

Sub 2-Year System-Wide True In-Car, In-Tunnel Broadband

Kristofer Spinka

Confidential proposal compiled on May 17, 2018

A "moving-cell" model whereby whole-train assemblies are provided with a minimum of 40Gbps of seamless data backhaul at full train speed throughout the entirety of the NYC Subway System allowing riders to utilize in-car WiFi and LTE as well as bandwidth for realtime UHD in-car video surveillance.

A. Applicant Information

Applicant Name: Kristofer Spinka
Contact Phone Number:
Contact Email Address:
Application Title: Sub 2-Year System-Wide True In-Car, In-Tunnel Broadband
Application ID #:
Applicant Type:

B. Executive Summary

Herein, we propose an end-to-end network architecture, most notably including a "moving-cell" type model, that provides low latency, high bandwidth and low retransmission rate broadband access for riders of the MTA's subway system. After extensive modeling of the performance, reliability, maturity and cost characteristics of possible designs, we have arrived at a solution based on a combination of existing and emerging technologies that can be readily adapted to meet the unique challenges that the subway tunnels in NYC present. At a high-level, since we are targeting user expectations in the year 2020, we are targeting a per-train bandwidth of 40Gbps allowing us to deliver Internet connectivity to riders via WiFi (802.11ax), LTE-A Pro (3GPP R13) and LTE 5G-NR (3GPP R15) on launch day: January 1st, 2020.

Deploying a system of this scale and capacity into a 24hr active environment requires special logistical engineering and prefabrication techniques in order to avoid significant service disruptions. Specifically, this project will be treated in much the same way as constructing a submarine fiber optic cable with some notable exceptions such as interlinked radios as opposed to optical repeaters. The foundation of this system can be conceptualized as three major components:

1. The NNI or backhaul: This is the in-tunnel and wayside cables and millimeter wave radios that transport the data to and from the trains.
2. The UNI or fronthaul: This is the in-car and in-train roof assembly that is a small network of cables and radios that services the end-user's mobile devices via WiFi and/or LTE and the MTA video cameras.
3. The MAN or metropolitan area network: This is the fiber optic network that connects stations to each other for redundancy, to Internet peering points and Internet provider gateways and to major wireless carrier facilities.

The in-tunnel and above ground wayside backhaul is designed as a prefabricated hybrid optical-electrical cable assembly with a series of strategically placed millimeter wave and, in limited cases, free space optical Remote Antenna Units that using CAD-based tunnel geometry will be tailored for each inter-station tunnel run in a way that allows for the preservation of line of sight in any bend across either axis (corners and changes of grade).

We are an independent research, design and engineering firm, which operates solely at the behest, and within the constraints, of our clients; allowing us the flexibility to pick best of breed technologies and deliver superior designs. We feel that we deserve to win this challenge because we have thoroughly explored the intricacies and pitfalls of designing a communications network of this complexity. From the train speeds, the high bit-rate service, the low packet loss all the way to the lack of clearance for both radio signal propagation and infrastructure attachment, the environment is challenging and unique only to NYC; there is simply no off-the-shelf solution that would produce industry-leading results. Additionally, a project that spans this many technical disciplines, at this scale, would take the typical firm on the order of 4-7 years to accomplish; we can do it in 2. This is as much a technology challenge as it is a logistics and prefabrication challenge and there is no combination that we enjoy more. As NYC residents ourselves, it is not only in our best reputational interest but also in our personal interest to make this a success.

C. Technical Overview

Components and Specifications.

Target Architecture Isolation Process. Some background on how we identified this architecture as the most viable solution is very helpful in understanding the design of the solution. We began by defining, and sometimes inferring, the solution's required operating characteristics, all of which are driven by user experience metrics:

1. UE (User Equipment) Distribution Projection for a service lifetime of 2020-2030
    - WiFi 802.11ax/ay: 45%
    - WiFi 802.11ac/ad: 45%
    - LTE-A: 45%
    - LTE-5GNR: 45%
    - Legacy LTE & WiFi: < 10%
2. Application Types
    - Streaming Video to Device (e.g., Netflix, YouTube)
    - Streaming Audio to Device (e.g., Voice Calls, Spotify)
    - Video Calling (e.g., FaceTime, Duo, Facebook Live)
    - Augmented Reality & Gaming
    - Messaging (e.g., iMessage, WhatsApp, SMS, etc.)
3. By 2020, end users will be accustomed to mobile response rates equivalent to an average wireless connection speed of approximately 200Mbps. Given the loading capacity of trains, and the fact that we will be servicing a "captive audience", we consider any solution that delivers less than 20Mbps per user during peak times as non-viable. Put another way, we are oversubscribing the network at a ratio of 10:1 due to our estimates of traffic patterns. However, as we are designing the fabric of this solution to be modular, we are confident that we will be able to meet or exceed this ratio as we explore upgrade paths beyond the initial release of 40Gbps per-train in 2020 at a fraction of the cost of the initial build.

