Optimizing ATSC 3.0 Datacasting spectrum utilization with Advanced Spectrum Resource Techniques

Advanced ATSC 3.0 Spectrum Resource Techniques

PEAK3

DigiCap

Introduction

The focus of this paper is to define, through analysis, several high-value approaches which use an orchestrated, network-wide spectrum manager to optimize the utilization of ATSC 3.0 spectrum for datacasting.

The analysis starts with a basic fixed allocation of dedicated capacity on a per customer basis as the notional starting point.

We will define an A-SRM (Advanced Spectrum Resource Manager) as cloud-native software which manages multiple transmission strings and their associated RF spectrum. These A-SRMs can be local, regional, or national based upon the number and geographical diversity of the managed resources.

The key optimizations we present here include:

- Statistical multiplexing of shared capacity
- NULL filling of Datacasting Data
- Application aware transmission
- Opportunistic bandwidth recovery with automation

We will also discuss some key attributes of an A-SRM architecture including:

- Spectrum Monitoring and Visualization
- Job scheduling and queuing across multiple transmission strings
- Automation framework

These techniques, when used individually and in combination, will result in significant optimization of the available bandwidth and enable monetization of ATSC 3.0 datacasting at scale.

Datacasting Baseline

As a starting point for the analysis of the definition, design, and value of an A-SRM, we define the following operational baseline as the following:

- Deployed National coverage of ATSC 3.0
- ATSC 3.0 Transmission strings as per the standard including local ability to support:
    - Linear Video PLPs with a fixed data rate at a fixed modcod
    - Datacasting PLPs with fixed data rate at a fixed modcod
- No central Spectrum Resource Manager
- Local Transmission strings Media server supports a programmatic API.
- Changes to capacity are done manually by an operator during typical working hours with a local GUI.

This notional baseline state results in the following:

- Requests for national and regional datacasting distribution must be delivered manually, which:

  - Does not scale across many requests to a large number of strings without significant increases in operator staffing (and cost).
  - Error prone when repetitive actions are executed manually.
  - Additional cost to support scheduling off-hours operational updates and optimizations.
  - Lacks network wide and real-time visibility of the current available capacity and utilization.
- Significantly inefficient (bits/Hz) and under-utilization (fill rate) of the spectrum

Notional Datacasting/Broadcasting Capacity

To model the impacts of various proposed A-SRM techniques, a notional datacasting channel is modeled.

ATSC 3.0 Spectral Efficiency

The ATSC 3.0 Modcod table below was selected from the list of all possible ATSC 3.0 modcods to define set of points that monotonically increase in SNR and spectral efficiency without overlaps. When appropriate, the lower order modulation scheme was used as lower order modulation schemes perform better with phase noise and require less backoff from P1db in the PA.

Modcod	Raw Spectral Efficiency (bits/Hz)	Spectral efficiency w/Overhead (bits/Hz)	SNR
QPSK 2/15	0.267		-5.06
QPSK 3/15	0.400		-2.97
QPSK 4/15	0.533		-1.36
QPSK 5/15	0.667		-0.08
QPSK 6/15	0.800		1.15
QPSK 8/15	1.067		3.44
16QAM 5/15	1.333		4.78
16QAM 6/15	1.600		6.27
16QAM 8/15	2.133		8.96
16QAM 10/15	2.667		11.73
16QAM 12/15	3.200		14.97
64QAM 11/15	3.667		17.87
256QAM 9/15	4.800		18.64
256QAM 10/15	5.333		20.5
256QAM 11/15	5.867		22.4

This table will be used as a reference throughout this paper.

Raw Datacasting Capacity

The following table is provided to provide context for the datacasting capacity available to an operator to sell for a given data rate allocation.

