

Internet, Video, Composition: Challenges and Opportunities

Leslie Daigle, *Thinking Cat Enterprises LLC*

Glenn Deen, *Comcast-NBCUniversal*

TL;DR

- Taken as monolithic chunks of data, movies and other video material present increasing challenges to the operational Internet. At the same time, digital development and delivery of video content relieves us from the strictly linear format.
- This talk will explore the possibility of considering video as a composite data structure, opening up opportunities for more successful handling of the data stream required to deliver it.

CONTEXT

Who remembers the “World Wide *Wait*”?

- Early 1990’s – text based World Wide Web launched, and revolutionized the world!
- Late 1990’s – image support introduced, fancy graphics appeared, and everything... got... SLOOOOWWW.
 - People would sit and wait for a web page to load
 - Incrementally
 - Unevenly
 - Eventually

... no waiting!

- And then broadband happened for most of us and things were good again!

PHEW!

That's not going to happen for video

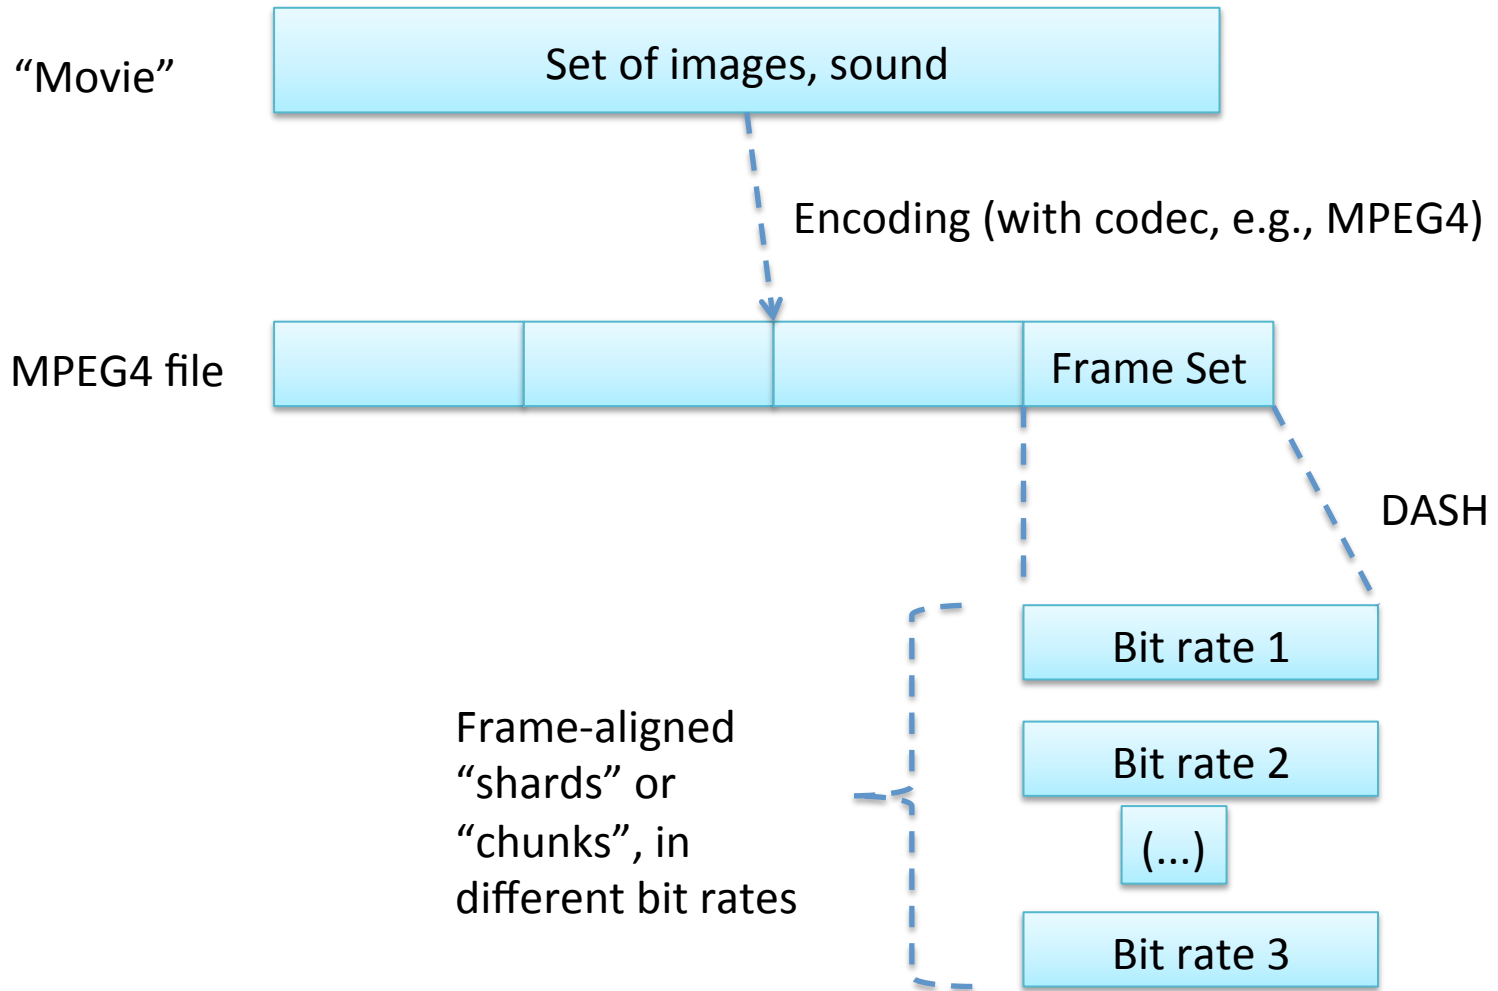
- We think of videos as being like photos only (much) bigger files, except
 - Extreme size
 - megabits per second, and gigebytes per hour of video
 - Sensitivity
 - when streamed for viewing, sensitivity to latency and dropped packets
 - Shear volume
 - The Cisco Visual Networking Index projects that by 2019 there will be nearly a million minutes of video per second transported by the Internet, a making up 80-90% of all IP traffic
 - How it met the Internet
 - “In contrast to the extensive end to end expert knowledge and engineering done to create the Web and email, Internet video has largely been an evolved cobbling and adaption exercise done by engineers with their focus on a few, or one, particular aspect or problem at a time, and little interaction between other parts of the Internet video ecosystem”. (draft-deen-daigle-ggie-00.txt)
- We can't grow bandwidth fast enough to solve the problem