Using these parameters to set the user experience stage, we were able to eliminate a significant portion of technology options. For example, we were able to eliminate any technology utilizing leaky feeder cables due to their inability to sustain high bit rates for any reasonable distance; even when treated as a MIMO system. We then moved on to tunnel geometry and characteristics to further contain the design to a viable solution:

1. Essentially zero overhead for signal propagation and/or equipment attachment.
2. Plugged-tube environment (steel train in tight tunnel) does not permit signal propagation directly to the next station when there is another train in the path. This will be further exacerbated when CBTC systems bring train spacing closer together, but is still an issue with fixed block signaling.
3. Multiple track configurations: single, double, triple and quad.
4. Maximum restricted speed: 55mph
5. "Rough" tunnel walls: the tunnel walls are not smooth concrete, they are of various construction techniques with decades of attachments and various reflecting and absorbing materials; unrealistic to model.
6. Volume and largest dimension of transceiver device on train and tunnel walls.
7. Avoidance of ISM bands that are utilized by CBTC and 2-way radio: 900Mhz, 2.4Ghz and 5.8Ghz.

With these constraints we were able to reduce the options even further, for example, we eliminated options utilizing radio frequencies below 10Ghz because of conflicts with CBTC and 2-way radio and extreme multipathing effects from the "rough" tunnel sidewalls, which will limit our effective bit rates, and are now left with only Line of Sight (LOS) solutions to consider. In order to narrow these further, we then consider the following:

1. Cost
2. Complexity
3. Harsh Environment Resistance
4. Serviceability & Maintenance Intervals
5. Reliability
6. Service Interruption for Deployment
7. Safety

With these final constraints, we were able to eliminate the remaining contenders for various reasons:

1. Free space optics, although one of our favorites, was mostly eliminated due to beam steering latency and device size with Galvo-based setups, target aperture size and steering angles with 2D MEMS setups and for all setups due to particulate accumulation on the protective transparent coverings. Despite these issues in terms of general coverage and deployment, we will employ FSO-integrated RRUs on select inter-station runs, where the environment is known to be highly controlled as it provides a redundancy where the RF environment is not suitable.
2. Back-to-back (and mesh style) tunnel wayside mounted radio or optical relays because they introduce too much latency due to O-E-O and Forward Error Correction (FEC) algorithms and do not allow for A/B system evolution, increase overall latency, require a power network like the other contenders and create extreme failure points in linear and zig-zag topologies.
3. UAV (quadcopter-based) relay networks were eliminated because flight time (even with a 15,000mAh battery, fuel excluded for fire safety), while carrying and operating the communications payload, would require swap outs every 15-20 minutes that would have to take place in stations resulting in a safety risk as there is very little space to operate a large enough drone in subway stations. It also suffers from cumulative FEC relay latency like any other mesh-based topology. This was a favorite here, but ultimately not fit for purpose.

4. LRUV (Light Rail Unmanned Vehicle) relay networks were eliminated because of relay latency and autonomous safety concerns especially in the areas of stacking and turning. Stacking is a concept in our RATT (Remote Autonomous Train Terminal) design that allows for multiple low-profile RATTs to stack on top of eachother in order to reduce inter-train spacing in areas of congestion. Turning is the case where the RATTs must reconfigure their leader/follower pattern when a train changes directions at the end of a line. This was also a favorite here, but ultimately there were too many risks associated.

And finally, we are left with one viable solution broken down as follows:

1. A hybrid optical-electrical cable assembly that runs on both sidewalls of every tunnel to provide n+1 redundancy.
2. Cable-integrated mmWave (millimeter-wave or mmWave as industry slang often has many meanings, in this context it refers to 20-90Ghz RF), patch-style, phased array, beamsteering and realtime beamtracking Remote Antenna Units (RAUs/RRUs).
3. Optical fiber-based fronthaul to the radios & backhaul over the MAN both to reduce onerous FEC latency, improve overall BER/PER, allow for system redundancy, A/B in-service upgrades, and to provide an avenue for revenue generation and system evolution.

General Design. From a logical perspective, the overall goal of our solution is to seamlessly tunnel an IP/Ethernet network between each train and the core of the network at a high bitrate and with minimal latency and loss. Focusing on presenting both ends (train and core) of these tunnels to each other as IP/Ethernet networks allows for the highest level of vendor interoperability and adherence to open standards. However, the hard work is in the layers in between to make this IP/Ethernet tunnel seamless. When formulating a mental concept of the solution, we find it best to think about a "moving cell", the train itself, as a mobile client connecting to various antenna sectors of a base station that are laid out linearly (tunnels) rather than triangularly. However, because we can guarantee only one "user" (train) per-sector at any given instant, we can do away with all of the latency, jitter and bandwidth limitations commonly found in LTE/3GPP systems. Of course each train is supporting thousands of users/devices, but air-link sharing for them is done via the local small-cells/WiFi APs within the train (UNI), so that the train to wayside link remains a single "user" air link at all times. As visualized logically in figure 1 below, the network has the following key features:

1. Redundancy across all dimensions
    - Train level: 4 front-face and 4 back-face beamsteering phased array radio antennas units with pre-generated steering codebooks.
    - Tunnel wayside: left-side run and right-side run hybrid cables.
    - Station level: both optical transport and DC power plant are n+1.
    - Inter-station level: end of line stations interconnected to form larger metro area rings.
    - Core, both primary and backup central datacenters are located close to the latency centroid of the NYC subway system
2. A/B in-service upgradability: ability to run existing network while upgrading components to support an enhanced network in place.
3. Component-based architecture with standard interfaces.
4. Adheres to established standards wherever possible.
5. Serviceability: A successful design of this complexity must include features that provide for infrequent and rapid repairs.
6. Each train assembly is a mobile cell that is transparently "attached" to the rest of the network via hardware-based VXLAN tunneling. This allows for extremely low-latency and abstraction from the wireless links so that multivendor interoperability is a non-issue.