Datarate (Mbps)	Modcod	Gbps/day	GBps/day	Gb/month	GB/month
0.1	Any	8.6	1.1	259.2	32.4
0.25	Any	21.6	2.7	648.0	81.0
0.5	Any	43.2	5.4	1,296.0	162.0
1	Any	86.4	10.8	2,592.0	324.0
2	Any	172.8	21.6	5,184.0	648.0
3	Any	259.2	32.4	7,776.0	972.0

The GB/month represents the total size of content sent over a datacasting PLP for a given Mbps rate sent continuously for 24 hours a day for 30 days.

Time to send a NRT object/file

The following table is provided as context for the length of time to send a file of assorted sizes across a range of data rates.

While most datacasting applications are considered non-real time, there are potential applications which rely on timely or “somewhat” timely delivery.

Datarate (Mbps)	Files Size (MB)	Time to transfer a file (s)	Time to transfer a file (min)
0.5	10	160	2.67
0.5	100	1,600	26.67
0.5	1000	16,000	266.67
0.5	5000	80,000	1,333.33
1	10	80	1.33
1	100	800	13.33
1	1000	8,000	133.33
1	5000	40,000	666.67
2	10	40	0.67
2	100	400	6.67
2	1000	4,000	66.67
2	5000	20,000	333.33
3	10	27	0.44
3	100	267	4.44
3	1000	2,667	44.44
3	5000	13,333	222.22

These examples are used to illustrate that small, dedicated channels might be used to improve channel (PLP) utilization but will result in long times to deliver the content which could be measured in hours and even *days*.

Baseline/unoptimized capacity utilization

The baseline ATSC 3.0 methodology is to assign a PLP for each customer, resulting in extended duration file transfer as described above and unused capacity per assigned PLP when the customer leaves the capacity unused.

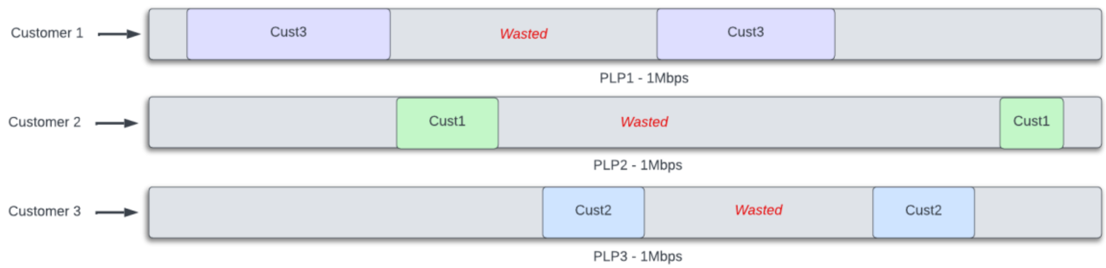

In this mode of operation, SLAs for a given customer are enforced by reserving a PLP at a given data rate at a fixed and worst case modcod. Idle capacity for a given customer is not usable by other customers/applications and is left unused.

The end-customer carries the burden of the wasted cost i.e., higher spectrum fees, or the operator loses out on potential revenue when the spectrum cost must be “right-sized” as compared to competing delivery technologies (4G/LTE, 5G, Satellite, etc.).

Furthermore, hard allocated capacity via a dedicated PLP does not allow for oversubscription of the capacity i.e., re-using capacity when the customer is not actively using it.

A-SRM Architecture

The following section describes a notional architecture for a network-wide orchestration A-SRM to use as a reference for the suggested and modeled optimizations which follow.

