Beyond 1.8 GHz

Evaluating Paths to 3 GHz HFC Networks

WHITE PAPER

Table of Contents

Introduction

Spectrum expansion in the cable industry has proceeded through a series of incremental upgrades. The progression is familiar: 450 MHz to 750/860 MHz, 1 GHz, 1.2 GHz, and now 1.8 GHz. Discussions about extending the network even further are not new.[1] Although the migration still remains years away, operators and equipment suppliers are actively evaluating what it would take to extend HFC networks toward 3 GHz and beyond.

Much of the industry's DOCSIS evolution over the past decade has focused on the CMTS, Distributed Access Architectures (DAA), and access technologies. Moving beyond 1.8 GHz expands attention to the physical HFC plant. The central question is no longer simply whether 3 GHz operation is possible, but how amplifiers, powering, network architectures, and the outside plant itself must evolve to make it practical.

This paper first presents a framework for evaluating the questions associated with extending HFC spectrum. The central section reflects the impact of increased attenuation at higher frequencies and examines several amplifier architectures currently under consideration. Then the paper concludes with some practical guidance on preparing today's networks for continuous spectrum expansion.

Paper Scope

While spectrum expansion involves the upstream and downstream, this paper focuses on the downstream path, where the principal engineering questions surrounding attenuation, amplification, TCP, and emerging amplifier architectures are currently being evaluated.

Spectrum expansion in HFC networks. Each upgrade has extended the downstream ceiling; 3 GHz is the next step under evaluation.

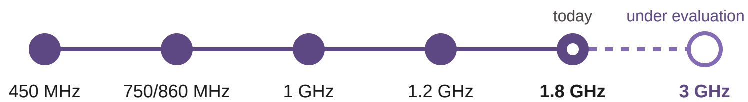

The Decision Framework

Every major HFC upgrade has introduced new questions. The move beyond 1.8 GHz is no different. Before examining the emerging amplifier architectures, however, it is helpful to think through the 3 GHz decision in a structured way. The framework below has three dimensions: engineering constraints, design responses, and decision criteria.

> **"Higher attenuation is what makes amplification central to extending HFC networks beyond 1.8 GHz."**

Engineering Constraints

Moving to 3 GHz has measurable consequences. As downstream spectrum expands, engineers must contend with higher attenuation, greater power and thermal demands, tighter mechanical constraints, and threats to signal integrity.

Attenuation. Signals naturally lose strength as they travel through coaxial cable. The critical point is that this attenuation increases with frequency: a 3 GHz signal loses much more strength over the same distance than a 1 GHz signal. (See Figure 1.) This physical reality is what makes amplification central to extending HFC networks beyond 1.8 GHz.

Figure 1: Coaxial cable attenuation vs. frequency

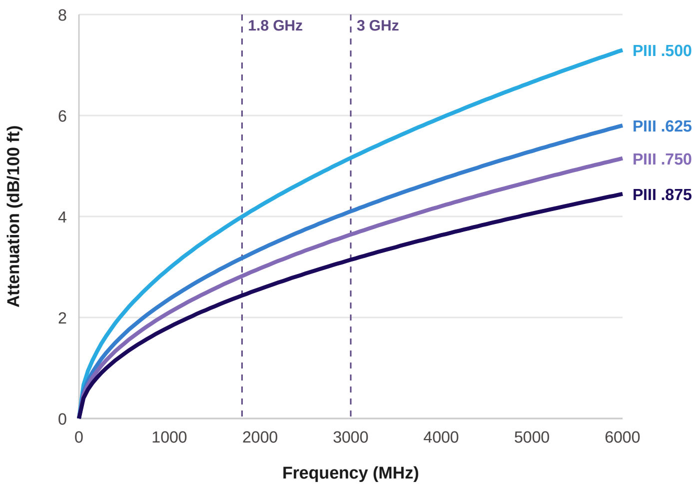

Attenuation rises with frequency for all hardline cable sizes. Values shown are representative of common PIII trunk and feeder cables, in dB per 100 feet.