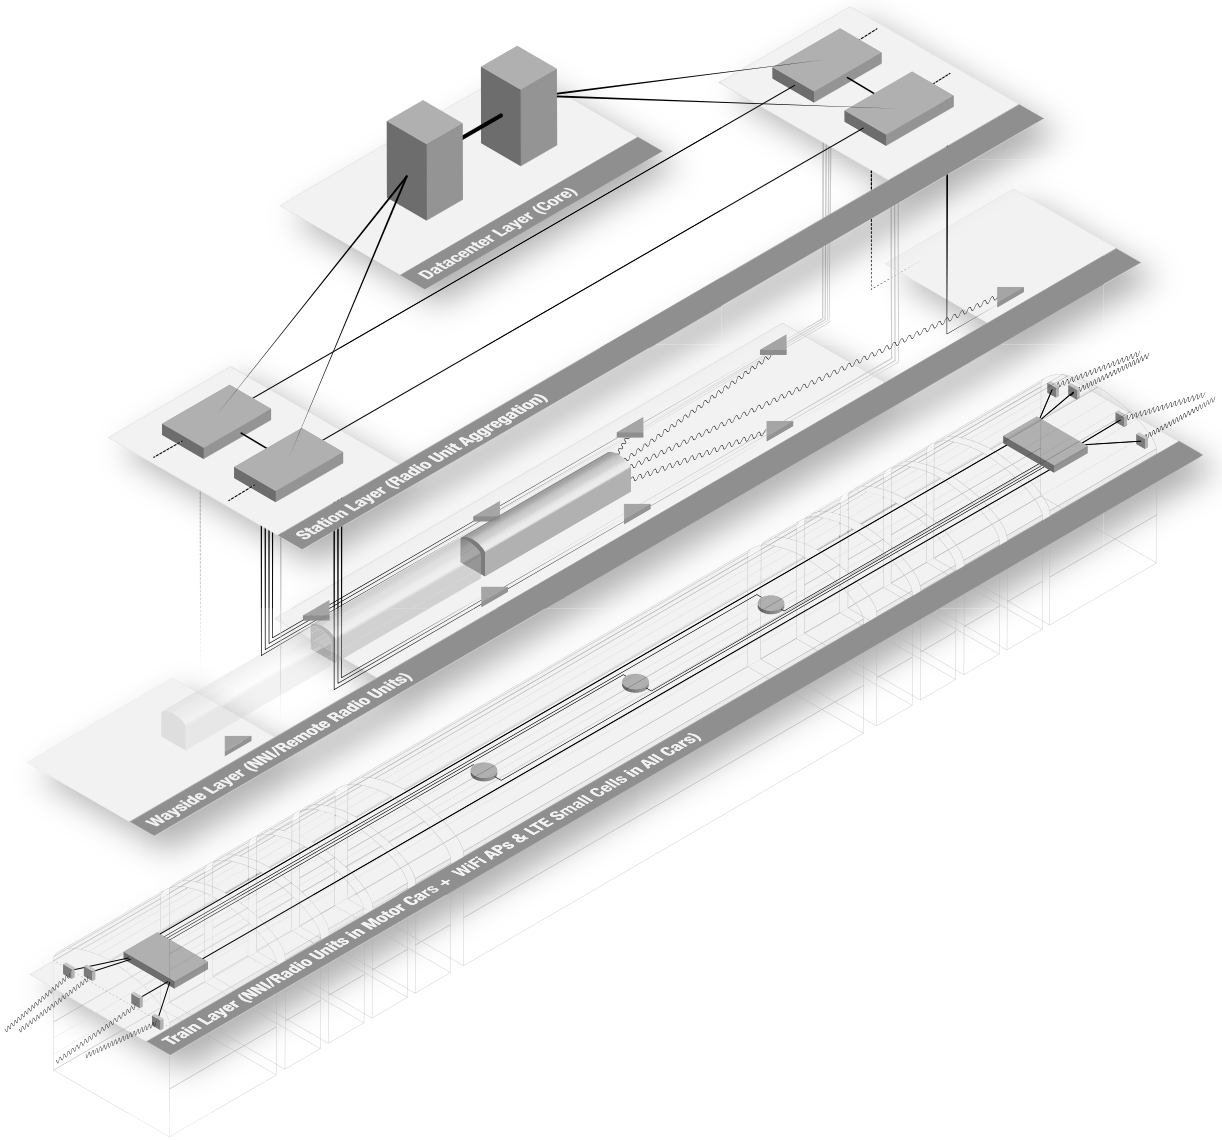


Fig. 1. Logical overview of entire network.

Physical Plant.