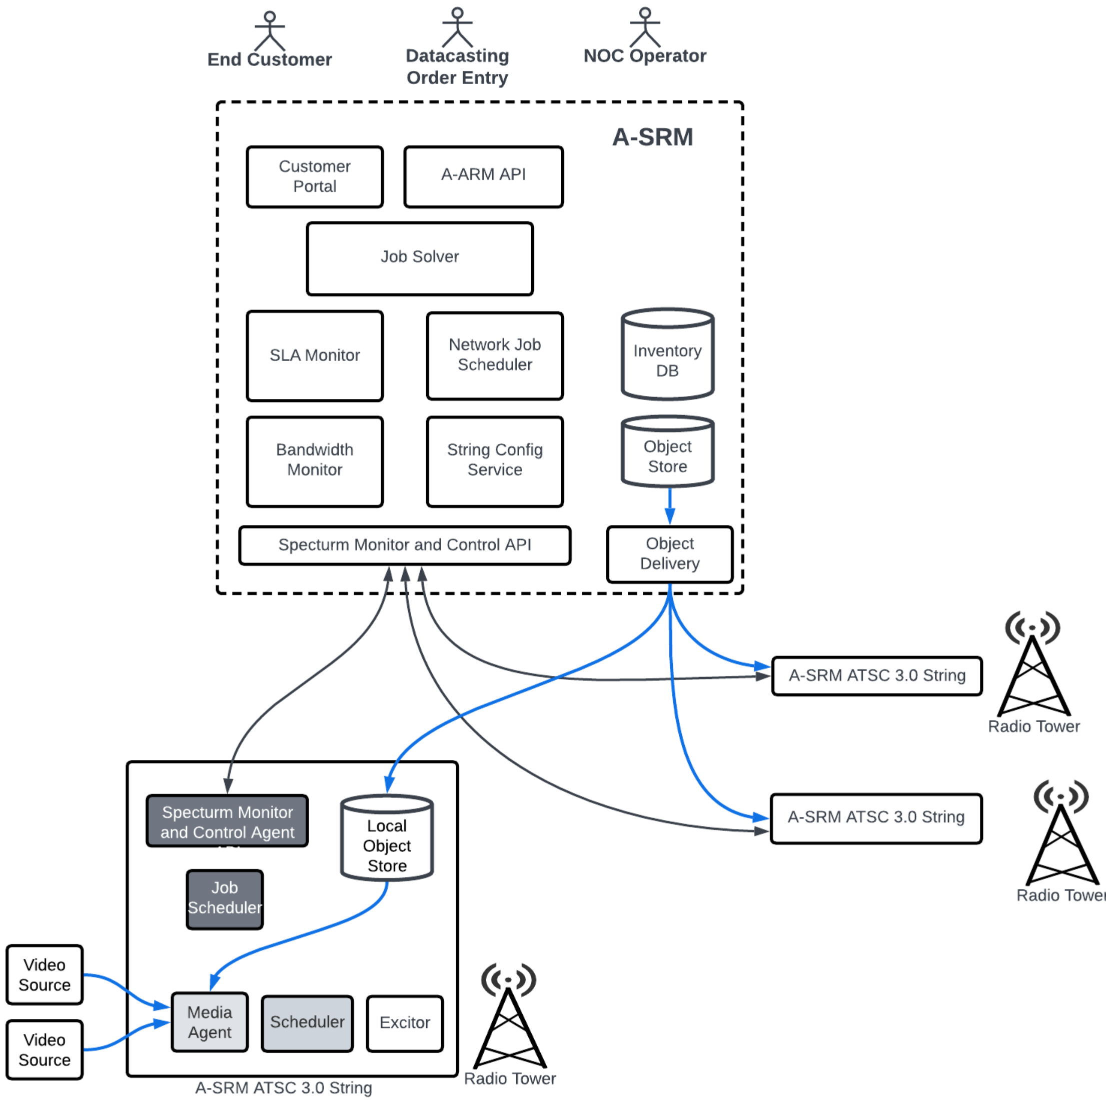

The A-SRM Cloud-native hosted components include:

- Monitor and Control API between A-SRM and the string local Scheduler
- Central inventory of network hardware and software components
- Monitors network utilization and application health
- Ability to add and delete PLP services for one or more sticks as an automated batch change.
- Object Store for Datacasting content
- Object Delivery Agent to distribute datacasting objects to local sticks.

Software Enhancements to local Media/Scheduler

- Multiple modcods in a PLP
- Local job scheduler to stat mux a PLP
- NULL packet Datacast injection

Based on the proposed A-SRM architecture, the following optimization techniques are proposed to enhance the available spectrum to increase utilization.