Just looking at 2015

- Number of creators grew incredibly fast
- Number of viewers grew even faster
- New resolutions & the need for more bandwidth
 - SD->HD was the big 2014 transition
 - HD->4K is here now
 - 4K-8K is around the corner
- New devices are capturing content at higher resolutions
 - Phones are now capturing at 4K
- New platforms are adding new ways to distribute
- 1::Many live video streaming (e.g. Periscope)
- 0 (Zero) New magic new technologies that make it easy to do

WHAT IS VIDEO AND HOW DO WE HANDLE IT?

“A Movie”



“Movie” production is complex

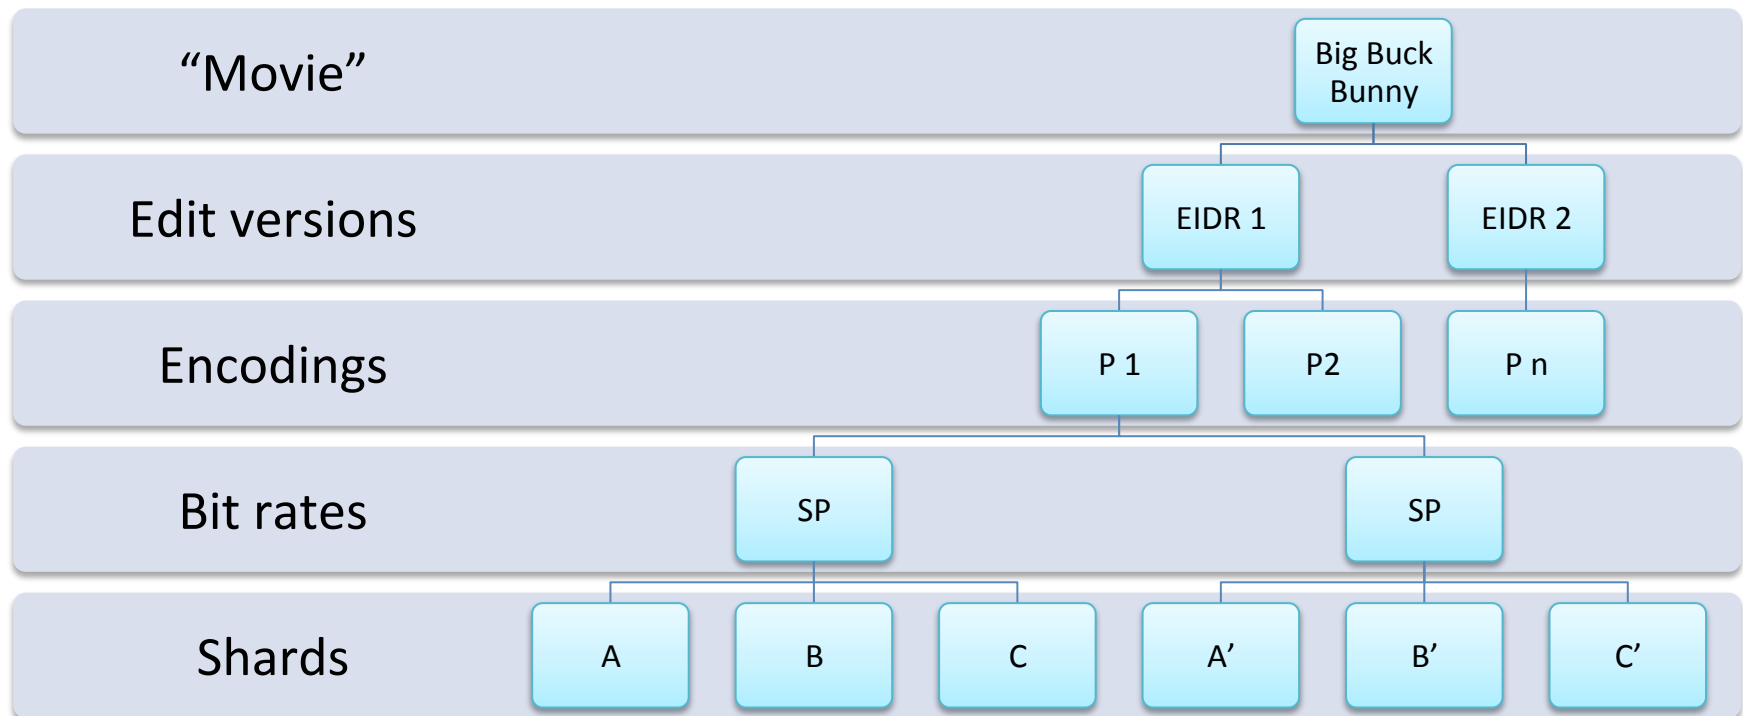
- Takes a lot of resources to produce multiple formats
- No single monolithic version of a movie
 - Not uncommon for a movie to have 80 packagings
 - Different edits (Han firing the first shot or not)
 - Slightly different content (depending on market)
 - Different bit rates, codecs
 - Audio tracks – alternative languages & encodings
- Produces lots of data that’s semantically similar

“Movie” distribution is complex

- Trying to reduce the issues of bandwidth by reducing the distance and number of times the data must be fetched
 - CDNs
 - Caching
- Takes a lot of resources and planning to get the video “close” to where it will be consumed
- Caches and CDNs are hard!
- And, we’re not even talking about user-generated content or contributions to content
 - E.g., multiple vantage points from sporting events

