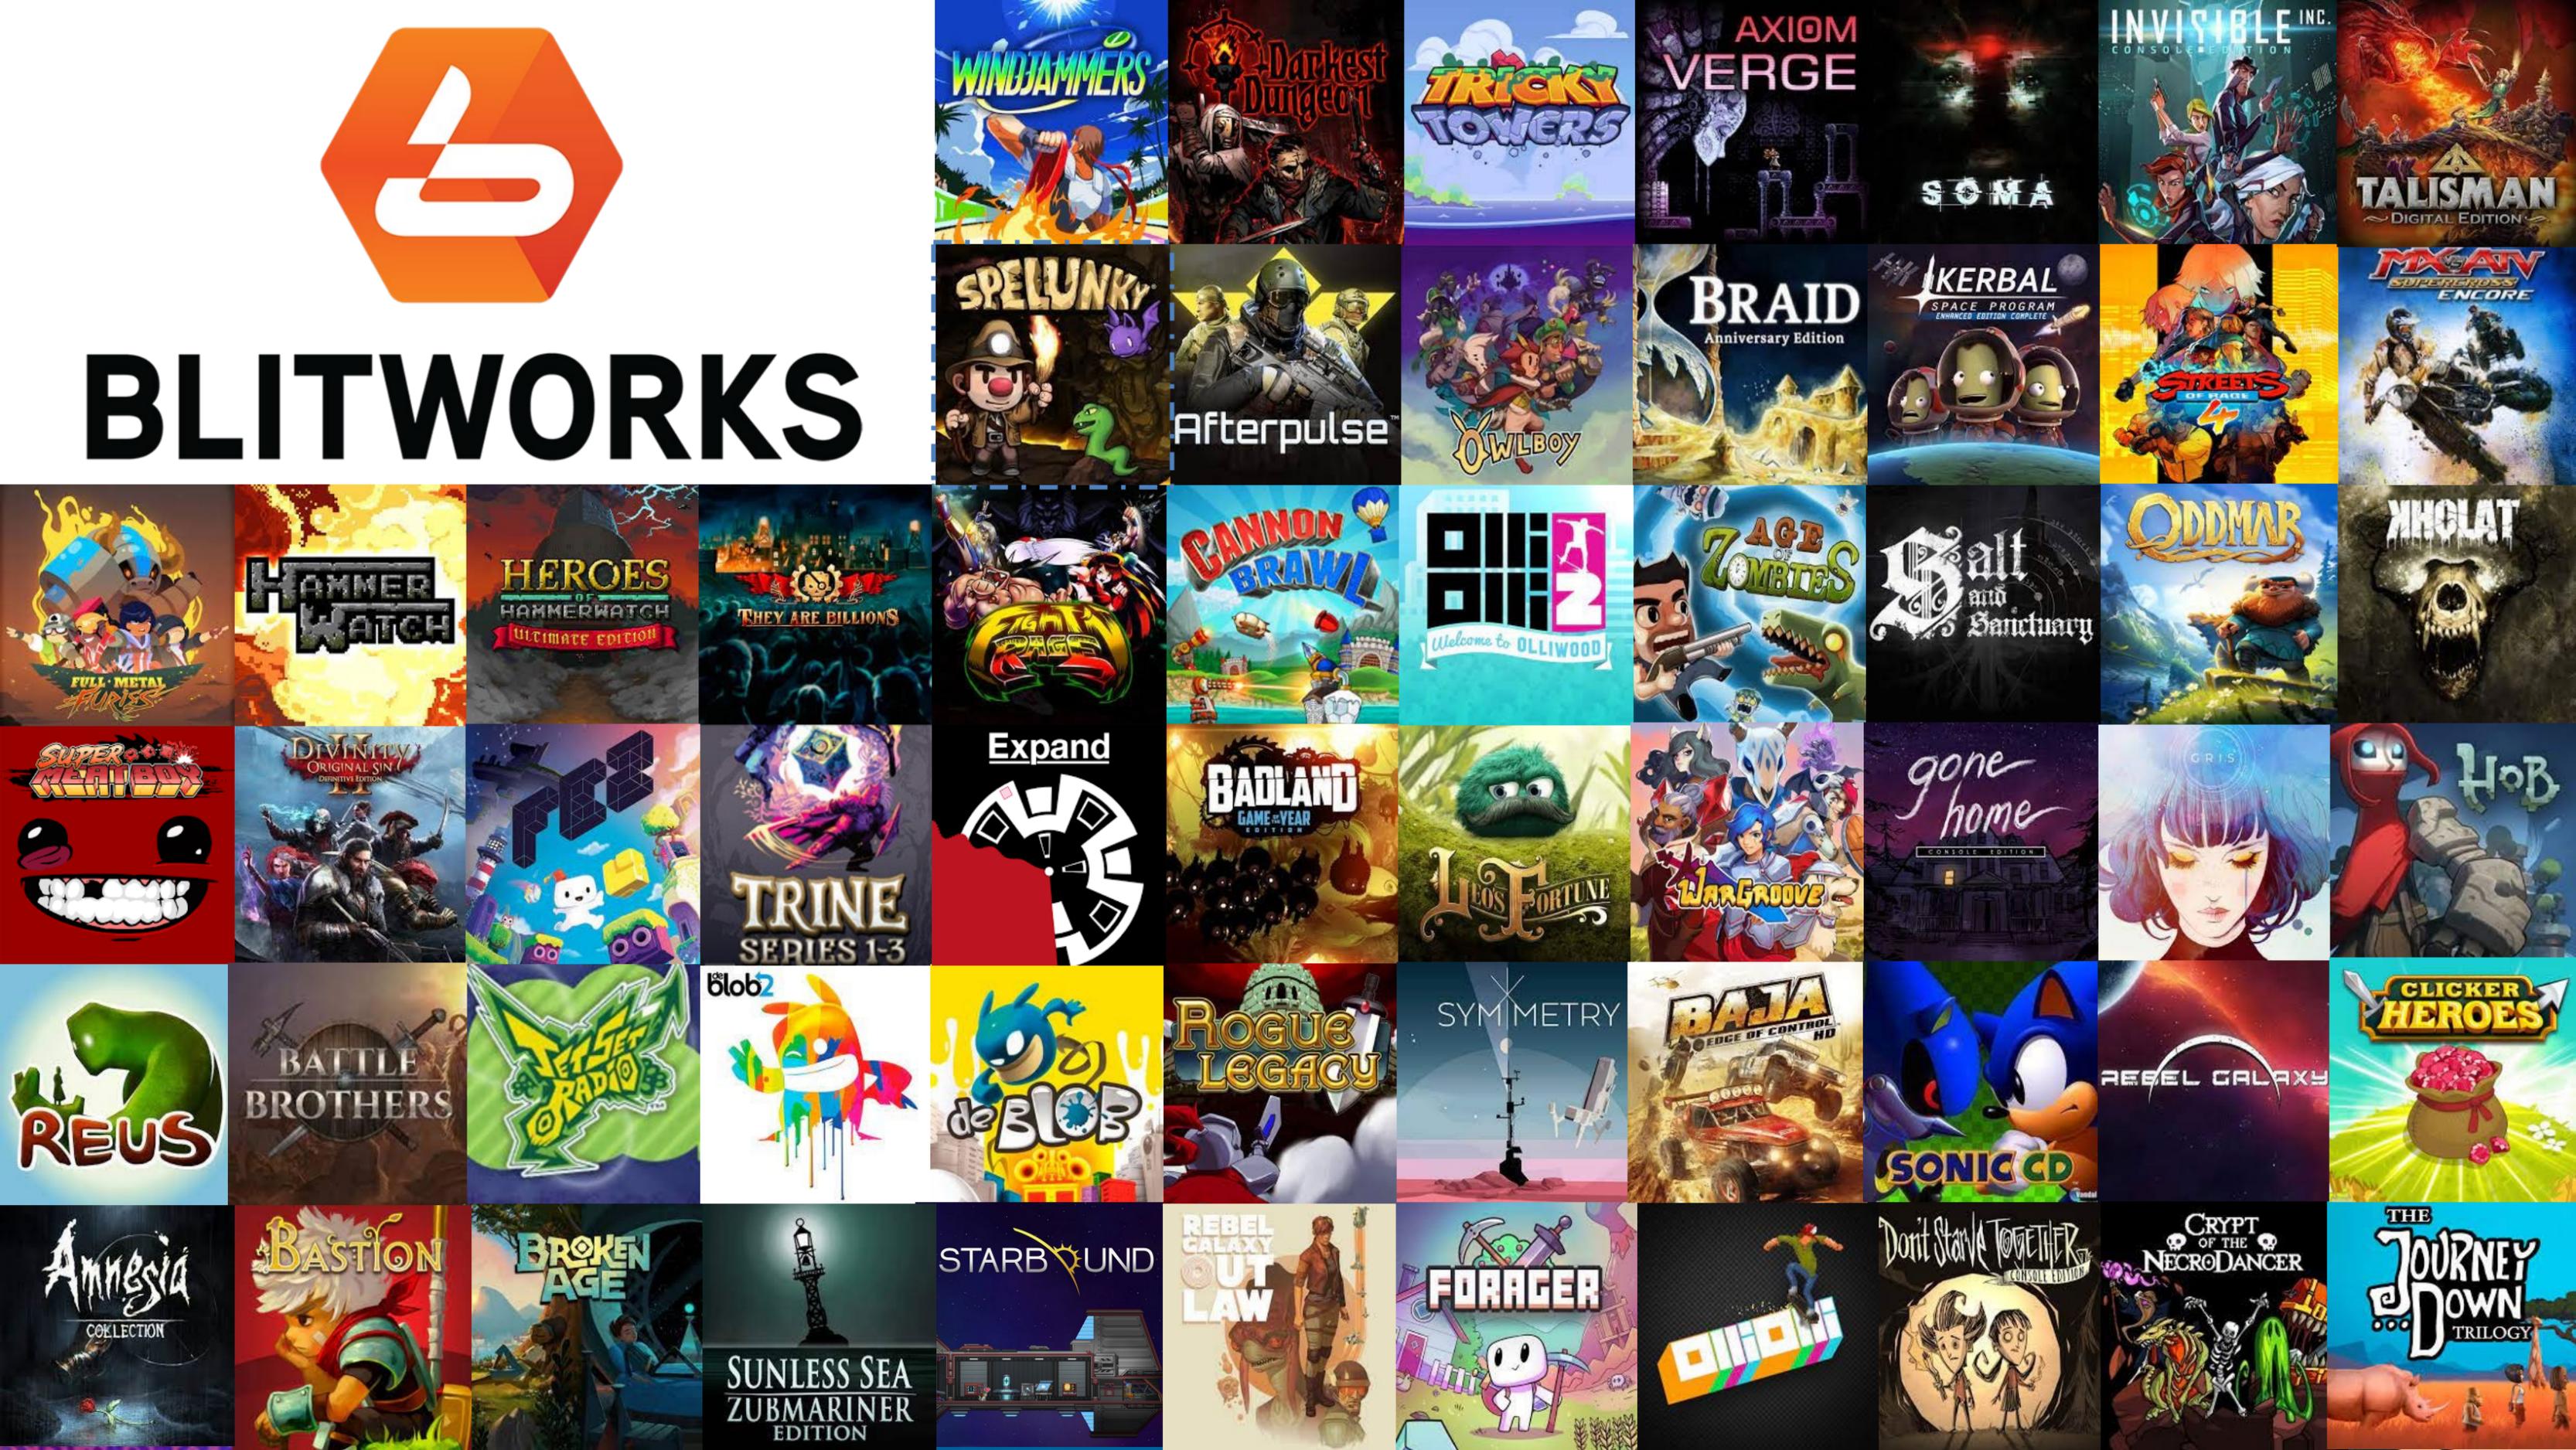


Universitat
Oberta
de Catalunya





BLITWORKS



Meet Spelunky 2

Is a:

- ❖ Precision platformer
- ❖ with changing room layout
- ❖ and short state-less runs

It's about:

- ❖ Emergent stories
- ❖ Game mechanics knowledge
- ❖ Skill (only you change between the 1st and 1000th run)



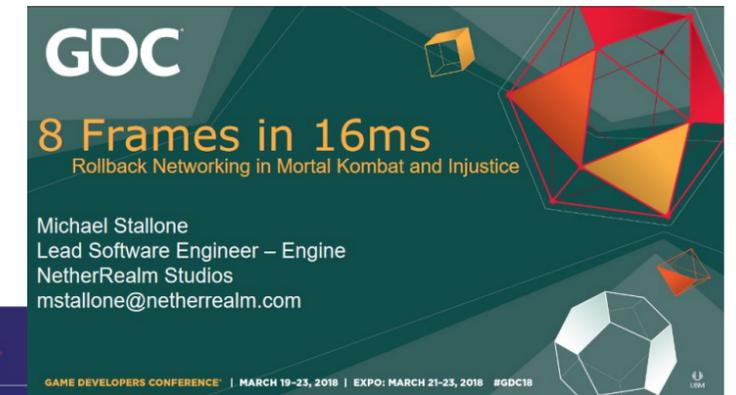
mossmouth[®]

Talk Goals

Talk focused on:

- ❖ Problems we didn't expect
- ❖ Netcode
- ❖ Servers

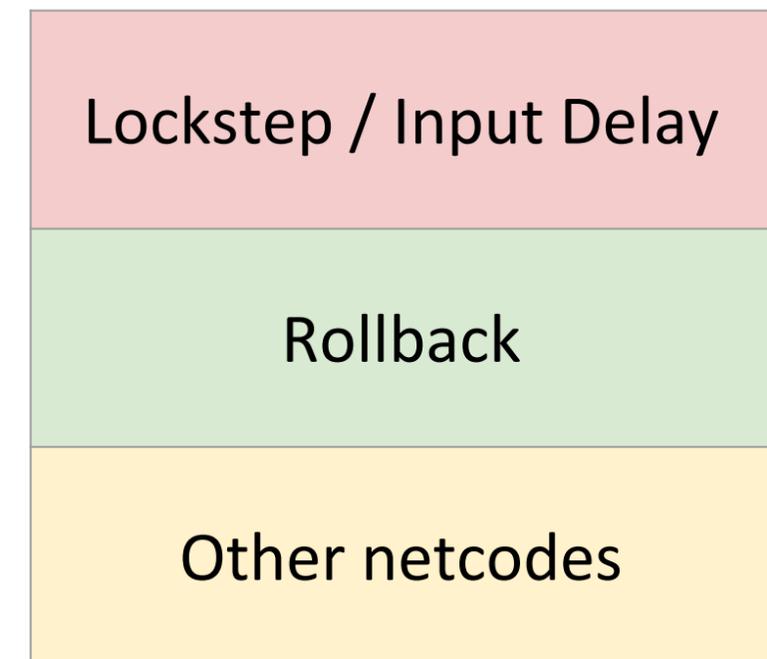
Alternative talk title: What I wish I had known or given more importance after watching preceding talks/literature but before starting to work on a server-hosted IAP-free deterministic online multiplayer for a game of the scope and complexity of Spelunky 2 as the only Network programmer of a really small indie team.