Optimization #1 – Bandwidth Sharing through Statistical Multiplexing

As a terminology baseline, statistical multiplexing is performed by switching systems in communication networks that merge data packets from multiple input lines and forward them to multiple outputs in a first come first serve or other scheduling discipline. In this way, many data flows can share capacity on a common transmission path.

For our discussion on ATSC 3.0 NRT datacasting, we propose the creation of a NRT Job Scheduler function which receives NRT Object/Files from the central A-SRM for multiple customers and multiplexes (shares) the PLP capacity based upon time scheduled and priority assignment.

A properly defined String Local NRT Job Scheduler would provide the following:

- Able to achieve 100% utilization of the shared NRT PLP
- Job Scheduler will prioritize customers based upon SLAs as defined by the A-SRM
- File transfer happen faster in time if the shared data rate is higher than per customer PLP data rate.
- Allows for efficient delivery of high priority jobs.
    - Existing lower priority jobs are paused, high priority jobs are sent, and then paused jobs are resumed.

Optimization #2 – Opportunistic Baseband utilization

Opportunistic utilization to 100%

Spectrum is not typically 100% full as this would lead to congestion and potential dropped information for critical services i.e., glitches in the TV channel services. As a result, there is potential for empty (NULL) BB frames to exist in the transmitted video stream.

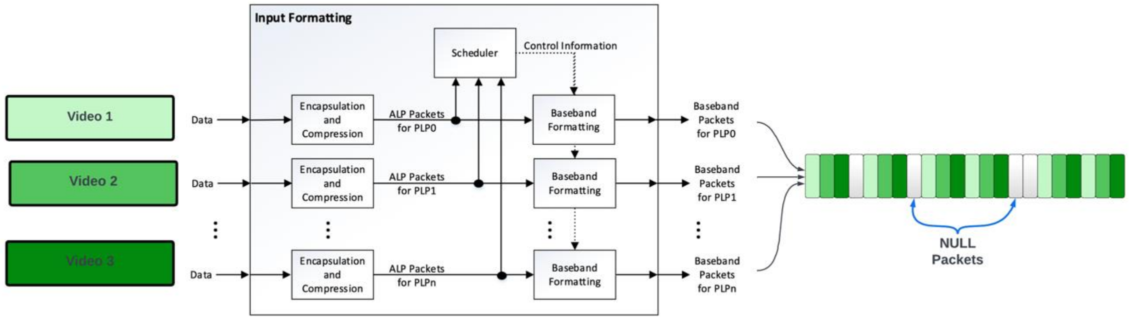

As NRT (Non-Real Time) Datacasting traffic is not typically time sensitive or impacted by jitter, opportunistic filling of otherwise empty/NULL BB Frames can be used to fully capture unused capacity.

In this proposed enhancement, a new type of PLP is defined in which datacasting content is 100% opportunistically sent rather than configured with a fixed data rate. The scheduler pre-builds an NRT Baseband frame and has it ready in low-level buffer memory to send whenever a NULL BB frame would have been sent during normal operation.

With this approach, the existing NRT files are streamed out without impacting the existing video service PLP operation i.e., no impacts to jitter, latency, or capacity.