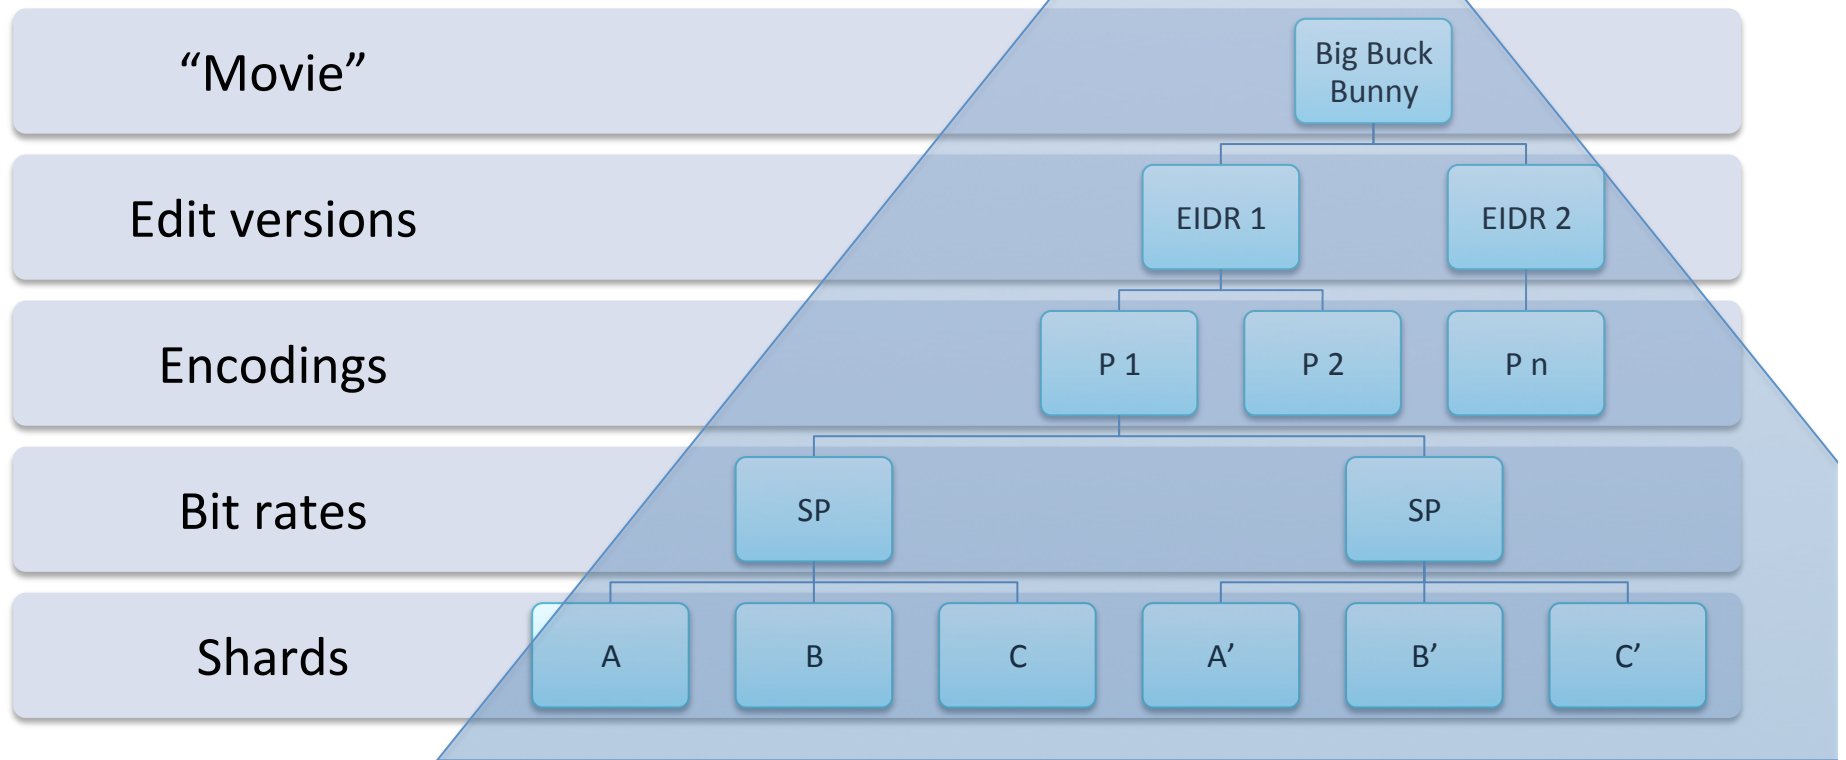
LOOKING AT THE PIECES PARTS