1. **Custom Hybrid Electro-Optical Cable,** no larger than ~25mm OD
    - 3 Conductor, Bipolar, Isolated, 380Vdc MTDC microgrid architecture (NEC 2017, Article 712 DC Microgrid compliant)
    - Loose sub-tube single-mode fibers (for cable assembly attached RRU fronthaul)
    - Dense ribbon sub-cable single-mode fibers for MAN and revenue generation
    - Spare conduits for upgrades and revenue generation

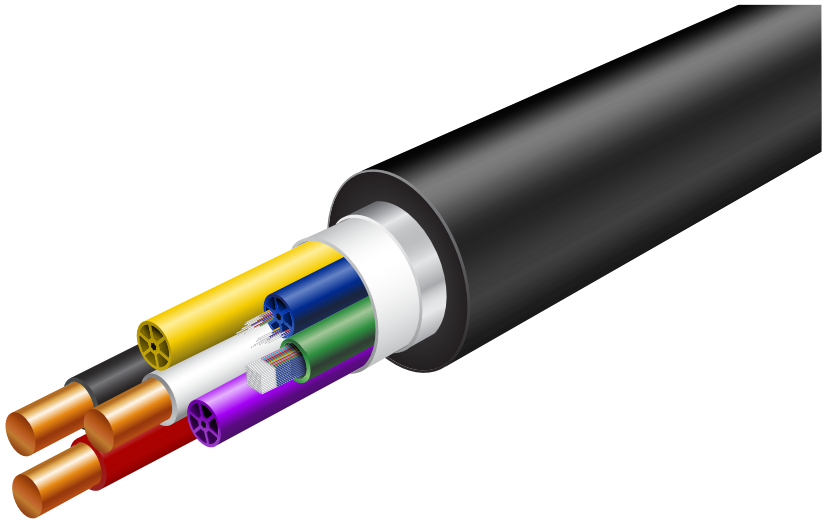


Fig. 2. Cable components visualization.

2. **Power Bus.** The power bus of the system must be designed for optimal availability and extremely limited serviceability, while factoring in the costs of long distance power transmission. Because the network is operating on a distance scale somewhere between a large datacenter and a large span microgrid, we source the components of our solution from those industries. We operated our design process under the assumption that powering the system from the third-rail was not an option. This necessitates the construction of an integrated DC power system as further pollution of the tunnel environment with AC would only contribute to RF noise. Perhaps more importantly than the noise is the goal to avoid in-tunnel construction apart from hanging the integrated cable. Therefore our proposed cable integrates all of the necessary power conversion, breaker and terminal functions as pre-molded, and factory tested, assemblies onto the cable so that wayside cable attachment is truly only installation requirement. With the ongoing standardization of DC microgrids, and the growing push for datacenter distribution via DC bus topologies, we believe this is the most cost-effective route as it allows us to be compliant with established safety codes as well as have the cost advantage of multiple sourcing.

    Grid attached AC-DC rectifiers, with battery backup, will provide redundant power by forming a multi-terminal DC microgrid from the stations on either side of an inter-station span, along both cable paths (tunnel left and right) exiting each direction of a station. The power bus design, between stations, will be based on a bipolar, isolated, 380Vdc, HRMG (High Resistance Midpoint Ground) architecture:

- AC to 380Vdc rectifiers in stations with battery backup
- 8 AWG 3-conductor system (with the exception of runs beyond 3km, where we will employ 5 AWG conductors)
- PLC (Power Line Communications) based O&M interface for breakers and system monitoring
- Standardized PoE++ (802.3bt type 4) interface at DC-DC converters feeding RRUs for out of band communications with the radio units in order to perform diagnostics, patching, hard cycling, etc.
- Optional contingency for price risk on high-efficiency DC-DC converters: 100m spaced, 48Vdc, solid state, cable integrated, DC-DC converters, each able to draw 400W to feed PoE++ (802.3bt type 4) compatible spans spaced on center, and then at 25m intervals down-cable and up-cable for a total of 5. This allows us to provide redundant power to the cable-integrated radio units at all levels while reducing the cost burden of small-scale, high efficiency, DC-DC converters and leaving the window open for a parallel A/B upgrade. Further, by not colocating the high to low voltage converters with the RRUs, we amortize our heat spreading more efficiently.
- With all network elements fully active, the system wide power draw can approach 3MW. However, the RRUs are designed to rapidly transition into low-power mode bringing the average system wide power draw to approximately 600kW (including all in-car UNI, wayside NNI and MAN equipment).

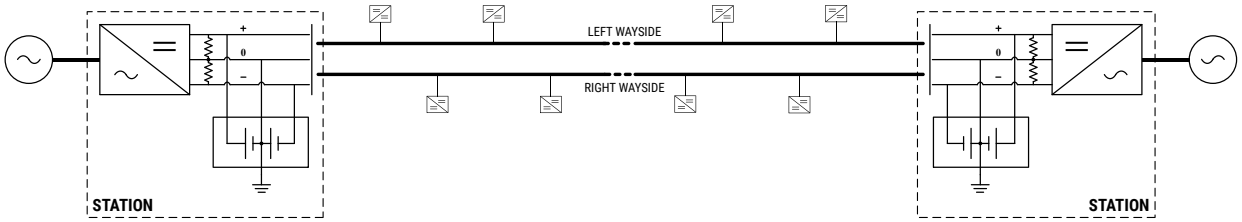


Fig. 3. Block diagram of MTDC microgrid.

3. **Station data closets.** We approached the station data closets from the perspective that there was space and AC grid power, but nothing more. This solution will require minimal space per station for the optical ROADMs, BBUs, fiber management and passives and the power system (rectifiers, voltage converters and battery backup) and will ideally be fed from independent AC power grid feeds. We plan to preassemble the racks, and deliver them to stations via a track-based service car where possible to avoid intensive in-station labor and configuration in the actual data closet. However, power and fiber splicing will need to take place in the station data closets.