Day -461: First contact with Spelunky 2

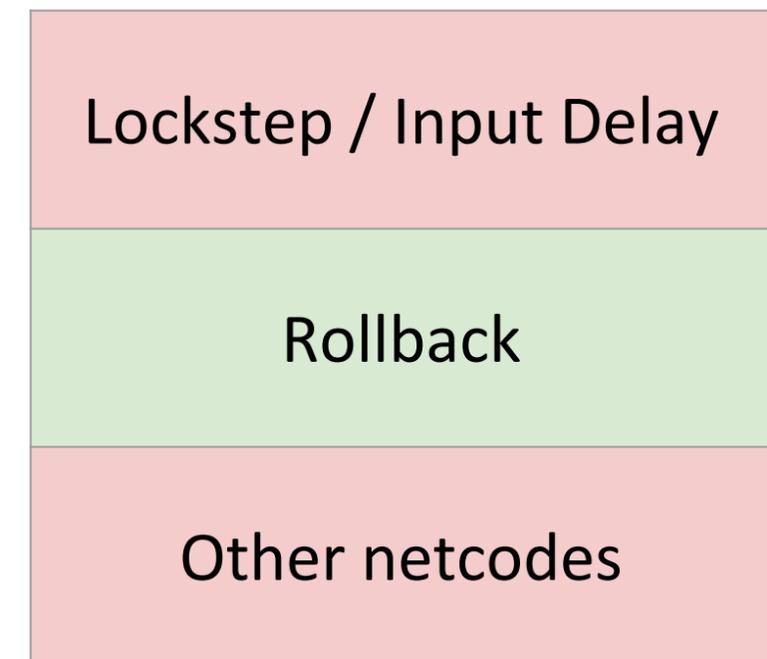
Conceptual Constraints

- ❖ Arena must be fair
- ❖ The whole huge procedural stage must update
 - Probably not a problem given Spelunky HD data
 - Can't skip work based on camera
- ❖ Punishing precision jumping
 - Must feel great to control
 - Positions must be exact



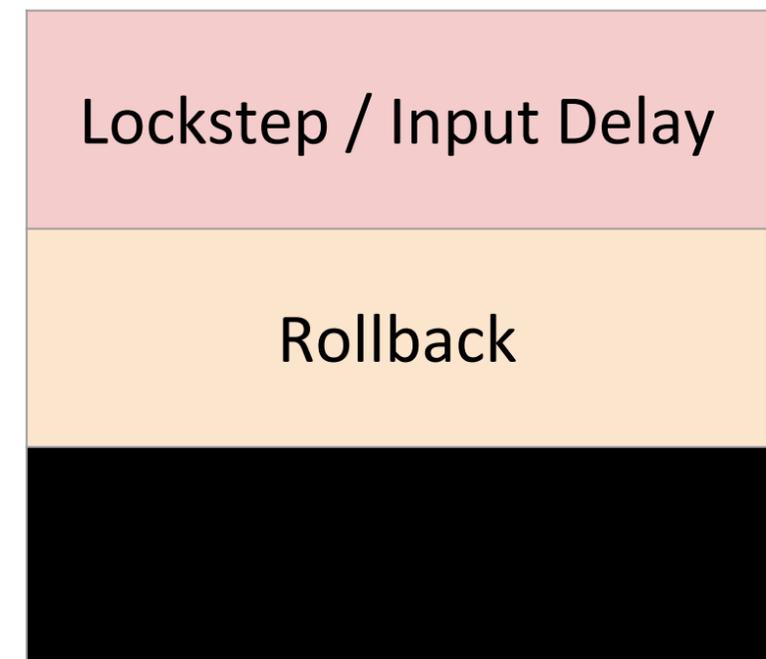
Staffing Constraints

- ❖ Roughly 8 engineering years to date
 - Including work done on all 3 platforms
 - Including all 53 patches
- ❖ 1-4 concurrent engineers
 - Engine
 - Gameplay
 - Multi-Platform
 - Online



Codebase Constraints

- ❖ Not cache-friendly or data oriented
 - Designed for faster trying out ideas and iteration
 - Performance “not bad”:
 - 3~6 frames in 16ms on PS4
 - ~12 frames in 16ms on PC
- ❖ Already deterministic:
 - Replay system already working
 - Mostly working rollback code



Rollback Netcode: The good

- Feels like local multiplayer.
- Codes like local multiplayer.
- Prevents state-cheating.

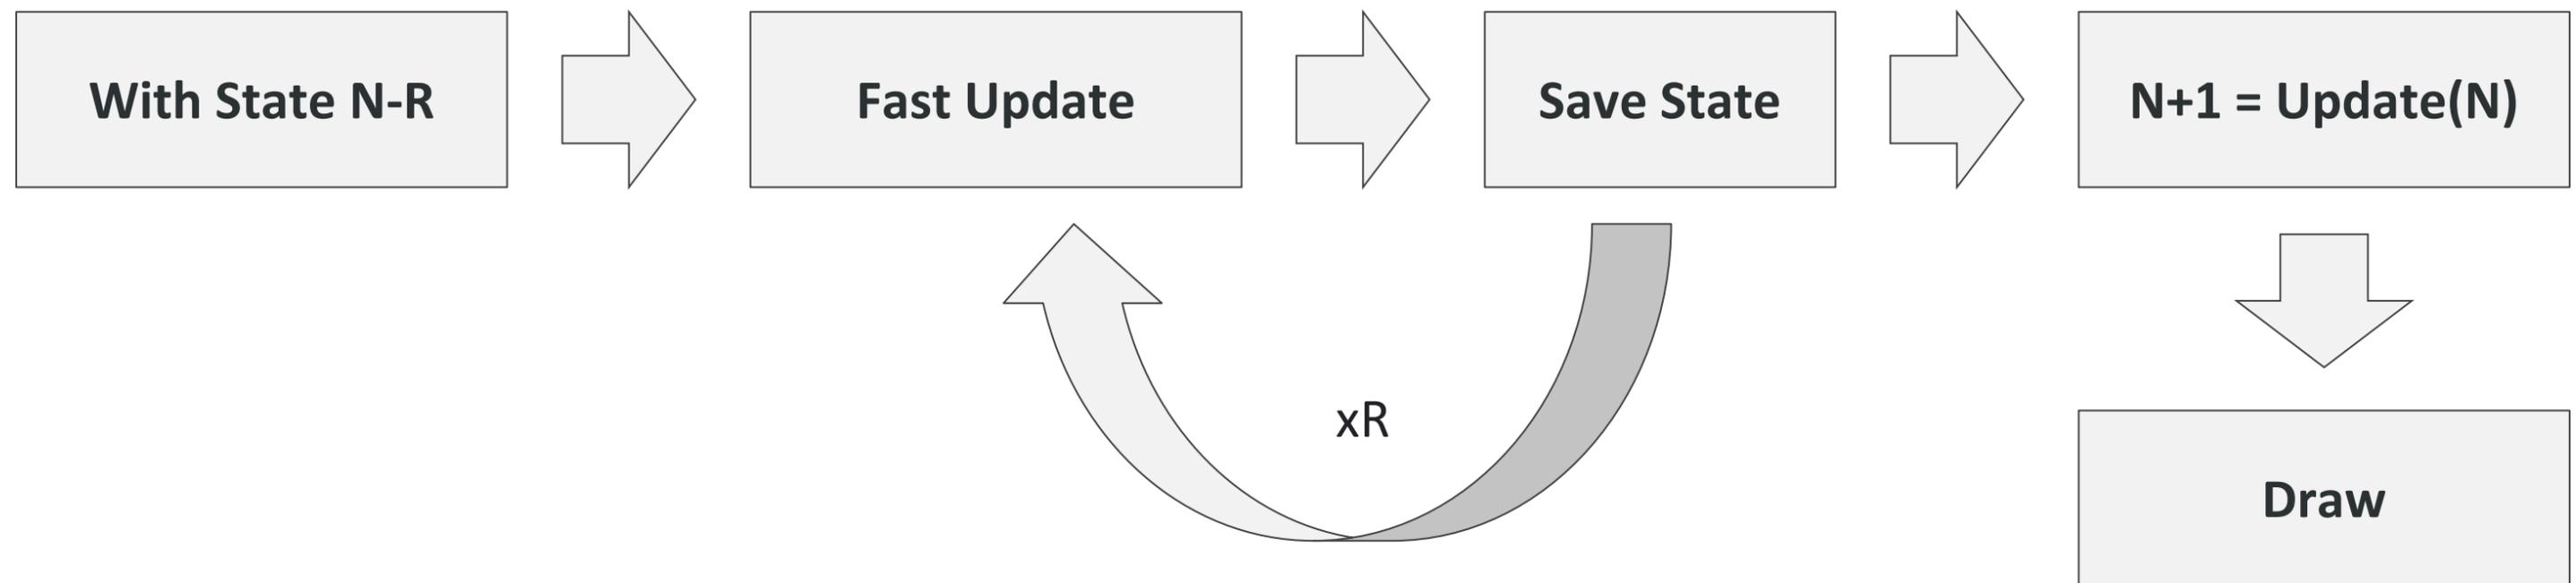


Hold!

Deterministic Rollback

Given $\text{State}^{N+1} = \text{Update}(\text{State}^N, \text{Input})$ we:

- ❖ Send inputs to other players
- ❖ When we receive inputs, we rewind and simulate with new input data