A Digital Content Tree

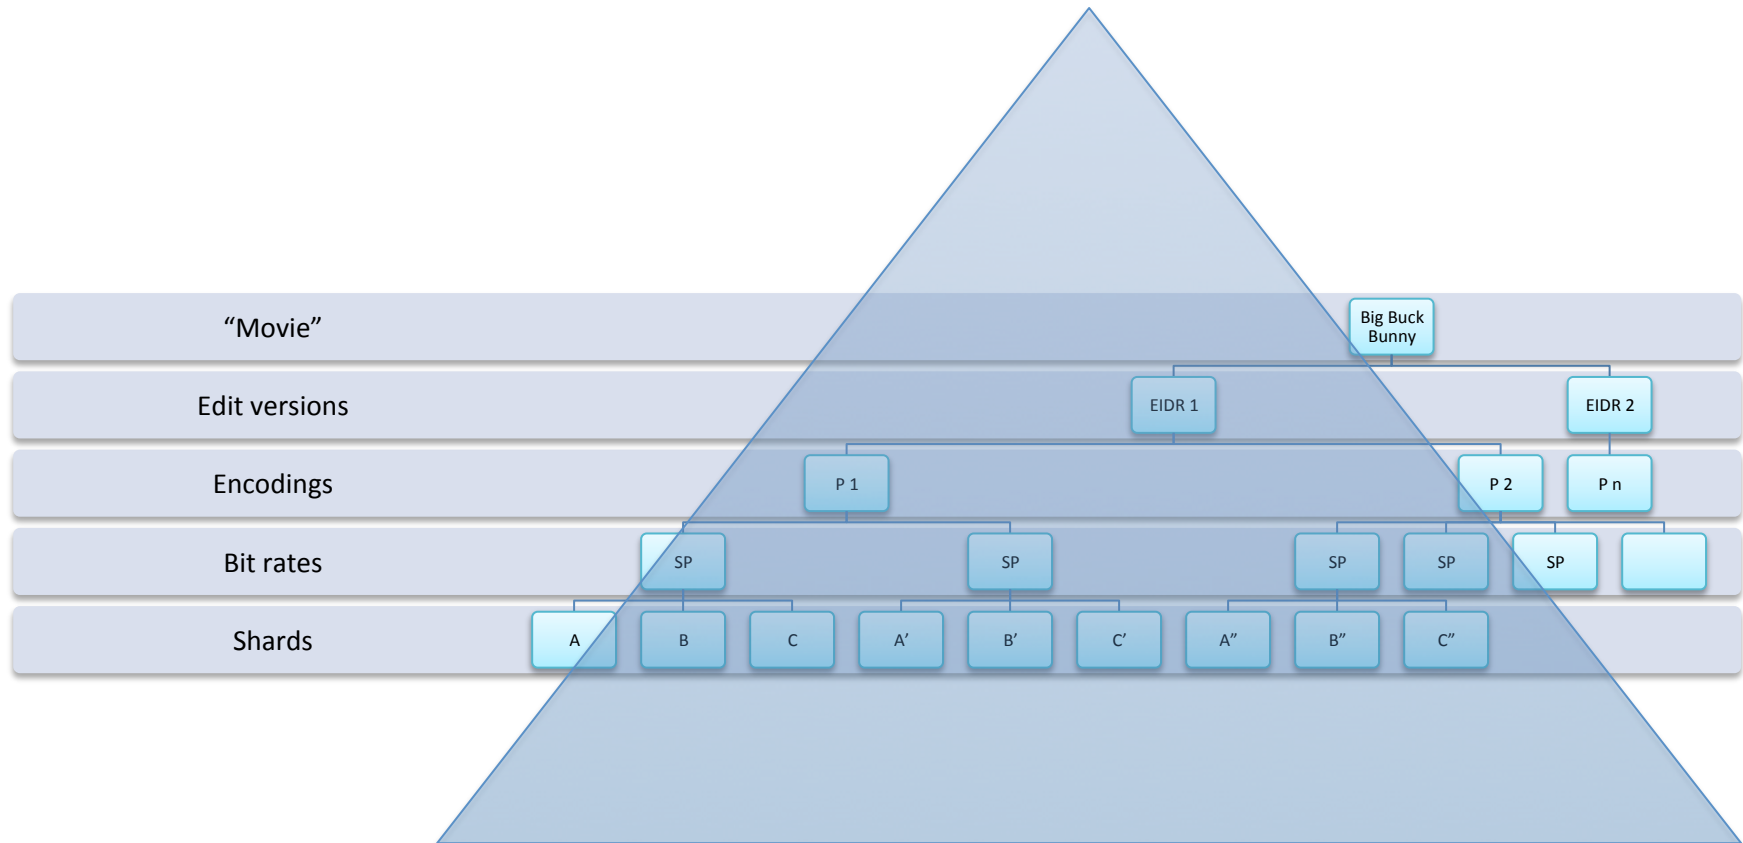


Thinking