Core-to-Station, Station-to-Station & Station-to-Radio.

1. Primary Core Site: 111 8th Avenue
2. Backup Core Site: 60 Hudson Street
3. Direct, passives only, fiber from each station to the core
4. ROADM-style rings (traveling both sides of tunnels) arranged in concentric ellipses starting from single-station spans radiating outwards to a maximum radius of core to end-of-line station. This allows for route-efficient multiple failure protection.
5. End-of-line stations will form larger optical rings with nearest neighbor end-of-line stations via non-tunnel paths to avoid service disruption on core to end-of-station dual cable breaks.
6. Due to the number of Remote Radio Units, we must focus excessively on cost optimization. Unfortunately, we are unable to use TDM-PON (Passive Optical Networks) due to their inherently unacceptable

latency (> 1 millisecond floor). The next best compromise is short-reach colored PIC-based (photonic integrated circuits) optics to the nearest stations.

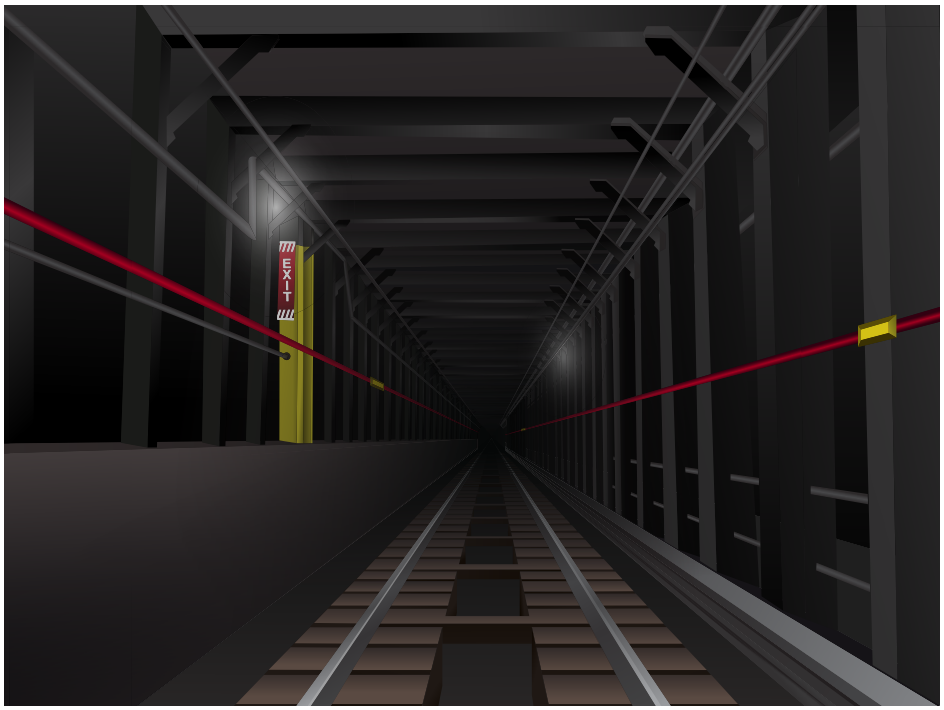


Fig. 4. Rendering of wayside integrated cable assemblies.

Wayside-to-Train. Perhaps the most important part of this solution is the wayside to train NNI. It must support maximum train speeds of 55mph (25m/s) while delivering on promises of true "broadband" connectivity (which we define as servicing 2,000 users per train assembly with a minimum aggregate per-train throughput of 40Gbps). This is not an easy task with today's technology; however, it is most definitely achievable.

1. Combination of mmWave radios, integrated into prefabricated inter-station cable runs, spaced 50m per cable (average). Cables are laid on both side walls of every tunnel with a 25m shift in logical RRU (Remote Radio Unit) "start" point.
2. Assuming average radio reach (pending ray tracing against point cloud or CAD data) will be approximately 100m, the effective bandwidth, redundancy and deadline to repair are all improved.
3. As CBTC rolls out to more NYC subway lines, we must be mindful of the minimum inter-train spacing (e.g., it is no longer assumed to be 1000ft). This could potentially obscure the rear-face radios of the first train and the front-face radios of the second train from having radio visibility to any wayside RRUs. For this reason, we explicitly utilize a train-face mounted, short-range (<100m), low-cost free space optical links between trains when they are in range. This allows the trains to pass signals through each other without an interruption in service. It should be noted that if there are existing "pileup" statistics, and/or CBTC minimum distance floors, then we can engineer tighter RRU spacing in congested areas rather than utilize the fixed spacing to avoid this altogether.
4. Beam tracking algorithm that utilizes an adaptive codebook based on AoA, AoD, prior radio contacts, expected radio contacts, accelerometer input and derived train velocity. This allows for the near elimination of link acquisition latency and a new level of link stability. It adapts to changes in tunnel features such as new equipment on tunnel walls, equipment alignment deterioration, etc.