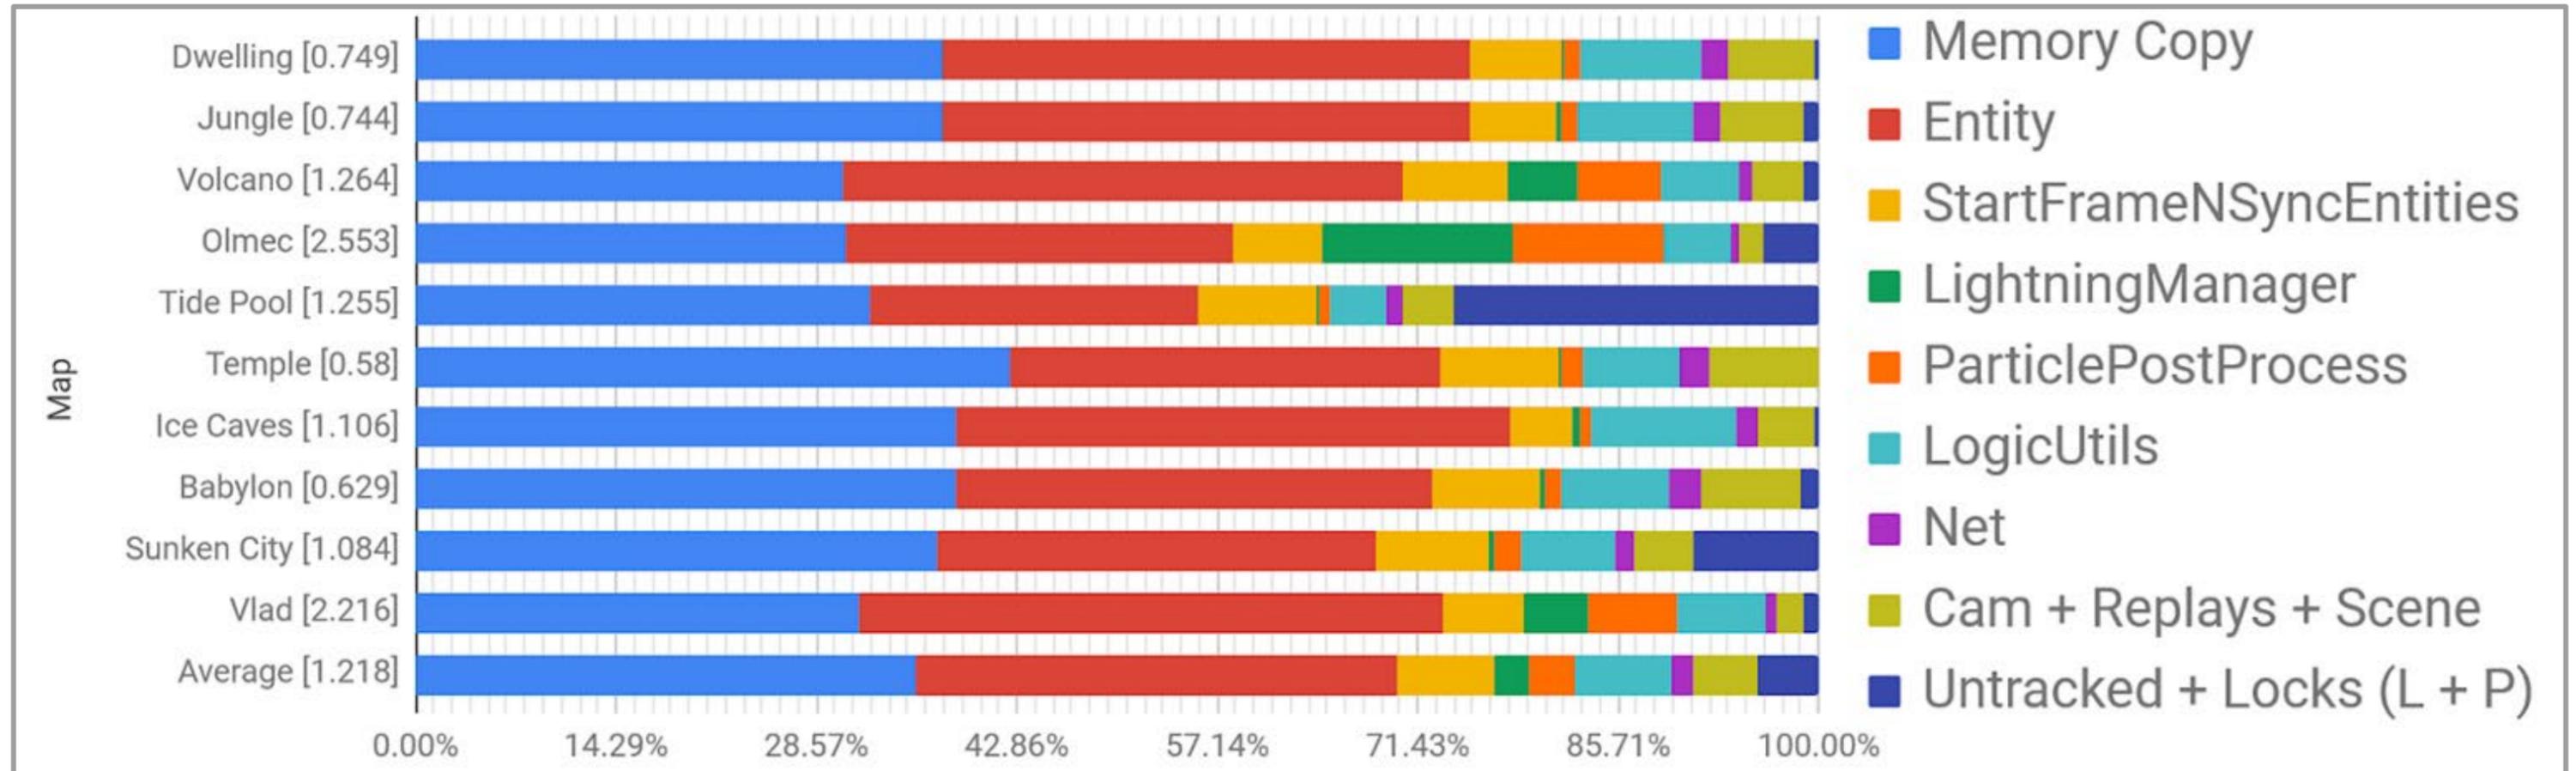


Rollback Netcode: The bad

- Performance requirements (8 updates in 16ms!!!)
- Multiple state copies need to reside in the heap
- Speed and time become entangled
- No hiding information from clients
- Game needs to be deterministic
- Harder resource management

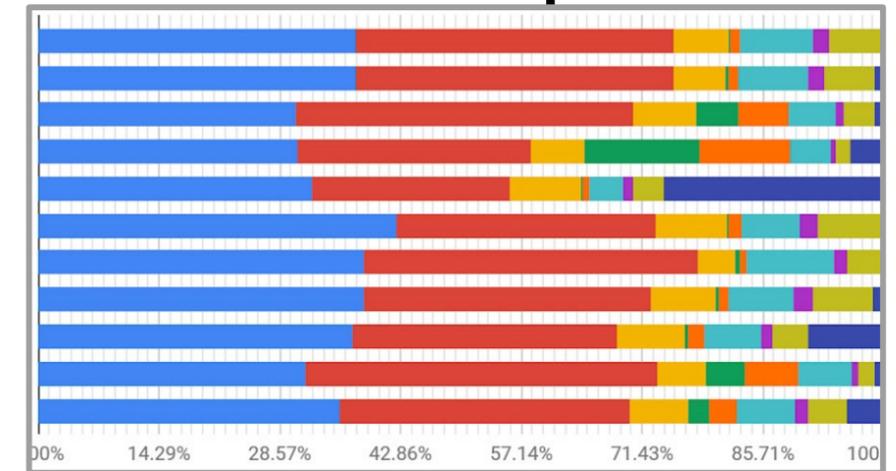


Performance Profiling



Optimization: Memory Usage

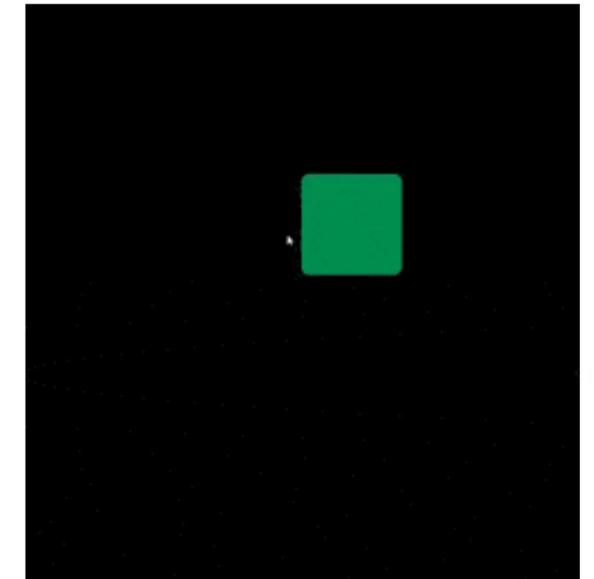
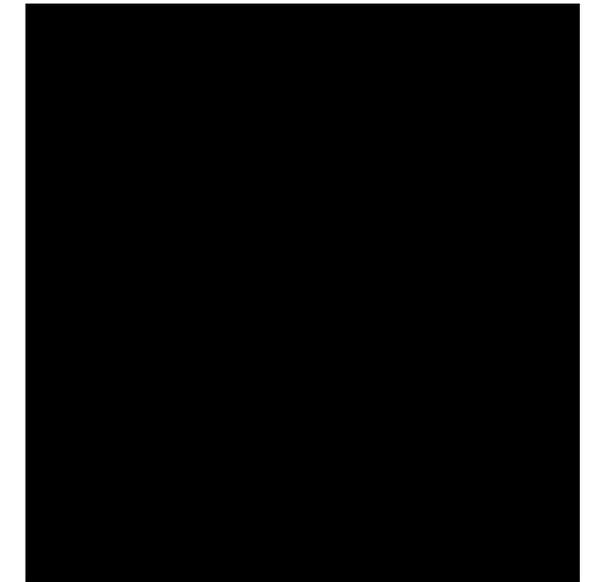
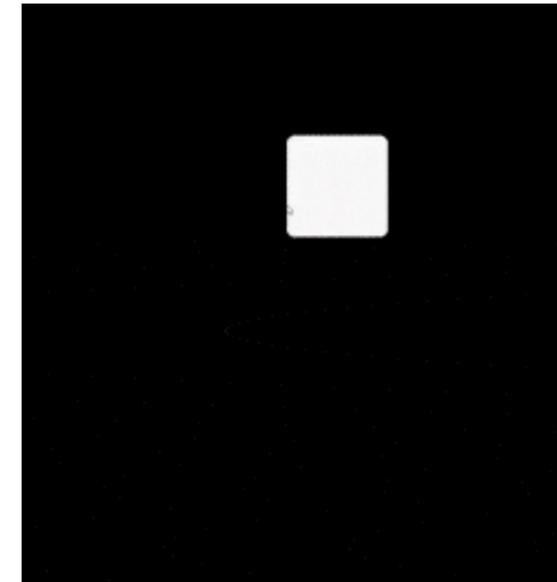
- ❖ Between 25% and 65% of time would be spent in memcopies.
- ❖ Already a focus from day 1:
 - Structures manually packed
 - Good memory allocator
 - A lot of systems outside of the deterministic state
 - 0 improvement with NT-SSE copy, movsb, threads...



Not much room for improvement?

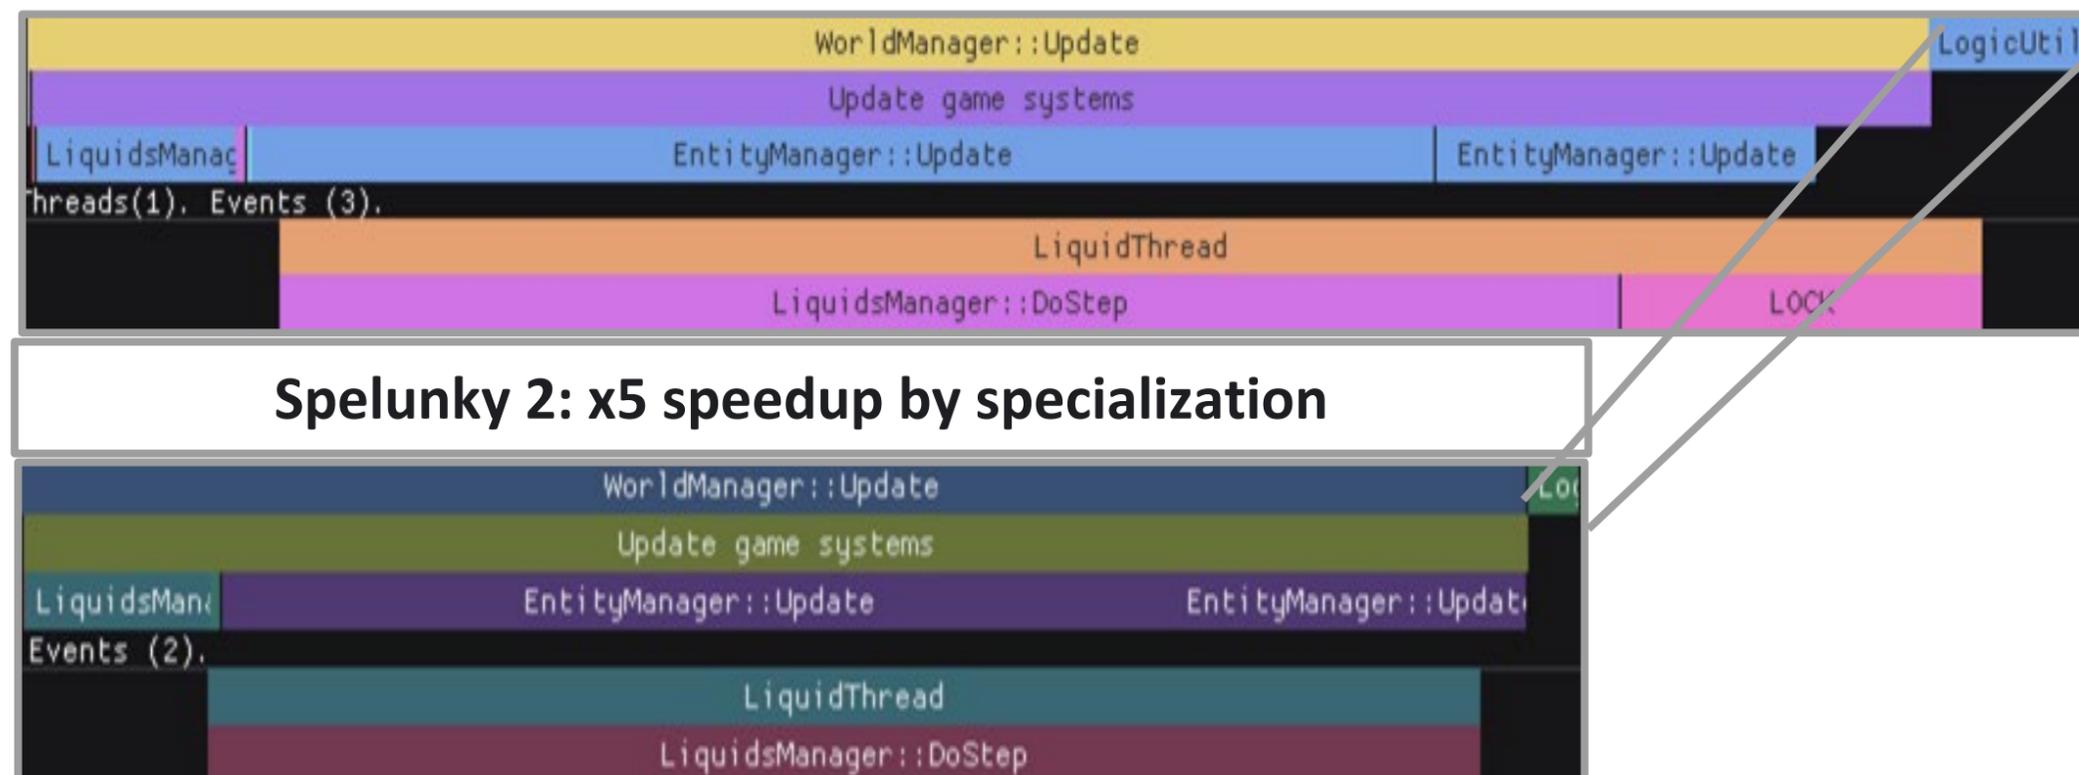
Optimization attempts: Liquids

- ❖ Liquidfun based:
 - SPH-based physics
- ❖ Liquid sim takes up to 4ms
- ❖ Used to run in a separate thread
- ❖ Sometimes 25% wait
- ❖ MPM-MLS physics
 - Granularity and feel