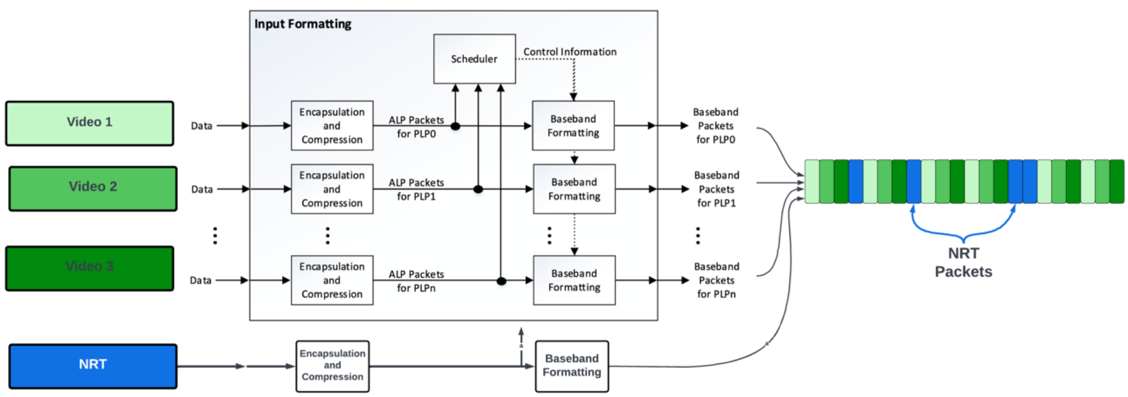

It is estimated that this approach would yield a 1 to 15% increase in available Datacasting capacity with no additional allocated spectrum or impact on existing video services.

Next Steps

To validate this proposed optimization, it is suggested that the existing NULL packet insertion statistics/telemetry be pulled from the deployed network to validate the premise and rate at which NULL packets are being transmitted.

In addition, it should be determined if the scheduler can be optimized to more completely fill a given BB frames and thus leave opportunity for more NULL packets rather than partially filling/smoothing all the frames and leaving the potential savings un-realized.

Optimization #3 – Application Aware Spectrum Utilization

In this proposed optimization, A-SRM and local scheduler is enhanced to be aware of the end customer's application.

As we look forward to the various datacasting use cases, a wide range of endpoints will be deployed. For example:

- Fast Mobility: Streaming content such as videos and adds to trains.
- Mobility: Upgrading software for EV cars
- Fixed Indoor antenna: Timing distribution for industrial applications
- Fixed with outdoor antennas.

As described above in our treatment of the modcods and robustness vs bandwidth, we make the simplifying statement about the various modulation schemes:

- QPSK is extremely robust, ideal for fast moving vehicles.
- 16QAM is robust, ideal for mobile reception.
- 64QAM is ideal for indoor antenna reception, good building penetration and semi-mobile.
- 256QAM is best for fixed outdoor antenna reception.

When sharing a PLP, if the scheduler is enhanced to maintain a relationship between applications modcod, the best of statistical multiplexing and best fitting the modcod to the application can be achieved.

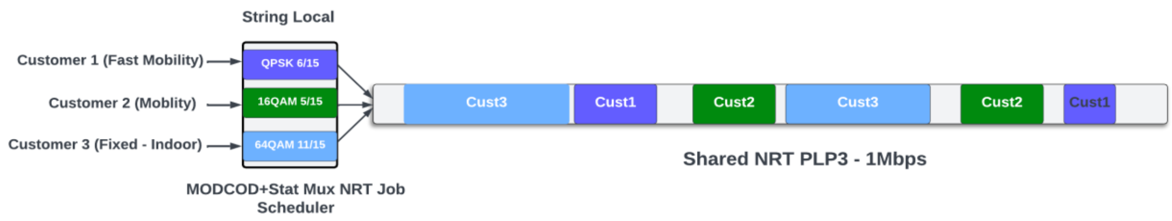

Figure 1Modcod and Stat Mux NRT Job Scheduler

Optimization #4 – Dynamic Spectrum Grooming

Once a network wide A-SRM is available with the ability to monitor and manage the capacity and services across the network, additional savings are available by dynamically adjusting spectrum usage.

This paper suggests that based upon the local network and service the option exists to adjust the fidelity of the existing streams during low or ultra-low viewership periods to free up additional capacity.

In particular, a given service could be lowered from UHD to HD, or HD to SD from 1am to 5am nightly to free up datacasting capacity. Datacasting file delivery is often not ultra-time sensitive which makes them a perfect fit to send in off-hours.

The A-SRM would orchestrate a lower fidelity service at a configured time as appropriate to the service i.e., based upon viewership, local time zone, etc.

A more detailed treatment of the bandwidth made available are addressed in a later section.

A-SRM Capacity and Value Model

This section provides a model for additional capacity for each optimization approach for a single transmission tower. The working assumption is this translates into a network wide increase in available capacity to use for datacasting revenue.