5. Super low profile side-wall trapezoidal prism-shaped protrusion (6cm), no large diameter antennas or other obstructions/attachments. (figure 4)
6. Continuous connectivity by way of multiple active links at all times; only a maximum of 2 out of 8 NNI train-side radios are ever involved in a handoff at any point in time.
7. DFS (Doppler Frequency Shift) mitigation algorithms utilizing known information on train velocity, front/rear-face approach vs fade, along with Direction of Arrival (sourced from antenna array elements)

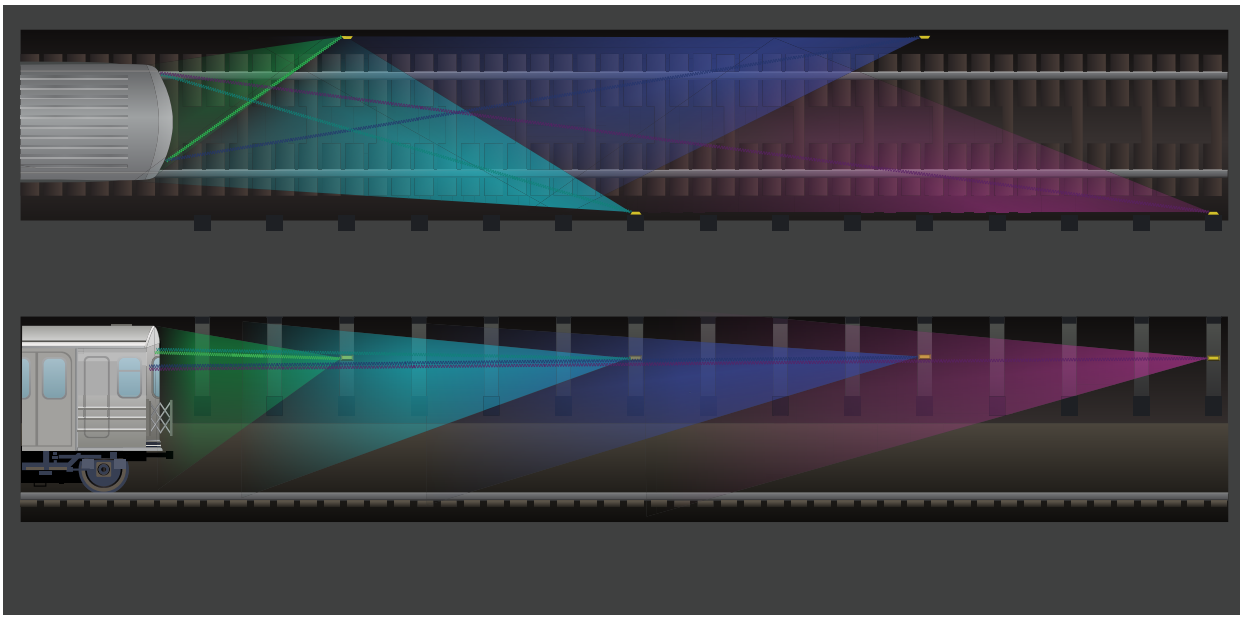


Fig. 5. Visualization of azimuth and elevation beam steering (not to scale).

Since there are many layers to the NNI radio stack, we will devote a substantial portion of the allotted space to the wayside Remote Radio Units (RRUs) and the front and rear train-mounted radios (TMRs) since they are technologically identical. We will use a network of cable attached Remote Radio Units with an integrated splice box that molds around the hybrid opto-electrical cable. These RRU's are mmWave, point-to-point modular units with integrated phased array, beam-steering, and beam-tracking antennas. They are pre-installed, verified and tested onto pre-engineered lengths of cable for deployment in the tunnels without specialized labor and without service disrupting or time-consuming occupancy of a tunnel. Because the radio topology is effectively a linear series of point-to-point (PtP) connections, we utilize 4 front-face and 4 rear-face train-mounted radio units interfacing, at up to approximately 100m of distance, with wayside, cable-integrated, RRUs at 50m effective spacing per tunnel side. This gives us an effective mutual (sides of tunnel) wayside RRU spacing of 25m, but with each radio covering a service interval of 100m. This is an important distinction because at speeds of up to 55mph (25m/s) we want to maximize the time between RRU acquisition and subsequent passing of this RRU in order to minimize handoff while at the same time ensure enough continuous coverage to maintain access and sufficient bandwidth even during RRU failure and replacement cycles. At top speed, each RRU will service a train for less than 4 seconds (although it should be noted that the average speed of the entire subway system is approximately 17mph). While the interval may sound like a small number, 20GB of data can be both transmitted and received to the train during this interval.

This brings us to the most important aspect of all: Handoff. Handoff latency and service disruption is one of the largest concerns in a project such as this. We are intentionally ambiguous as to the delivered "layer split", as there are many ways to deploy a massive RRU/RRH (sometimes called a C-RAN topology). The

reason we are ambiguous here is solely due to the cost of commodity chipsets at the time of ordering, but we are prepared to support the most cost efficient layer-split, which we will discuss later in this proposal. In any case, since we are constrained to line-of-sight (LOS) communications due to the tunnel geometries and bitrate requirements, we can establish that the RRU spacing (in the worst case) is unchanging at 50m along an inter-station run, both sides, alternating, giving an effective 25m spacing. Using this, we can estimate that we will deploy a maximum of 45,000 RRUs. At this scale of active, in-tunnel network elements, the overlapping redundancy is critical as failure rates, even at 1%/month, represent 450 RRUs and those failures should be seamlessly covered by other RRUs long enough for a service window, which may only be practical once per month to avoid impacting train service schedules. Fortunately, in this design, the RRUs are modular blocks that snap into and out of the hybrid cable, not requiring skilled labor (e.g. splicing, electrical, RF, etc.) to effect the swap of a failed unit with a new one.