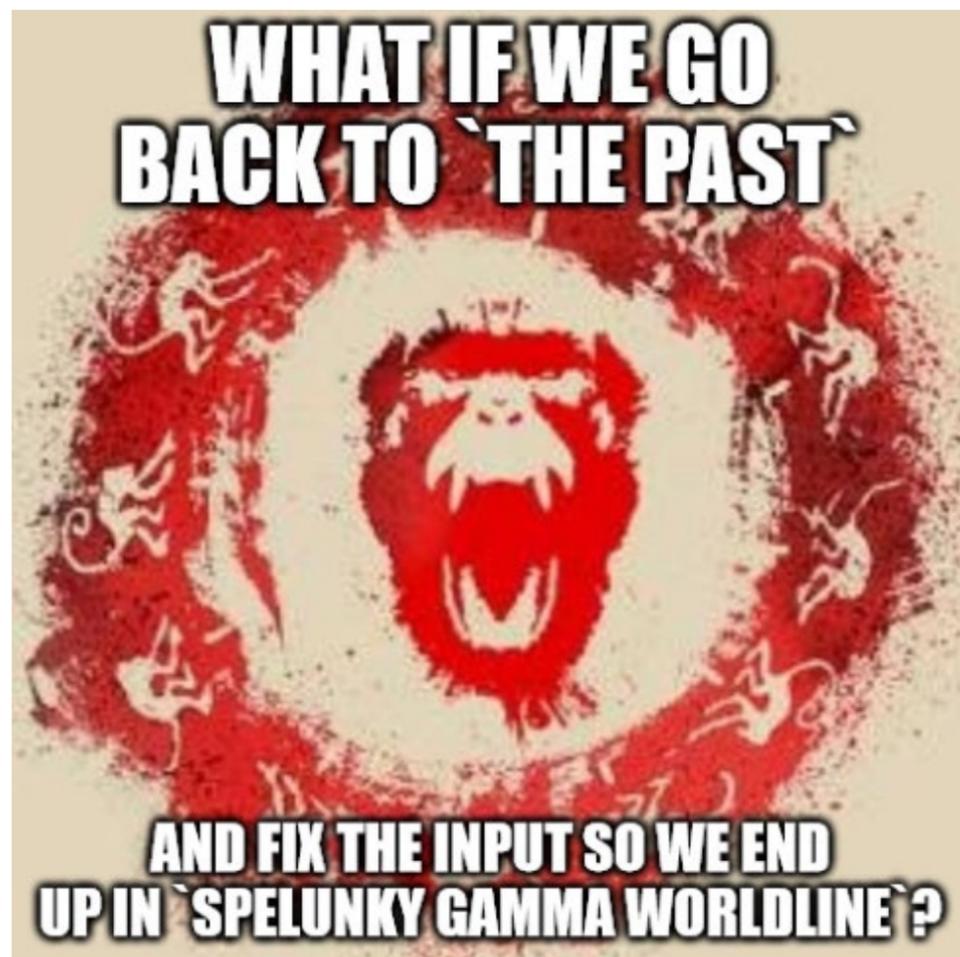
Optimization: Garbage Collection

- ❖ Cache performance: x5 improvement replacing with robin-hood integer probabilistic filter.