Our approach will be to model the potential available GBs of available capacity for each approach.

A-SRM Statistical Time-Division Multiplexing on Shared NRT PLP Value Model

As a baseline for the analysis of ATSC 3.0 Statistical TDM channels, we describe typical datacasting sources as:

- Non-real time datacasting (software upgrades, movie files, etc.)
    - Small: <10MB
    - Medium: < 1GB
    - Large: >1GB
- Real time (less than 5 minutes) such as
    - Weather maps for first responders
    - Situational updates
    - Typically smaller in size but sent more often.

The value of statistical TDM increases as the number of input sources increases and will closely approximate a completely full/utilized channel.

The time to deliver the content will typically be included in the SLA and is included in the model below. Note that the SLA could also be used as an additional parameter to price services i.e., faster delivery comes with premium pricing.

A prioritized statistical TDM will allow for low-latency input sources to be sent which explains why were able to combine all services into a single Statistical TDM channel below.

While mathematical models exist for statistical TDM, ultimately, the benefit for a given ATSC 3.0 datacasting channels is defined by the number of data sources, utilization (how often each source utilizes the channel) and the customer's required delivery time.

As such, the value of statistical multiplexing would be the:

Saved/Available GB = Full Channel – Sum (Per Source Unused Capacity at required delivery time)

As a reference, the following is an example of (5) datacasting PLPs which deliver capacity on a fixed basis to meet a customer's max delivery time for each input source vs. a shared and perfectly sized statistical TDM PLP which services all input streams.

Customer	Files (MB)	Occurence (per month)	GB/Month	Max Delivery Time (hrs)	Required Channel size (Mbps)	% Utilized (per month)	GB Unused
Large Software Update	10,000	1	10	24	0.93	3.3%	2320.0
Medium Software Update	2,000	2	4	24	0.19	6.7%	448.0
Real Time Update	5	50	0.25	0.1	0.11	0.69%	286.0
Movie Files (HD)	3,000	10	30	24	0.28	33.3%	480.0
Movie File (4K)	28,000	10	280	48	1.30	66.7%	1120.0
				Total	2.80	22.1%	4654.0
Stat MUX Channel	All above	All above	324.25		1.00	100%	0

Analysis: In this example, a 2.8Mbps service is required to satisfy all inputs while 1.0Mbps is required when Statistical TDM is used for a savings of 1.8Mbps (2.8-1.0) over the naïve approach.

However, to meet the 48hour delivery time of the “Movie File”, a 1.3Mbps channel must be commissioned which results in a 1.5Mbps savings over a non-statistical TDM approach. All other “Max Delivery Time” requirements are meet with the shared channel.

A-SRM Null fill Value Model

Using a 3Mbp video service as a basis, the chart below shows the potential additional GB per month which would be available to monetize based upon the percentage of available BB frames which would have otherwise gone out empty/NULL.