As mentioned earlier there are multiple actionable layer-splits that pertain to deploying RRUs/RRHs in this scenario, where the RF (or IF) signal arriving and departing from the RRU is either transported to a more central location via fiber (pure RoF) to be processed in a BBU (baseband unit) or it is processed partially in the RRU (CPRI/eCPRI) or it is processed completely in the RRU into a packet/frame stream before transmitting via backhaul (IP/Ethernet). All of these schemes have pros and cons at high bitrates, but due to the volatile cost structure of the components involved, we provide a solution for seamless handoff under all splits.

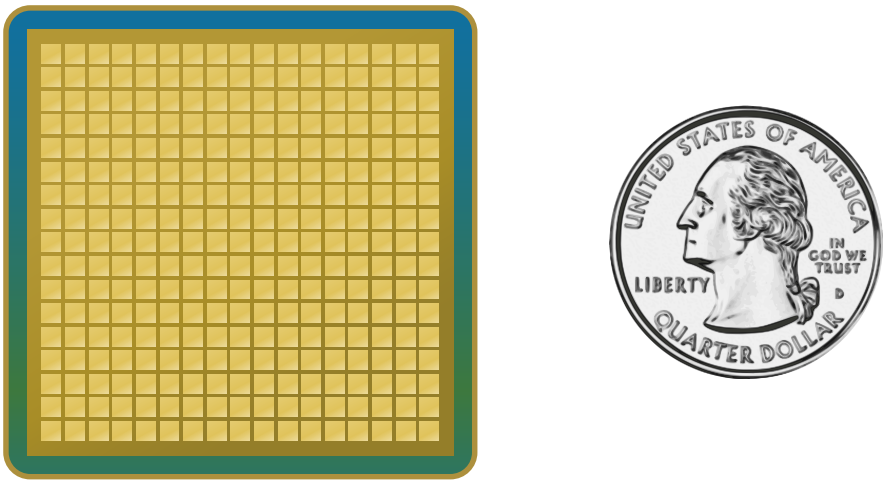


Fig. 6. Illustration of phased array antenna size.

There are 4 antennas per train-face operating in support of each other. On an intra-face basis, they are providing for seamless handoff and larger bandwidth utilization (MIMO arrangement). On an inter-face basis (front and rear), however, they are providing for diversity, meaning that we are always transmitting (and receiving) two copies of every packet/frame (PDU); one through the front-face and one through the rear-face.

We employ a scheme very similar to the PDPC-layer in LTE/3GPP so that the first PDU, in sequence, to arrive "wins", but for the purpose of not only ordering, but also diversity. While this might seem excessive, the biggest quality of service degradation in any communications system is packet loss as it either necessitates retransmission or causes dropout.

Back to the question of handoff, in the case of RRUs that are fully integrated APs, we will be implementing a layer 3 lossless handoff that is capable of draining the entire upstream queue potential (data in flight: cables and switching buffers and delayed by FEC algorithms) towards an RRU about to be passed (leave service) before it loses its layer 1 (radio) link with a train. It should not be underestimated how much data is actually "in-flight" throughout the network at any instant at high bitrates. We must calculate the OTDR-based fiber optic lengths as well as sum the switch-port queues, the various FEC buffers (delays) in a given path plus add in implementation specific delays such as radio modulation overhead in order to determine when an RRU, that is about to lose its link due to fading (or is peaking because it is about to be passed) should transition into queue draining mode only. This, of course, will be done in realtime and satisfy that no in-flight data is lost during handoff. This is a hybrid of user initiated and network initiated handoff that is pre-trained and adapts to changes over time.

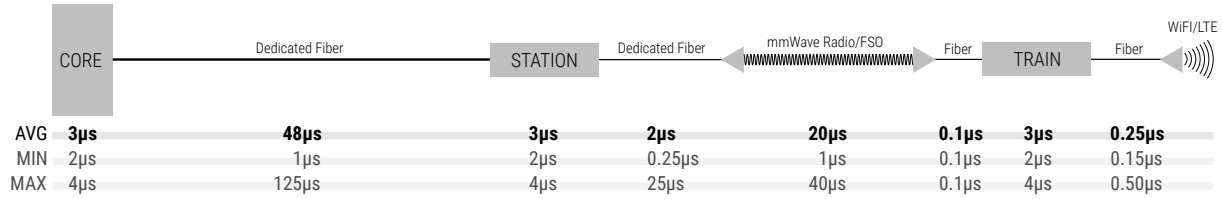


Fig. 7. Visualization of latency path.