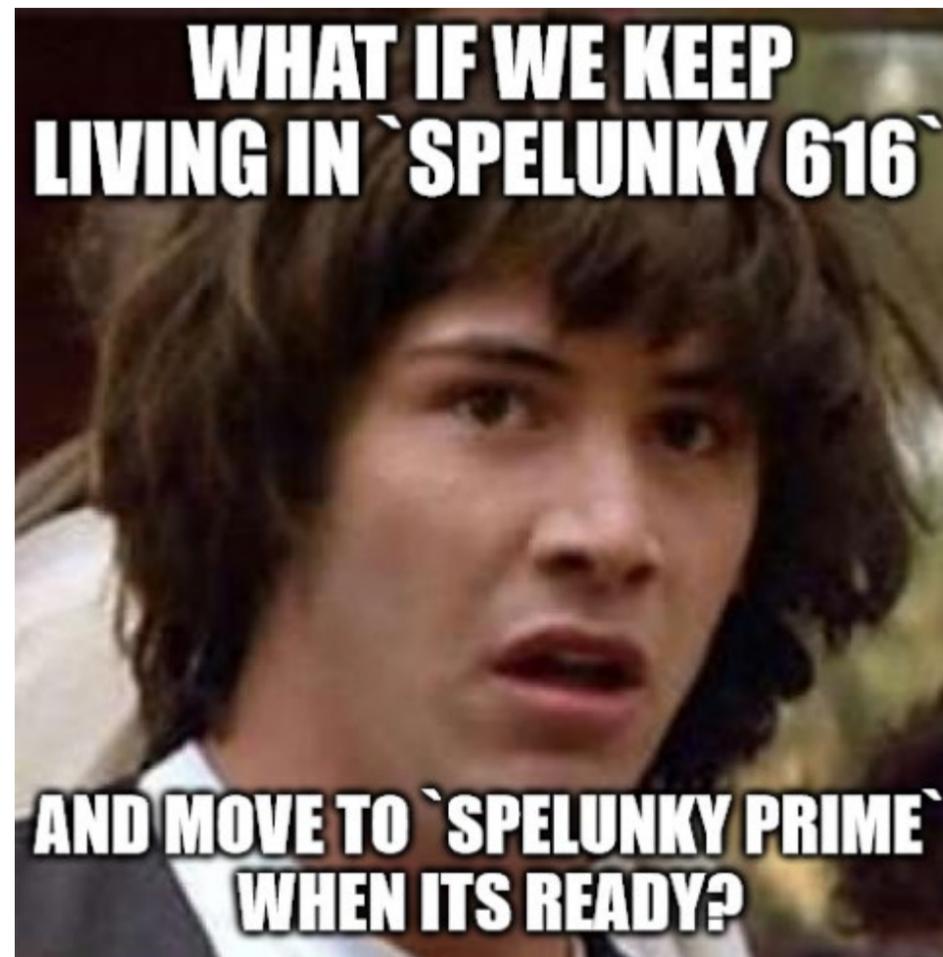


Day -403 (2019/08/09): Propagation Idea

Rollback



Propagation

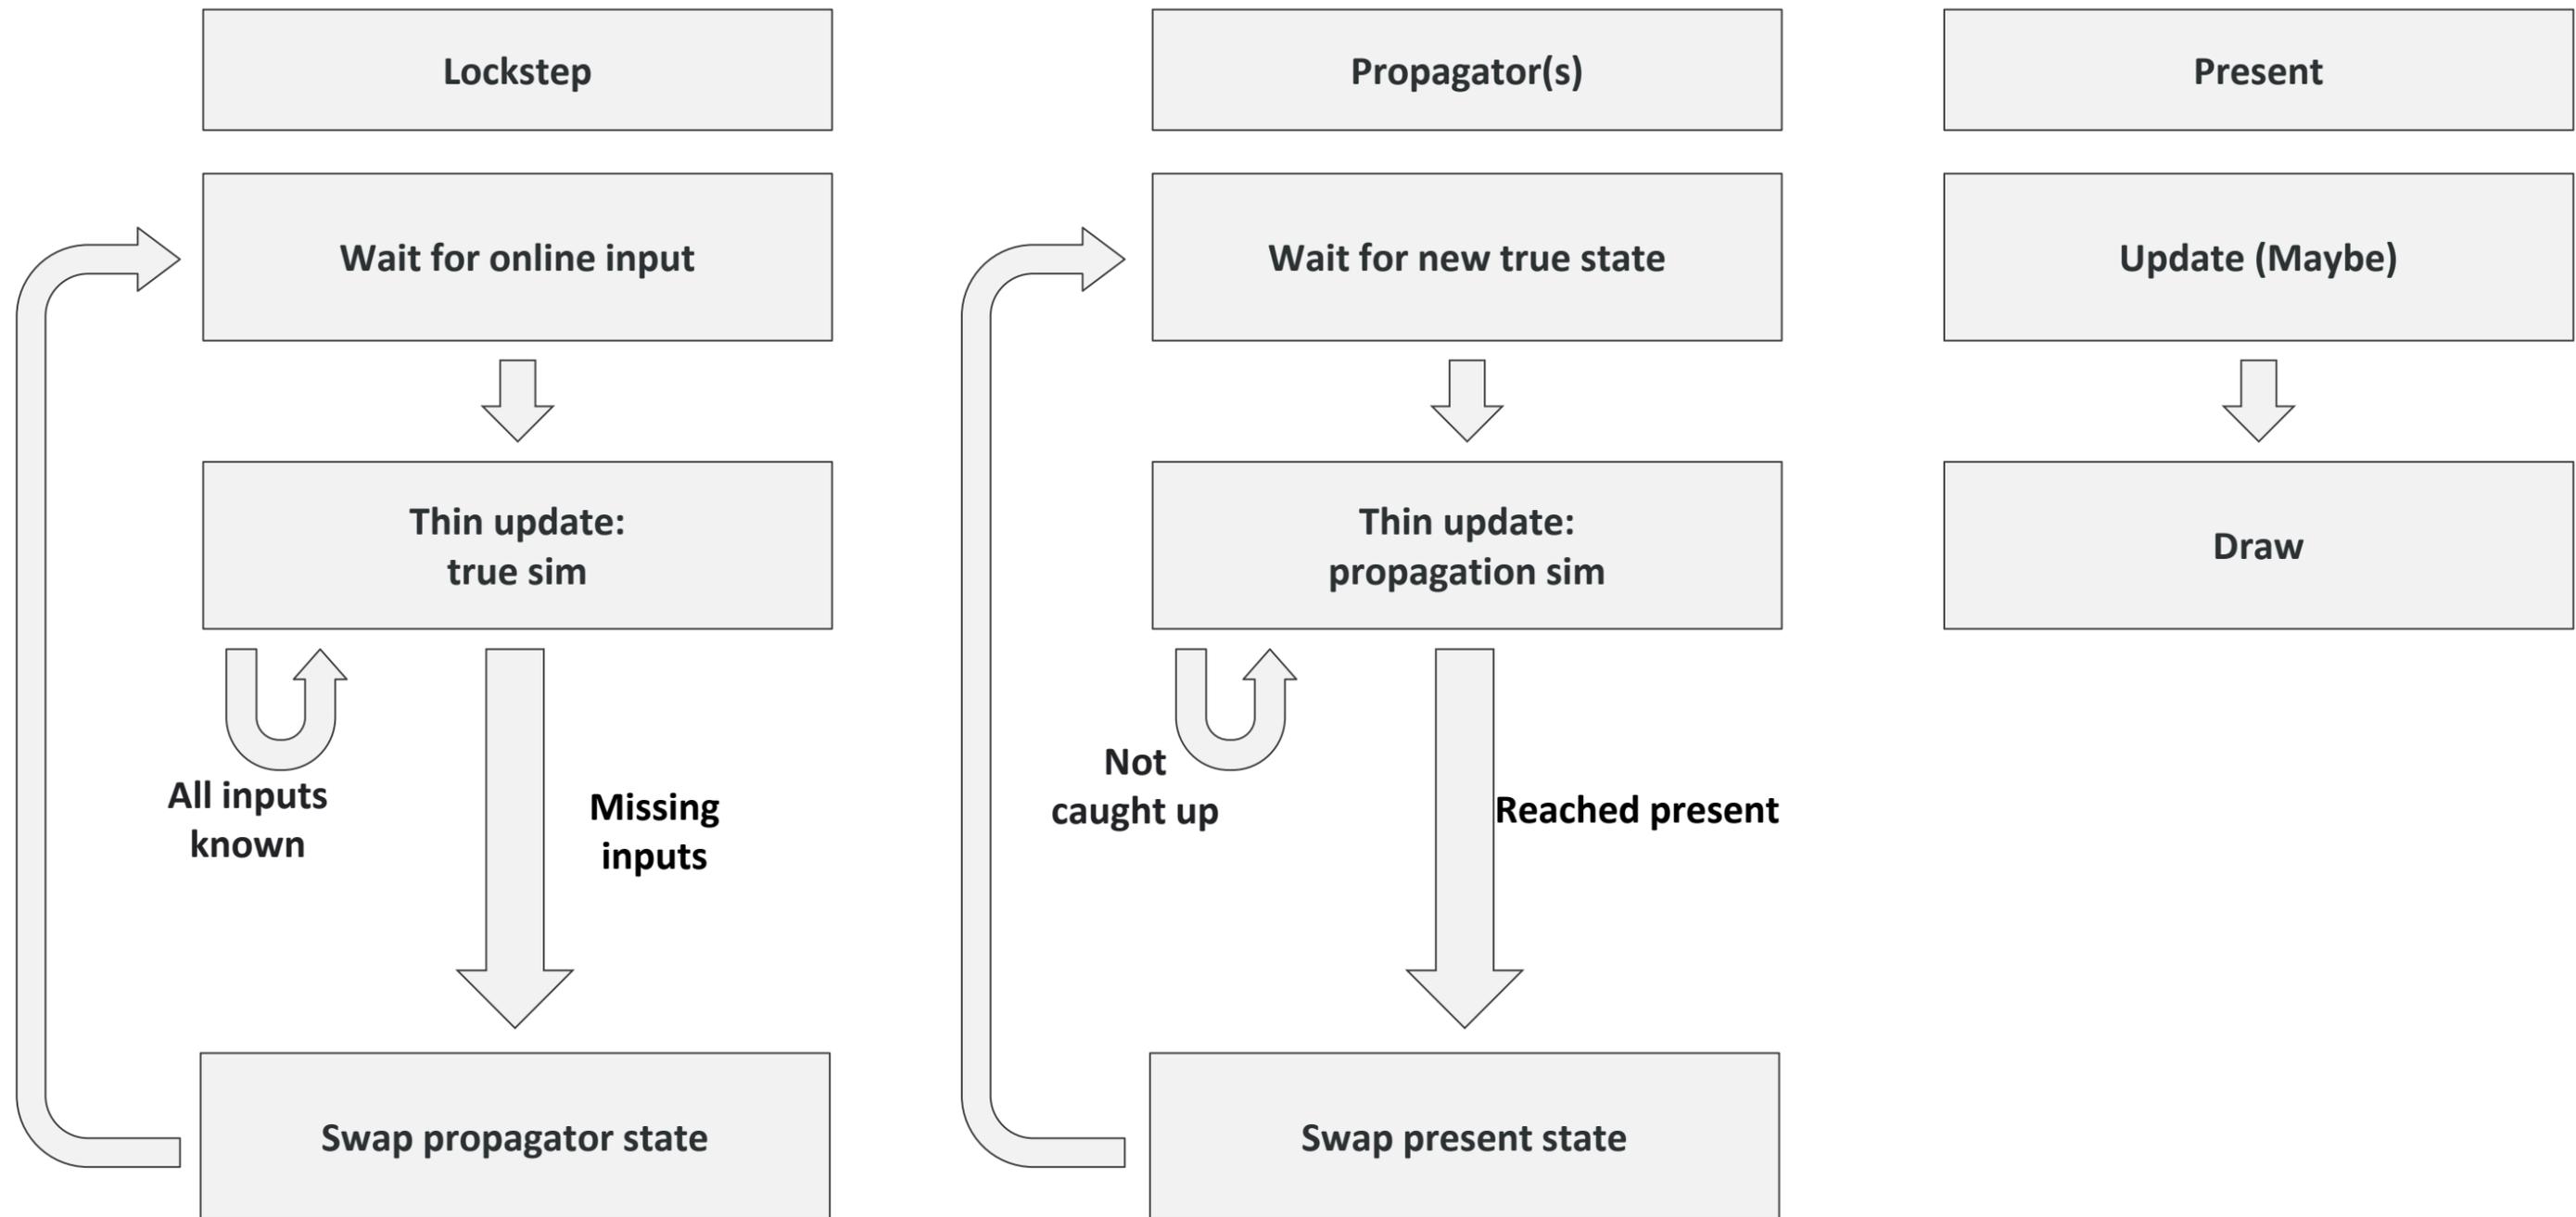


Deterministic Propagation: Implementation

- ❖ Present “no-sync” simulation: Your traditional update.
 - Uses “everyone’s” last known input.
 - Assumed to be 100% fake.
- ❖ Lockstep simulation:
 - Waits for everyone’s inputs to run a frame.
 - Assumed to be 100% true.
- ❖ Propagation simulation
 - Starts with the latest truth
 - Guesses inputs until present is reached.



Deterministic Propagation



Ideal Netcode Features	Non-deterministic	Input Delay Lockstep	Rollback	Propagation
Trivial to multithread	/	/	/	✓
Not a hassle to other devs	✓	Determinism	Determinism, Resources	Determinism, side-effects, resources
Constant netcode scale	X	✓	✓	✓
Built-in anti-state-cheat	Implementation	✓	✓	✓
Built-in anti-info-cheat	Implementation	X	X	X
Not-bad cache usage	✓	✓	X	✓
Space not linked to time	✓	✓	X	✓
Feels like local	Implementation	X	✓	✓
Ease of optimization	Implementation	1	3	2
Graceful degradation	✓	Desyncs, Pauses,	Desyncs, Full Drifting, Extrapolation*,	Desyncs, Much less drifting, Extrapolation*,
Maximum latency support	Implementation	Delay (~50ms)	Performance (~200 ms)	Arbitrary

* Extrapolation may be mitigable through “output delay” by buffering and interpolating cherry-picked entities.

And then we decided to move on

- ❖ Good idea, but:
 - Requires reworking a lot of changes.
 - Performance is not a problem yet.
 - There is a lot to be done.
- ❖ We:
 - Decide to focus on other stuff first.
 - Believe it will be required for Switch.



Good Decision: Going above 60hz

- ❖ Arbitrary frame-rate interpolation
- ❖ Interesting for pretty much everything
 - Brings the previous idea to the next level:
 - Can you imagine playing a platformer at 20 FPS?

Step:

```
pos0 = pos1;
```

```
pos1 = pos1 + vel * step;
```

Draw:

```
position = pos0 + (pos1 - pos0) * clamp(timeSinceLastStep / step, 0, 1)
```



60UPS → 144FPS

Game launches on PS4

Day 1 (2020/09/15): Rollback, the ugly

Degrades awfully, very binary performance (all determinism in general)



Comparing with Full Metal Furries

Traditional degradation: Was something off in FMF? Probably, was it noticeable? No.

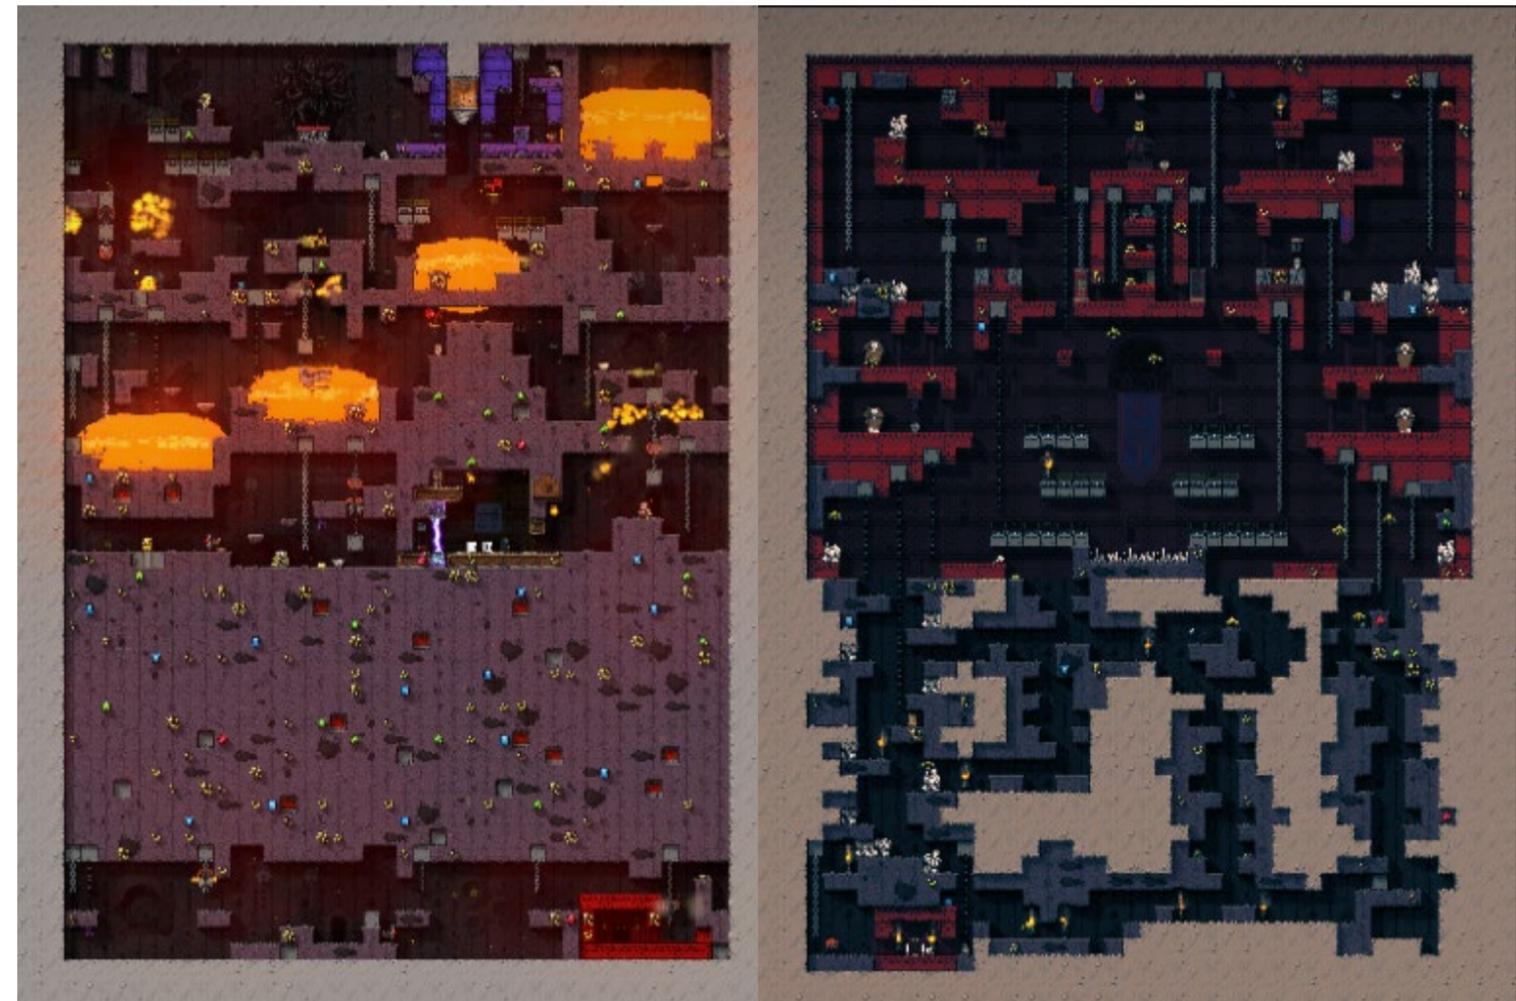


Day 2 (2020/09/16): Investigation

- ❖ Latency causing performance spiraling out of control, happens to some people even in Europe and America because the internet is not good for everyone every day.
- ❖ We investigate what could be happening by setting a dev server in Australia.
- ❖ NA & EU Servers, but no australia, brazil, africa, japan, korea...

The pink elephant in the room

- ❖ The whole huge procedural stage must update
 - Destructible terrain
 - Real-time liquids
 - Both layers
 - The usual: Collisions, pathfinding, explosions...
- ❖ Procedural gen
 - Worst case very worst.
- ❖ Not so bad on PS4 Pro.



This is one level

Day 3 (2020/09/17): Propagation (1C)

- ❖ We decide to work on propagation Netcode ASAP in order to mitigate “turn-based” Spelunky multiplayer.
- ❖ Propagation Netcode is implemented on a single thread.
 - Expected to work right if we can run 4 updates per frame
- ❖ Objective: Mitigate current worst-case scenario
 - Run at most N rollback frames + 1 fake frame
- ❖ Memory copying overhead makes it unplayable.
 - What the fruit is going on?
 - Profiling confirms memory bound on some levels.