Power. Historically, plant upgrades to higher frequencies have required amplifiers to deliver more RF output to overcome increasing cable loss. But adding power also increases heat, reduces efficiency, and may require changes to network powering and amplifier architecture. As a result, power allocation and amplifier efficiency become central design considerations.

Heat. Higher RF output and more complex electronics generate additional heat that must be dissipated in sealed outdoor housings. An open question is whether existing amplifier housings can continue to support future spectrum expansion.

Mechanical design. Higher frequencies place tighter demands on the physical design of outside plant components. Small variations in seizure pins, connectors, housings, and other mechanical interfaces that had little impact at lower frequencies can measurably affect insertion loss, impedance matching, and overall signal integrity.

Signal Integrity. Maintaining signal quality across the outside plant becomes more challenging as operating frequencies approach 3 GHz. Modern DOCSIS networks rely on OFDM and higher-order QAM modulation to maximize spectral efficiency, but these techniques require cleaner signals and better CNR to achieve their full performance.

Design Responses

The constraints described above help define the design space within which engineers are exploring a variety of responses. These generally involve amplifier philosophy, digital enhancement, network architecture, and migration strategy, with increasing emphasis on smarter, more connected outside-plant equipment.

The first question is how much work each amplifier should perform. Some approaches seek to increase total composite power (TCP), the combined RF output of all signals leaving the amplifier; others reduce TCP and distribute gain more broadly throughout the network.

Digital signal processing (DSP) represents a complementary engineering response. Rather than replacing amplifier architectures, DSP techniques may improve linearity, efficiency, and signal quality, depending on the architecture being evaluated.

Engineering responses also extend beyond the amplifier itself. Amplifier spacing, distributed gain, and cascade design influence how attenuation is managed throughout the network and how efficiently spectrum can be extended.

Finally, engineers need to keep migration in mind. Modular upgrades and reusable housings are two existing ways to preserve future options.[2] More connected actives, such as 1.8 GHz smart amplifiers, are another. However else engineers respond, boosting visibility and network intelligence will continue to be prerequisites to moving forward.

Decision Criteria

Any approach to extending spectrum must also satisfy operational and economic realities. For operators, the evaluation comes down to three practical questions: Can we afford it? Can we deploy it? And will customers notice?

Beyond the price of 3 GHz amplifiers, operators must consider taps, passives, power supplies, housings, labor, network engineering, testing, training, and the ongoing energy costs of upgraded equipment. Upgrade longevity also matters: How long will current investments remain useful? Then there are opportunity costs to consider, including fiber-deep deployments or fiber replacement.

A new architecture may call for replacing taps or passives, relocating amplifiers, or upgrading power infrastructure. It would also need to accommodate the diversity of HFC networks, where operators often manage multiple generations of plant. The viability of any 3 GHz architecture depends on how practical and economical it is to deploy across millions of passings.

As for customers, what matters is not spectrum itself, but faster speeds, lower latency, greater reliability, and new service capabilities. If today's networks can satisfy demand for years, premature investment weakens the business case. But if another major capacity expansion is coming, designing with future spectrum in mind makes sense.

Three practical questions

1. Can we afford it?
2. Can we deploy it?
3. Will customers notice?

Emerging Amplifier Architectures

What about specific, amplifier-based approaches? This discussion begins with wideband amplifier philosophies, both high- and low-TCP. Then it turns to segmented, channelized, and hybrid architectures that distribute amplification and digital processing in different ways. Together, they illustrate a range of engineering tradeoffs.

> **"One approach to extending HFC networks toward 3 GHz is to preserve the traditional wideband amplifier"**

Wideband Amplification

One approach to extending HFC networks toward 3 GHz is to preserve the traditional wideband amplifier, in which one RF path amplifies the entire downstream spectrum. The central question is whether one amplifier can continue to both deliver sufficient TCP and maintain acceptable linearity, efficiency, and power consumption.