monolithically...

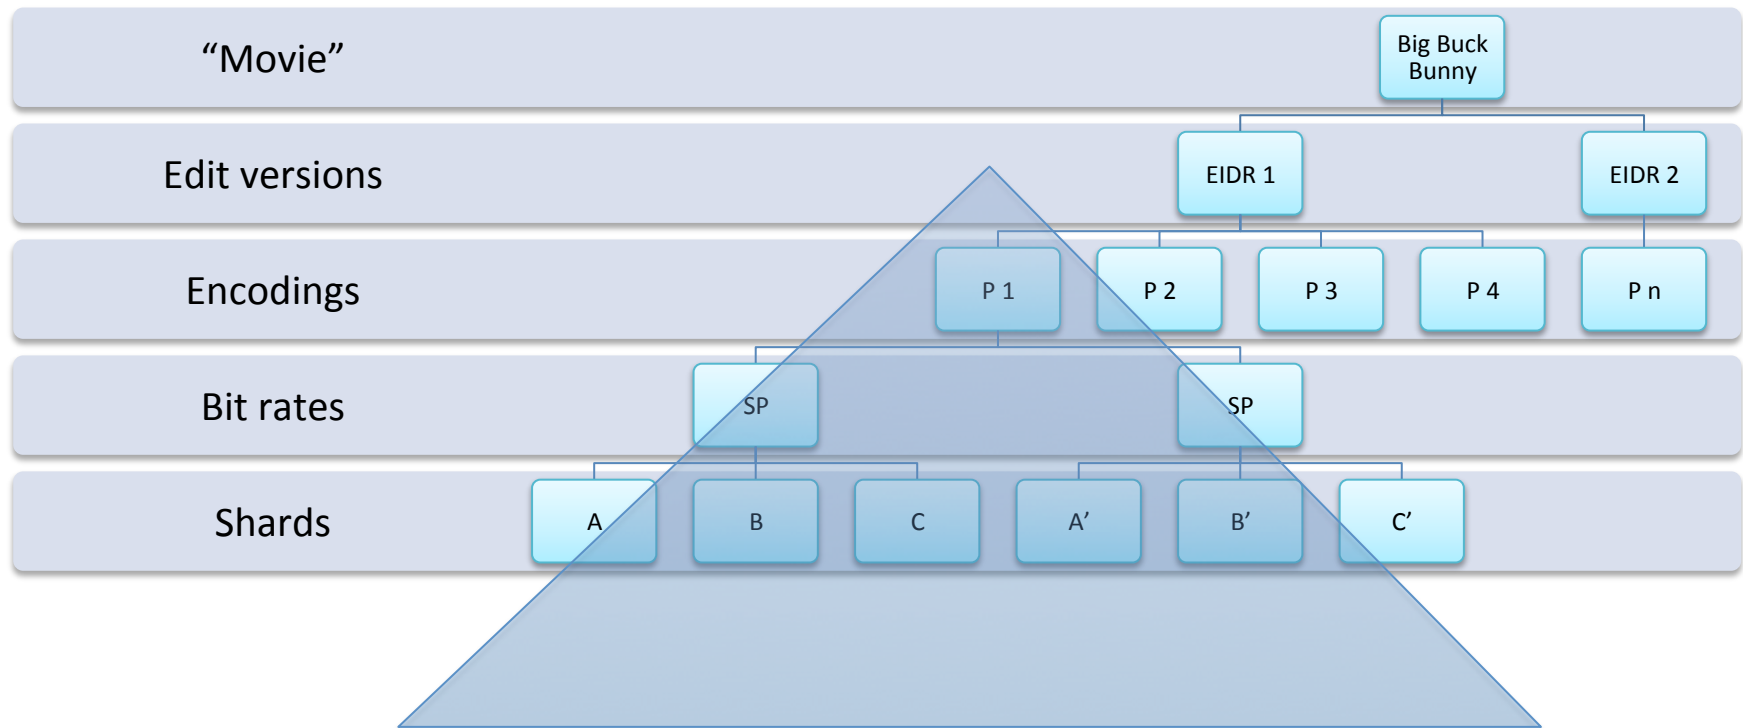
“Caching a movie”