Budget	4f
Dwelling	12
Jungle	13
Volcano	19
Olmec	33
Tide Pool	28
Temple	12
Ice Caves	23
Babylon	12
Sunken City	18
Caveman	17
Black Market	23
Vlad	31
Duat	15
Mothership	23
Hell of Ushabi	18
Palace of Pleasure	17
Tiamat	20
Eggplant World	23
Hundun	20
Leprechaun (Volcano)	23

Day 4 (2020/09/18): Virtual Magic

We start abusing memory mapping functions in order to avoid copying the state.

We use non-cached memory bus as storage for inactive game states.

Still not enough to run 4 frames in 16ms in two specific locations.



Still not perfect

16ms not enough for 4 updates + draw in some themes and levels.

PS4 <> PS4 Pro causes clock drifting, PS4 is not able to catch up and game can't slow enough.



Good Decision: Input Delay Choice

- ❖ 3 frames by default
- ❖ Reasonable for a lot of people

The screenshot shows a forum thread with the following content:

NWDD · 5m · *edited 5m*
Developer
Most games seem to decide the input delay for the player, but we wanted people to decide the tradeoffs to be done when playing online :)
We just decided to set 3 to the default as it seems to be good enough for most

Xanagear 01/04/2021
Yeah, I actually prefer the input delay to the visual stuff on stream makes for a better viewer experience I think, and input delay has never been a huge killer until I miss my dice throw

[Redacted] · 8m
I really appreciate you guys for implementing an input delay setting. A lot of fighting game developers implement GGPO or other types of rollback netcode into their games and then don't let you configure input delay, which defeats the whole point of implementing rollback when the game decides you should experience input delay anyways...

[Redacted] OP · 5m
You guys are truly doing the lord's work ❤️

NWDD 12/17/2020
it's something you can set up in options > online

Colin Northway 12/17/2020
whooaaaaaaaa

NWDD 12/17/2020
it's a tradeoff of rollback that helps mitigate how you see other players so for example, if you set it to 0 you'll have no input delay

Colin Northway 12/17/2020
omg! I had no idea! This is going to make it way way way more fun for me

The multi-instance problem

- ❖ Our update function was not a pure function
- ❖ Existing optimizations to shared data outside logic state.
- ❖ Requires surgical changes to pretty much every system:
 - HUD
 - Camera
 - Physics (Liquids)
 - Particles
 - Allocators
 - Lighting
 - Data structures
 - Sound and Music

Stats

102 files changed:

-1643

+1499

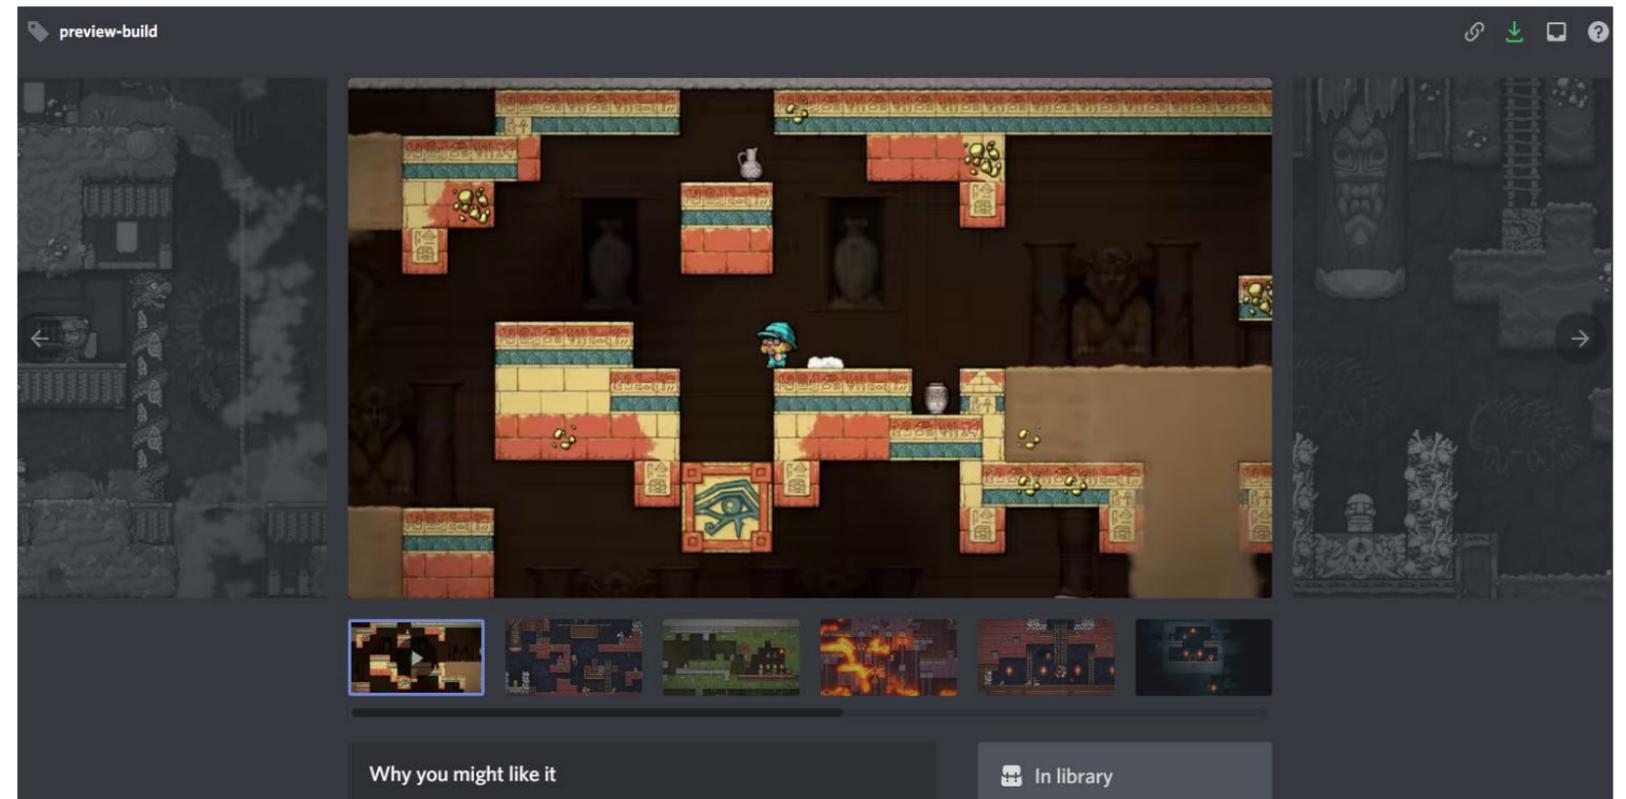


GDC[®]

GAME DEVELOPERS CONFERENCE | July 19-23, 2021 | #GDC21

Day 79 (2020/12/03): Online PC Preview

- ❖ First public version with multi-instance
- ❖ We used discord for this:
 - Already used by the community.
 - Playable depending on role.
 - `Dispatch` is great.
 - Fastest updates.
- ❖ Overall a very good decision
- ❖ Very useful feedback
- ❖ Daily patches



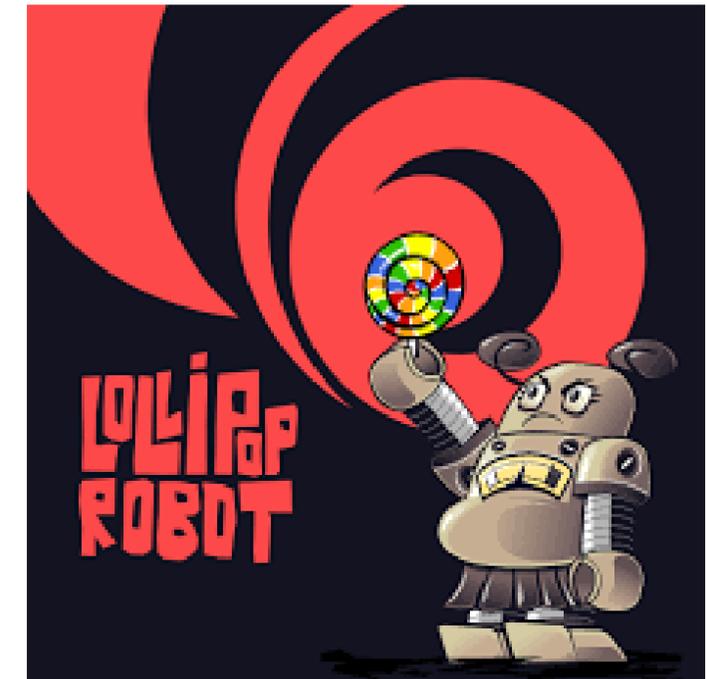
Good Decision: Preview as a reward

We didn't know how many issues would come up, not enough resources to build another server fleet to allow everyone in, so we mostly wanted to reward as many people as we could.

- ❖ Awesome people creating cool stuff (Artists / Modders / Youtubers / Streamers):
 - Ended up building a lot of hype, which was good and bad.
 - Would be nice to give them early material.
 - Wanted to thank them for their support.
 - Compare streams side by side.
- ❖ Daily users of “existing online solutions” (Steam remote play / Parsec):
 - Improve their experience.
 - Compare vs streaming.
- ❖ Problem setups (low-end pc players, not from NA/EU or with bad internet):
 - Validate worst case scenarios.
- ❖ Very active multi-community members:
 - To improve transparency & communication

Good Decision: Adding online ██████████

- ❖ Ghost bug:
 - QA spent a lot of time investigating Olmec.
 - Only “soft-desyncs” happened.
- ❖ Happens a random preview afternoon:
 - Two people, one streamer ([HectiqueX](#))
- ❖ Special build, 6 hours later fixed!



The 'Olmec' desync



Not Deterministic

```
if (screenShakeEnabled) {  
    → screenshake(rand(screenShakeStrength));  
}
```



```
auto value = rand(screenShakeStrength);  
if (screenShakeEnabled) {  
    → screenshake(value);  
}
```



Not Deterministic

```
int rand(int min, int maxInclusive) {  
    if (maxInclusive <= min) return min;  
    //generate the actual random number;  
}
```



Good Decision: Adding online replays

	Replays	No Replays
Reproduce issue	O(1), automatic	variable, manual
Iteration times	gets faster, uids	variable, long
Fix Verification	almost always	only if fully reproducible
Bug reporting	fast, happen once	slow, find steps
Race Conditions	frame rerolling	hell
Release as debug	✓	X



Bad Decision: Online first on PS4

- ❖ Wouldn't have released with online on PS4 if we knew.
- ❖ Had a lot of backlash on PC
- ❖ PS4 Feedback:
 - Very few feedback
 - Not accurate
 - Dev. Limits:
 - Replays
 - Permission