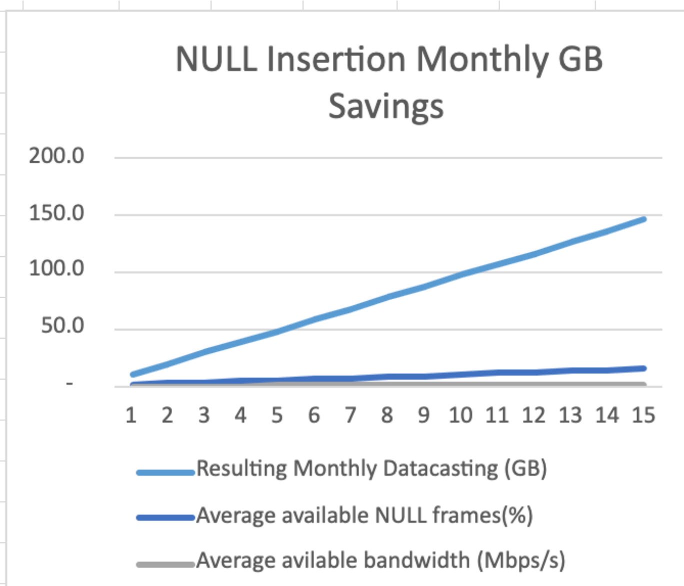

Null frame optimization Savings															
Average available NULL frames(%)	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Service size (Mbps)	3	3	3	3	3	3	3	3	3	3	3	3	3	3	3
Average avilable bandwidth (Mbps/s)	0.03	0.06	0.09	0.12	0.15	0.18	0.21	0.24	0.27	0.3	0.33	0.36	0.39	0.42	0.45
Resulting Monthly Datacasting (GB)	9.7	19.4	29.2	38.9	48.6	58.3	68.0	77.8	87.5	97.2	106.9	116.6	126.4	136.1	145.8

A-SRM Application Aware Value Model

Using a naïve approach as a baseline, all traffic across a range of applications use the QPSK modcod.

Application	Modulation	Coderate	Spectral Efficiency (bits/Hz)	Datarate Needed	Hertz/Symbols Used (1MHz)
Fast moving vehicles	QPSK	4/15	0.533	1Mbps	1.9
Mobile Reception	QPSK	4/15	0.533	1Mbps	1.9
Indoor antenna	QPSK	4/15	0.533	1Mbps	1.9
Fix outdoor Antennas	QPSK	4/15	0.533	1Mbps	1.9
				Total	7.50

Using the application aware approach, applications with different coding gain requirements are sent on per application the highest possible modcod.

Application	Modulation	Coderate	Spectral Efficiency (bits/Hz)	Datarate Needed	Hertz/Symbols Used (1MHz)
Fast moving vehicles	QPSK	4/15	0.533	1Mbps	1.9
Mobile Reception	16QAM	6/15	1.600	1Mbps	0.6
Indoor antenna	64QAM	11/15	3.667	1Mbps	0.3
Fix outdoor Antennas	256QAM	10/15	5.333	1Mbps	0.2
				Total	2.96

As can be seen from this example case, the amount of spectrum saved to deliver 4Mbps of traffic across a range of user applications is roughly 60%.

The benefit of this approach will vary based upon the type and mix of user traffic required but would scale for regional and national service offerings.

It is suggested that a per customer assignment of application to modcod could be a viable approach to simplify deployment.

A-SRM Dynamic Spectrum Grooming

While additional bandwidth saving grooming approaches exist, we choose to focus on the approach described above. The network wide A-SRM will orchestrate a lower fidelity broadcast fidelity during periods of low to ultra-low viewership. This grooming could also be configured to behave differently for different channels.

Video Speeds	Mbps
UHD	6
HD	3
HD to SD Savings	1.9
UHD to HD Savings	3

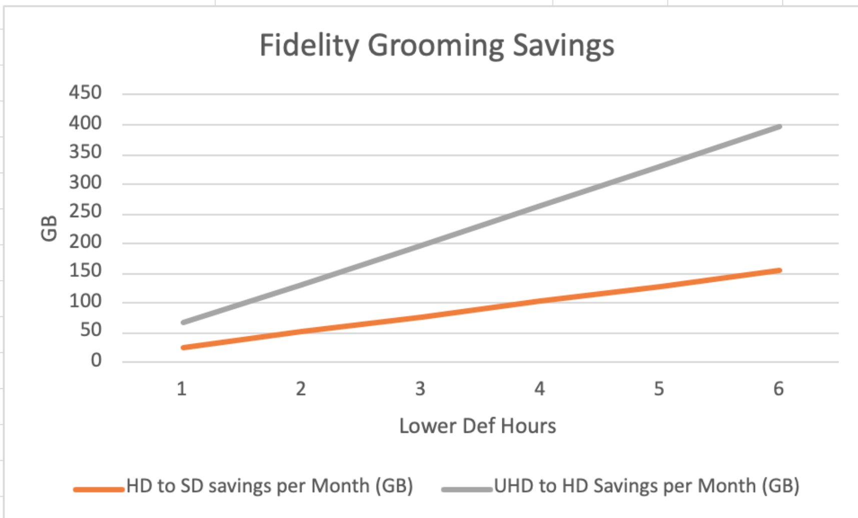

Dynamic Fidelity Scheduling per Service Savings						
Lower Definition period (hrs)	1	2	3	4	5	6
HD to SD savings per Month (GB)	25.65	51.3	76.95	102.6	128.25	153.9
UHD to HD Savings per Month (GB)	40.5	81	121.5	162	202.5	243