“Caching a version”



“Caching an encoding”



Using the parts to advantage...

Video as a path

- Choices about which encoding (bit rate, resolution) may vary as the user is watching and conditions (devices, network) change
 - A->B->C
 - A->B'->C'
- It's more like a traversal of a path of possible shards of video

What and where do you cache stuff?

- Monolithically: cache all of it
 - Too big!
 - Wasteful!
- If you think of it as a composite: cache the next set of shards, because you know the structure
 - Back to dealing in small chunks, which is what the Internet has been dealing with for years

A theory, which is mine...

COMPOSITE TYPES

Internet Application Infrastructure Architecture

- To date
 - Has dealt with identifiers, lookup systems, etc
 - Assumes that data elements are units
 - Get
 - Post...
- Rather than thinking of (application) data as single blobs, or with structures specific to the application (e.g., HTML for WWW), we can think of composite application resources
 - Each resource includes a manifest of pieces of the data, with possible variants dependent on the application layer's preferences.
 - We can define actions on composites
 - E.g., “Streaming” means retrieving and processing successive pieces in real time.
 - “Equivalences” are expressed in the manifest — where the semantic content of two pieces are the same, but the application layer characteristics differ (e.g., bit rate, resolution, access controls)

Handling composites

- Might have a sparser “tree” of content to cache
 - because semantic equivalences are known
- Know where there are “logical breaks” in the structure
 - possibly, at the network layer

Conclusions

- We can't treat video as monolithic chunks of data
- Need to “think smarter” about how to handle it in an Internet-friendly way
- This is *one* aspect of the work underway on “Glass to Glass Internet Ecosystem” work
 - IETF – discussing BoF possibilities
 - Deen, Glenn and Leslie Daigle, “Glass to Glass Internet Ecosystem Introduction”
 - <https://tools.ietf.org/html/draft-deen-daigle-ggie-00>
 - W3C – GGIE Task Force discussed use cases
 - https://www.w3.org/2011/webtv/wiki/GGIE_TF

Idaigle@thinkingcat.com

THANK YOU!