High TCP wideband

High-TCP architectures concentrate the engineering challenge within the amplifier by increasing output capability. The objective is to extend downstream spectrum and minimize changes, avoid additional amplification stages, and preserve a single RF path. Compared with channelized approaches, High TCP also avoids guard-band losses.

The tradeoff is that TCP rises dramatically as downstream bandwidth expands. Extending today's typical amplifier RF output profile to 3 GHz would increase TCP by more than 10 dB over current amplifier levels.[3] Amplifier linearity becomes more difficult to maintain, power consumption increases, and existing network powering may prove inadequate.

Engineers are also evaluating whether DSP can help overcome some of these limitations. Techniques such as digital pre-distortion (DPD) may improve amplifier linearity, allowing higher composite output without a corresponding increase in distortion. However, they introduce additional circuitry, power consumption, and complexity, and their role in future High-TCP architectures remains under evaluation.

These considerations have prompted engineers to explore whether similar network performance can be achieved through lower-TCP architectures that distribute amplification more broadly throughout the network.

Low TCP wideband

Rather than increasing the output capability of each amplifier, Low TCP limits the composite RF output expected from any one device. Operating at lower TCP allows each amplifier to run with greater linearity, thus reducing thermal stress and potentially lowering station power requirements.

The tradeoff is that gain must be distributed more broadly throughout the network. That may require additional amplification stages, like low-gain, low-output booster amplifiers, or modest amplification integrated into taps.

The Low TCP approaches that distribute gain more broadly naturally shorten the distances between amplifier locations. Closer spacing reduces loss between amplifiers but also increases the number of actives. The optimum spacing depends on plant attenuation, amplifier output levels, total power consumption, and desired operating margins.

Narrowband and Channelized Architectures

Rather than asking a single amplifier to support the entire downstream spectrum, narrowband architectures divide that spectrum into multiple frequency bands or channels. By reducing the bandwidth handled by any individual amplifier, these approaches may ease some of the composite power and linearity challenges associated with wideband operation. But there are tradeoffs: greater hardware complexity, additional filtering, and more sophisticated signal processing. (See Table 1.)

Split-band amplification

The simplest departure from a traditional wideband amplifier is to divide the downstream spectrum into two frequency bands, each served by its own amplifier. This split-band approach reduces the bandwidth, and therefore the composite power, that each amplifier supports. The two bands are separated using RF filters, amplified, and then recombined before being transmitted through the network.

The result may be improved amplifier linearity and fewer of the challenges associated with very high TCP operation. The tradeoff is additional RF hardware, including filters and combiners, some loss of usable spectrum in the guard-band regions between adjacent bands, and the potential for higher power consumption due to needing an extra power amplifier for each amplifier output port.

“Rather than asking a single amplifier to support the entire downstream spectrum, narrowband architectures divide that spectrum into multiple frequency bands or channels.”

Channelized amplification

Channelized architectures extend the split-band concept by partitioning the downstream spectrum into multiple narrower frequency bands, each served by its own amplifier. Because each amplifier supports only a fraction of the total downstream bandwidth, channelization further reduces composite power requirements and allows each band to be optimized independently.

That optimization may include different gain settings, frequency-response compensation, or DSP techniques. Partitioning the spectrum may also provide greater flexibility as spectrum allocations continue to evolve. The tradeoff is increased complexity: more power amplifiers, filters, combiners, DSP, and control software. Additional filtering also reduces usable spectrum, as with split-band.

Hybrid architectures

Hybrid architectures combine analog and digital techniques, applying each where it offers the greatest advantage. For some of the downstream spectrum, traditional analog RF circuitry continues to perform the proven amplification functions. DSP is introduced for the remaining spectrum to improve signal conditioning, linearity, or distortion correction.

The objective is to preserve analog performance and apply DSP, but only where it delivers measurable benefits. Digitizing selected channels rather than the entire RF path may improve efficiency. Selective DSP can also help control power consumption and reduce implementation risk.