Server: Growing up our fleet

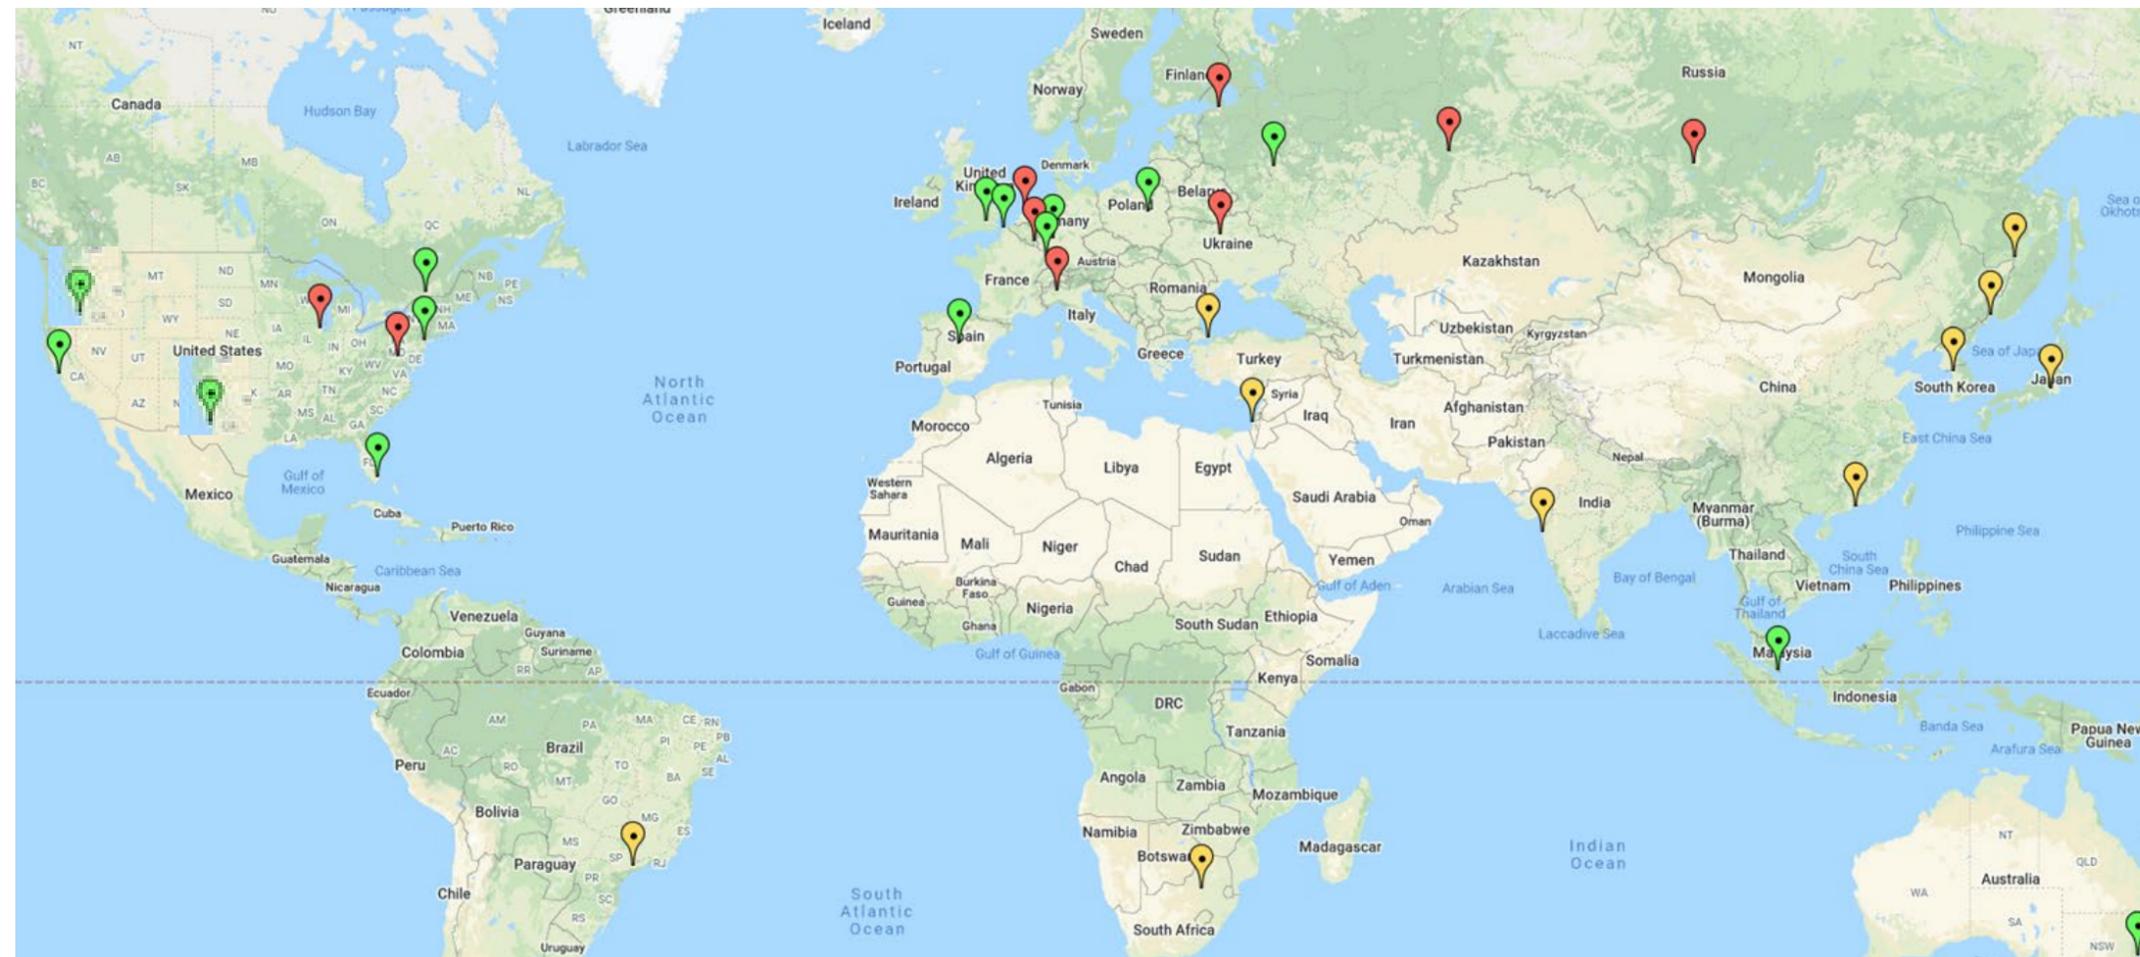
“The edge is a dumpster” - Glenn Fiedler

1. Bad experience with FMF / HW.
2. Chicago server
3. Extra servers (14)
4. Final fleet

Bandwidth important.

Many providers/services:

Vultr, DO, GCore, Hetzner,
BunnyCDN, OVH, linode...



Server: Tech

- ❖ C/C++ Hot-reloading dynamic library
 - Around 20% disconnect on update
 - Only one interruption of service
- ❖ FreeBSD:
 - 4 processes (including server)
 - low maintenance
- ❖ Simple protocols:
 - Mostly fixed buffers
 - Fully unreliable even matchmaking
 - Desired state
 - Download the whole daily leaderboard (HTTPS / CDN)



Mixed Decision: Hosted instead of P2P

- ❖ The “Bad”:
 - Waking up with nightmares →
 - We still plan to add self-hosted solution.
 - The path of most resistance.
- ❖ Overall Better experience:
 - More guarantees regarding quality of service.
 - Explicit timeline for latency compensation.
 - +10% NAT players.
 - Cross-play.



Day 99 (2020/12/13): Netcode Beta

- ❖ Not superstitious but maybe we should've waited for the 14th*?
- ❖ Google went "down":
 - Discord rich presence broken
 - Took multiple builds and conversations:
 - Published build was corrupted:
 - Steam invites
 - Desyncs
- ❖ Too much hype:
 - We had to make our discord pay to chat.
- ❖ A neighbour to one of our most used servers gets DDOSed

Madame, Monsieur,

Nous venons de détecter une attaque sur l'adresse IP [REDACTED].

Afin de protéger votre infrastructure, nous avons aspiré votre trafic sur notre infrastructure de mitigation.

Toute l'attaque sera ainsi filtrée par notre infrastructure, et seul le trafic légitime arrivera jusqu'à vos serveurs.

A la fin de l'attaque, votre infrastructure sera immédiatement retirée de la mitigation.

*Actually it was the 15th, but the story improves so much with this change

Day 110 (2020/12/23): PS4 Propagation

- ❖ PS4 feedback confirms improvement
- ❖ This feedback was the best xmas present

Been playing with a friend the past few days.

Not had a single case of visual lag or glitching of any kind whatsoever. I'm in UK and friend is in Italy, it's been absolutely flawless.

Welp i just tried online for the first time because i decided to take a break from studies, safe to say it was really fun

online IS fun

Online is vvvvv fun

Thanks you so much, online works great now, meand my gf are having a blast with this lovely game!

↑ 2 ↓ Reply Give Award Share Report Save

the online is awesome

i found a random dude and it was so fun

I played an online match of Spelunky 2, awesome sauce

has anyone been able to play with friends yet

we got tons of randoms before and it was actually kind of fun

yes and it was really fun

Having a blast playing with randoms

Today at 10:22 PM

This online has brought a new depth to this game, it is genuinely incredible I had a riot playing earlier today thank you spunky devs

👍 7 ❤️ 3

Hey I wanted to apologize for being kind of a dick about PS4 online. Tonight my buddy and I played the update and had the best time. Favorite game ever.

1

↻

❤️ 8

↑

Other than invites not working you guys did a really good job with the online! It's almost lagless and it works really well. Thanks for making my GOTY even more...GOTYish.

Replying to @mossmouth @BlitWorks and 3 others

Thank you so much for fixing multiplayer Derek!!!!

I can invite my friend but it just loads then kicks us out, pretty awesome 😊

Randos are pretty fun tho

online is great! you did a good job 😊

Derek Yu @mossmouth · Dec 24, 2020

It's been a frustrating year, so I totally get it. Much appreciated, though. Glad the update worked out for you!

1

↻

❤️ 11

↑

Today at 10:35 PM

The online is incredible just needs text chat

yea its super fun

I had a fun match with randos

Online is fun period just dont expect serious play

playing with randos is so much fun

Online has been drastically improved, and I ended up playing for a few hours today with virtually no lag. Great experience, won a couple times. Thanks

👍 10

I don't know if they actually changed anything but online PS4 feels amazing now.

Online is fun

Good Decision: Online Multiplayer (for us)

- ❖ Think if it fits your game:
 - How long will it take to implement?
 - What could you implement in that time?
- ❖ A side feature, designed from day 1:
 - For some people it's the "only" mode
 - Barrage of requests will never end.
 - The arcade view for online:
 - Not trying to optimize retention
 - No need for accounts
 - Just as if playing in the same machine
 - The world is IAP



Game analytics

❖ Measure server impact:

- Merge a few nearby servers in Europe and America (based on matchmaking data)
- Multiple spikes in Lobbies
 - After adding a Seoul server:
 - Corresponding to times when [Nokduro](#) played



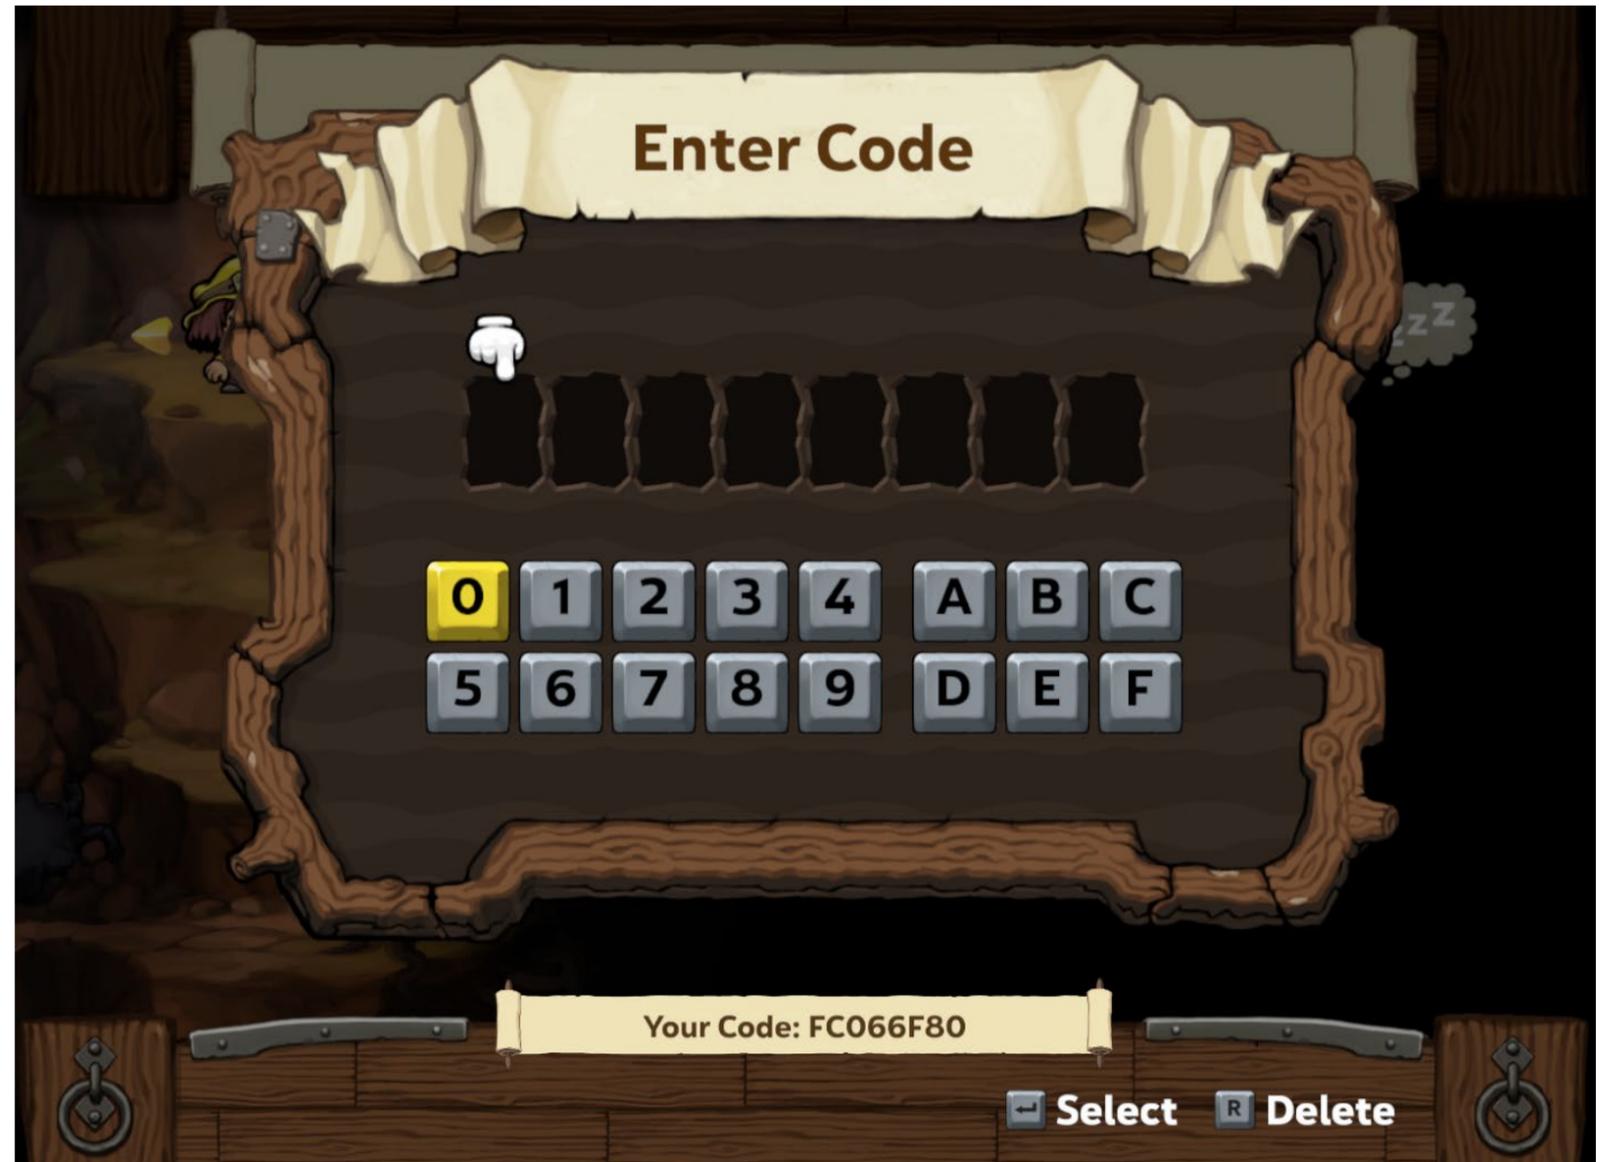
Game analytics

- ❖ 200~700 CCU 24/7: Three people complain about lag since “last week”.
- ❖ Scheduling updates.
 - In case shit happens.
- ❖ Satisfying to look at.
- ❖ Bandwidth costs:
 - AWS EC: \$1000/mo + servers
 - Others: \$150/mo (server incl.)
 - Including Amazon lightsail
- ❖ Fleet costs:
 - Junior programmer 1yr
 - 10yrs. of Server fleet