As modeled above, significant additional capacity for datacasting can be made available with the use of an orchestrated A-SRM.

Additional Ideas for further study

Receive Return Channel

If an endpoint has a return data path, such as LTE/5G, the current received SNR could be sent back to the A-SRM. This information could be used to accurately adjust the associated modcod based upon the worst-case endpoint.

Conclusion

This paper suggested several significant bandwidth optimizations which allow additional bandwidth to be harvested from existing networks while greatly enhancing any dedicated datacasting capacity.

The ability to take advantage of the bandwidth optimizations suggested in this paper depends on the ability to deploy them at scale using a network wide Advanced SRM which can monitor bandwidth while automating tasks.

The Broadcasters who have invested large sums of capital into their infrastructure aiming to demonstrate the advantages of the ATSC 3.0 technical architecture and its business value proposition to the commercial edge marketplace, must abide by the number one rule and focus of scale… the customer journey.

Proving their solution and service works on a regional scale is not going to move the dial when competing with incumbent telecoms. National coverage and a "seamless" service offering, regardless of which Broadcaster the owns a market a customer operates in or is driving though, the time of day its being delivered, the customer count possibly relying on that tower, etc., is critical for success.

This requires the "solution stack" to include the necessary pieces to provide the promised service and experience and thus scale will occur. While SRM may be omitted when leveraging ASTC 3.0 locally, it is hyper-critical that it be integrated as your starting quarterback when national accounts are being engaged and rolled out.

Please reach out to our team (below) to learn more.

About PEAK3

PEAK3 has a long-term engagement with the spectrum owners to offer this alternative data highway to its many technology partners, its enterprise clients, and to its engineering teams to further develop novel applications where the ATSC 3.0 value proposition can enable the edge-device community. From hardware architecture through the application layer, the PEAK3 team has a rich history of successful enterprise edge deployments and IT system designs and operations.

We recognize the unique value proposition ATSC 3.0 provides in delivering secure, efficient, data-delivery methods to the edge. The foundation of our business model is Data-Streaming as a Service for organizations wanting to efficiently get data from one point to many.

PEAK3 provides a standards-based, open, end-to-end, nationwide, wireless, IP, multicast network. In simple terms, we provide a cost-effective datacasting pipe for Internet Service Providers, public and private cloud providers, and any organizations operating large edge device architecture.

About DigiCap

For the past 20 years, DigiCAP has used advanced software development to make television and telecom systems easier to use, less costly, and more profitable. DigiCAP's ATSC 3.0 air chain, DigiCaster, has now been deployed in over half of the ATSC 3.0 market deployments in the US. DigiCaster contains all the software components for ATSC 3.0 transmission, including ROUTE/MMT packager, NRT, ESG, multiplexer and broadcast gateway (or scheduler). DigiCAP also has been providing major telecom providers, broadcasters, and government agencies with end-to-end software solutions to optimize their multimedia business operations.

DigiCAP manufactures the HomeCaster device, an ATSC 1.0/3.0 home gateway product, and supplies broadcasters and system integrators with B2B solutions.

One application of HomeCaster is transmitting learning materials with ATSC 3.0 datacasting. HomeCaster has been deployed in 14 states in the US as a distance learning solution to K-12 students and proved especially useful during the COVID-19 pandemic. DigiCAP is working with broadcasters and system integrators to continue to expand its use cases for other business applications including IOT, precise location positioning, digital signage, among others.