The tradeoff is again greater system complexity. Hybrid architectures combine analog RF, digital electronics, and software. That combination increases integration challenges as well as the cost of DSP hardware, high-speed converters, firmware, and control software. Whether hybrid emerges as a preferred approach remains to be seen, but its appeal lies in balancing mature analog techniques with targeted digital enhancement.

Table 1: Split-band, Channelized and Hybrid Amplifier Architectures

APPROACH	BASIC CONCEPT	POTENTIAL ADVANTAGES	PRIMARY TRADEOFFS
Split-band	Divide DS spectrum into two bands, each with its own power amplifier (PA)	Reduces TCP per PA while retaining relatively straightforward analog architecture	Additional PAs, filters, and combiners per output port; guard band losses
Channelized	Partition DS into multiple narrower bands, each amplified independently	Further reduces TCP per PA; enables more selective DSP	More PAs, complex filters with guard band losses, DSP complexity
Hybrid	Combine analog RF amplification with selected DSP	Balances proven analog performance with targeted digital optimization; avoids digitizing entire path	Similar tradeoffs as Split-band and Channelized; greater design and integration complexity

Preparing for Continuous Spectrum Evolution

The preceding discussion of amplifier architectures suggests there may be more than one path to 3 GHz. But if divergence leads to different product types, a viable market may not emerge. Avoiding that undesirable outcome will likely fuel a drive towards consensus.

At the same time, flexibility remains a virtue. Whichever architectural direction the industry adopts should be able to fit with a variety of plant designs. For their part, operators need to remain adaptable. Here are some recommendations for doing so:

Design for Flexibility

Because several architectural approaches remain under consideration, operators should prioritize investments that preserve future options. Hardware that can evolve, amplifier locations that support future upgrades, and infrastructure decisions that avoid limiting future architectures all advance that cause.

Modernize Strategically

Operators should evaluate their overall outside plant to ensure that any current modernization efforts support multiple future migration paths. That includes powering, amplifier locations, housing capabilities, cascade length, and opportunities for modular upgrades. Standards-compliant taps and passives that support 3 GHz (at least in their housings) are a good place to start.

Leverage Connected Infrastructure

Operators can also prepare today by taking advantage of the smart and connected capabilities already available in 1.8 GHz platforms. Building operational experience with software tools, monitoring, and network intelligence will help ease future transitions regardless of the direction ultimately adopted.

Common Goal, Diverse Paths

Extending HFC networks toward 3 GHz is a daunting initiative. It requires operators to work within more stringent constraints, reconsider amplifier architectures, and balance performance against operational and economic realities.

As this paper has shown, the industry is exploring multiple approaches, from high-TCP and distributed-gain philosophies to channelized and hybrid architectures. No single solution has yet emerged as the clear winner, and ongoing research, demonstrations, and trials will continue to shape the way forward.[4]

The good news is that by preserving flexibility, modernizing the outside plant strategically, and planning smartly for continuous evolution, operators can position their networks to capitalize on whichever architectural approaches prove most viable. The destination is clear, even as the paths remain under active development.

Ongoing research, demonstrations, and trials will continue to shape the way forward.

Notes and References

1. See “Blueprint for 3 GHz,” Chapman et al., SCTE Cable Tec, 2019. In the paper’s sub-title, the authors underscored a driving motivation: “Getting 25 Gbps PON-Like Performance Out of HFC.”

2. A good example of optionality is the SCTE’s creation of standards for hardline taps and hardline passives (ANSI/SCTE 264 2021 and ANSI/SCTE 265 2021) that called for housings capable of working to 3 GHz, while the RF circuitry (commonly in a pluggable faceplate) was mandated to work to 1.8 GHz.

3. Power is measured on a logarithmic decibel (dB) scale. A 10 dB increase corresponds to roughly 10x the RF power, making it an order-of-magnitude increase.

4. At SCTE TechExpo25, CableLabs announced plans to begin industry work enabling DOCSIS data transmission up to 3 GHz. Work on an optional 3 GHz Annex to the DOCSIS 4.0 specifications is currently underway.