Day 111 (2020/12/24): Invite Codes

- ❖ Good Decision
- ❖ Very useful for communities:
 - Great for streamers:
 - Don't show by default
 - Other platforms and chats
 - Rich presence disabled
- ❖ Will be useful for cross-play
- ❖ Copied / Pasted from clipboard



General problems

- ❖ Desync:
 - Something not deterministic enough.
 - Ultra-rare.
- ❖ Extrapolation working as intended:
 - Happens every now and then.
 - Players are not used to this.
- ❖ Clock drifting:
 - Central authority provides timeline.
 - Shorter/Longer frames: 57 to 63 fps.
 - Some players will feel the stutter:
 - Regret: Interpolation wasn't designed to merge frames.
- ❖ **Ambiguous feedback**



Biggest Problem: Ambiguous feedback

Expectation:

- ❖ The game desynced and we ended in a different 4-1

Reality:

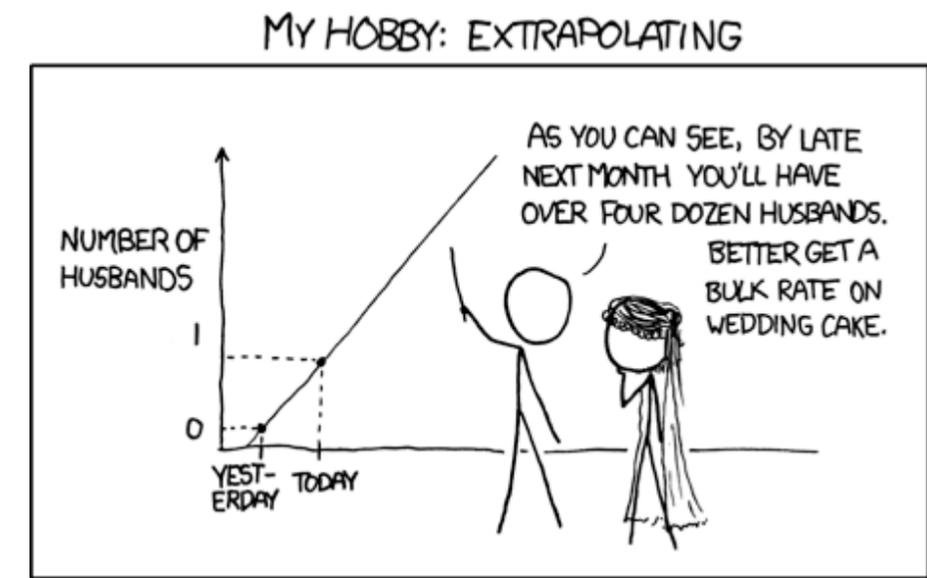
- ❖ “The game desynced” → “The game crashed”
- ❖ “The game desynced” → “The game runs in slow-mo”
- ❖ “The game desynced” → “The game stutters too much”
- ❖ “The game desynced” → “I disconnected from the game”
- ❖ “The game desynced” → “The rollback worked as intended and I saw a misprediction being fixed”
- ❖ “The game desynced” → “My machine doesn’t meet the requirements for rollback and clocks drift”
- ❖ “The game desynced” → “The game doesn’t feel smooth because of input delay”



Works in pretty much any direction.

Open Challenges: Extrapolation

- ❖ Extrapolation doesn't look good
- ❖ It's something built in in rollback & propagation.
- ❖ Possible solutions:
 - We already have input delay, why not to add Output Delay?
 - Requires additional semantics
 - Delays outputs (of the update function)
 - Buffer last N states of remote clients
 - Paint (Now-OutputDelay)
 - Becomes a cheap way to simulate interpolation



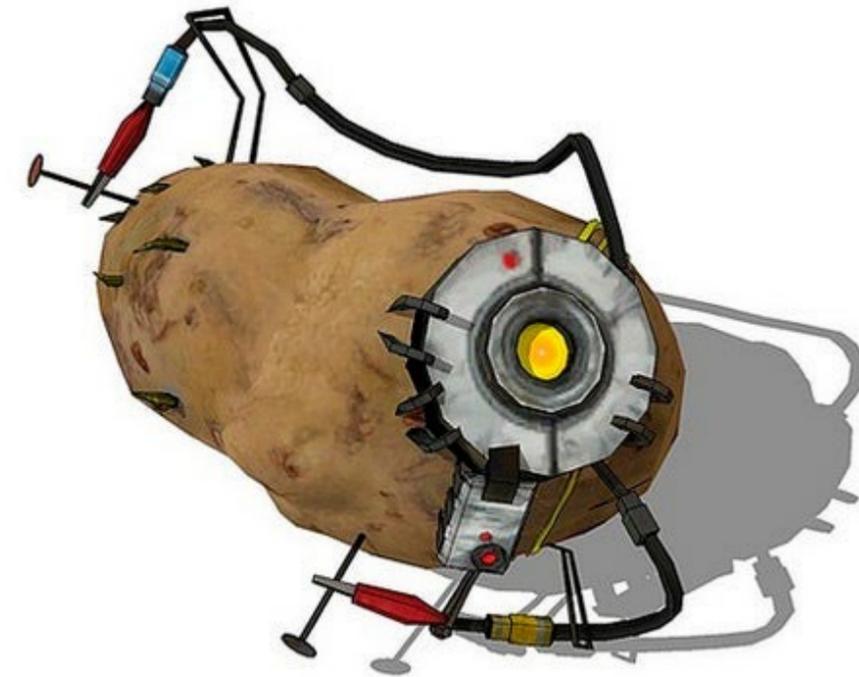
Open Challenges: Matchmaking

- ❖ Most matchmaking info won't apply
 - Online player-base not big enough
- ❖ Regional matchmaking:
 - Great experience
 - Playerbase fragmentation:
 - No-one to play with off-peak
 - Mitigated an awesome community
- ❖ Non-regional matchmaking:
 - Reasonable matchmaking times
 - Worse play experience



Open Challenges: Below Minimum Req.

- ❖ Detecting this reliably
- ❖ 32/36 hells confirmed as of Early 2021
- ❖ Handling this in Precision platforming?
 - Slowing down more?
 - Pausing?
 - Making this public?
- ❖ Not only low end: Great Intel Laptops
 - One “fast” core
 - Very short burst
 - Battery vs Plugged
 - Thermal throttling
 - Windows updates



Unexpected

Around 8% of people have micro-cuts of more than 3 seconds:

- ❖ Wi-Fi background scans
- ❖ Cheap routers
- ❖ ISP cuts?

Cross-Play requirements:

- ❖ Designed for IAP games?
- ❖ Designed for GaaS?

Chat support:

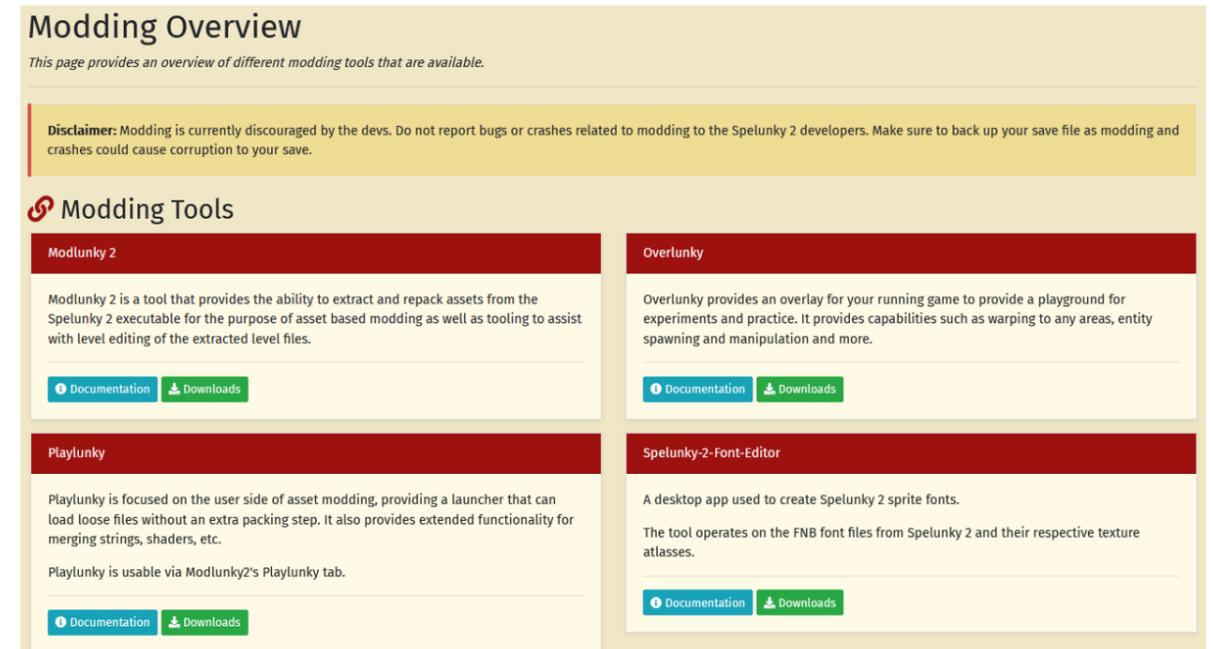
- ❖ Legislation makes sense but sucks

And... Some people can't connect to the servers !?



Open Challenges: Modding

- ❖ No official API or support yet
 - A lot of reverse engineering work
 - 800+ Mods
- ❖ Regular reports of bugs we fixed:
 - Online made evident the source.
- ❖ Out of three hard desyncs we analyzed:
 - 2 were confirmed to be caused by modding.
 - 1 wasn't sure if someone was modding.
 - Very time-consuming process.
 - Ideal world: “automatic match recording”
 - Too rare to be worth it: Arena, Cross-Play and Disconnects



Modding Overview
This page provides an overview of different modding tools that are available.

Disclaimer: Modding is currently discouraged by the devs. Do not report bugs or crashes related to modding to the Spelunky 2 developers. Make sure to back up your save file as modding and crashes could cause corruption to your save.

Modding Tools

- Modlunky 2**
Modlunky 2 is a tool that provides the ability to extract and repack assets from the Spelunky 2 executable for the purpose of asset based modding as well as tooling to assist with level editing of the extracted level files.
Documentation Downloads
- Overlunky**
Overlunky provides an overlay for your running game to provide a playground for experiments and practice. It provides capabilities such as warping to any areas, entity spawning and manipulation and more.
Documentation Downloads
- Playlunky**
Playlunky is focused on the user side of asset modding, providing a launcher that can load loose files without an extra packing step. It also provides extended functionality for merging strings, shaders, etc.
Playlunky is usable via Modlunky2's Playlunky tab.
Documentation Downloads
- Spelunky-2-Font-Editor**
A desktop app used to create Spelunky 2 sprite fonts.
The tool operates on the FNB font files from Spelunky 2 and their respective texture atlases.
Documentation Downloads

Community modding tools & mods at spelunky.fyi

Guillermo NWDD

Linkedin: [in/xnwdd](https://www.linkedin.com/in/xnwdd)

Twitter: [@xnwdd](https://twitter.com/xnwdd)

BlitWorker

@BlitWorks

Undergrad thesis tutor @UOCuniversitat

Game console prog. @ENTI_BCN/@UniBarcelona



Universitat
Oberta
de Catalunya