Latency, Jitter, Multipathing, Reordering and Handoff. Excluding the end user device itself and in-car access layer (WiFi/LTE), which are shared air link mediums, the entire system is designed to provide a one-way average IP-level latency from train to network core (8th Ave. & 14th St) of less than 80µs and a worst case latency of less than 200µs (A-line, Far Rockaway). This is a magnitude better than existing services that users are accustomed to. For example Verizon FiOS's Gigabit service exhibits a one-way latency of approximately 1.5ms, 15 times worse, due to the use of a TDM-based PON architecture. Because this network is designed with point-to-point radio links that are immediately transitioned, in-tunnel, to a fiber optic network, there is, at most, a 100m air interface that is not divided by way of time slots; they are full-duplex, dedicated links, from wayside to train. This reduces the question of jitter over the air interface to a function of frame/packet size, modulation technique, channel bandwidth and FEC which can produce different resulting latencies thus creating jitter. In the most extreme cases it will be no more than 40µs, but less than 1µs for the average packet/frame. But again, to give perspective, we're talking about microseconds here and are not remotely close to a millisecond. In terms of the non-air interface, we try to reduce the number of O-E-O (Optical-Electrical-Optical) and switching decision points such that there is only one station-hop in the path. This station hop is a full "layer 3" Ethernet/IP switch with a latency range of 2-4µs. The reason for not optically backhauling each individual RRU is due to the cost of optics and fiber in a setting where even if all rolling stock was running, it would only cover approximately 10% of the revenue track. Therefore it is more efficient, and a negligible latency hit of 2-4µs to aggregate at the station level. Further, because all paths through the MAN are undersubscribed, there is never any input or output queuing at the switch ports thus allowing for constant latencies and jitter at the MAN layer.

Each train has two logical faces (front and rear), each with 4 active radio links at any time. In order to

eliminate end user traffic disruption during handoff, we must maintain continuous connectivity at all times as the train travels down a section of track. In order to avoid disrupting end user packet flows, we must design the system such that we don't drop packets or reorder them by way of multipathing effects or handoff. In order to accomplish this, we must rely on an intelligent handoff mechanism. There are a few very important details on how this solution achieves a low-latency, zero disruption handoff:

1. Only 2 out of 8 radios are ever involved in a handoff at any point in time, traffic is routed around the radios that are entering a handoff cycle.
2. Wayside RRUs are connected to the same aggregation node at a particular station to avoid any common case routing updates (macro-cell handoffs) under normal train operations. Although we do this very quickly (< 25ms), we optimize to avoid it by way of forcing these macro handoffs (a.k.a. routing updates) to take place while the train is at rest in a station.
3. Acquisition and beam alignment to a successive RRU is not a global elevation and azimuth search, it is an evolving codebook (industry slang for lookup table or mini-database) that learns and adapts. This lowers the latency of RRU acquisition to about 15ms thus increasing the useful service span-time/distance for each RRU.
4. Establishing MAC-level (layer 2) link takes an additional 10ms.
5. Convergence of the entire network in the event of a non-common train path (e.g. train skipping a station, unscheduled reversal, etc.) requires 25ms, but during this, we do not drop packets nor reorder them.
6. The in-train switch forwarding tables (FIB), for radios that are about to enter a handoff, are updated before the radio goes offline to send traffic to the most recently acquired wayside RRU. This takes approximately 1ms and does not disrupt traffic.
7. When an RRU is about to be passed (or faded out of service), inbound routes to this radio are withdrawn so that flows are redirected to one of the other 7 radio links before the current RRU is passed. We begin this process with ample time (several milliseconds) before the handoff so that no packets are dropped and all upstream queues and data in-flight is drained.
8. Packet reordering is avoided entirely in this solution, but careful attention should be paid to why that is. Throughout the entire MAN, while multipathing is employed for load-sharing and redundancy, a flow will only ever traverse a single path, so there is no issue there. However, when traveling towards the core there is a scenario where, because the speed of light through fiber is 30% slower than it is through air, that a newly established RRU link could have a worst case 383 meters (new RRU is 100m towards core from train's front face, train is 183m long, farthest tail end RRU that is handing off is 100m from train rear face) worth of latency improvement over the tail end of the train's furthest RRU. This translates to about 600ns of gain where packet reordering could take place during a radio handoff (i.e. flows on the farthest link are transitioned to the closest link). To put the 600ns in perspective, the user would have to be engaged in a flow approaching 500Mbps in order to experience a reordered packet, but regardless, we prefer to avoid it entirely. For this reason, we restrict the flow handoff to the same train face with the exception of complete radio failure on an entire train face. This reduces the maximum latency gain to 100 meters of distance, which is roughly 100ns in gain of air vs fiber. A reordering event at this scale would require an application flow in excess of 1Gbps which is an unlikely scenario given the access interface (WiFi & LTE), but still, we eliminate it as follows: when traveling towards the core, RRUs are fed from the departing station. This forces packets to travel over the fiber MAN slightly beyond the train (from the perspective of the core) and then towards the RRUs in a manner that is "chasing" the train so that no gain when switching to an air path that is physically closer to the core will result. While PDPC-layer buffering and reordering can easily correct, it is always best practice to eliminate reordering at the physical layer where possible.

UNI: Train-to-User. The UNI is a very dense and RF-noisy environment. For this reason, we focus on keeping all user signals within the train and provision sufficient access points within each car to maintain a quality of service that will truly be usable. For an outsider, it may seem excessive that we choose to deploy 3 access points per-car, but when it comes to CSMA (Carrier Sense Multiple Access) frequency sharing MAC protocols like WiFi (802.11), there are real life limits that aren't often reflected in specification sheets.

Inter-car and inter-consist network linkage is fiber-based and exterior linkage cabling consists of IP67-rated housings over MTP/MPO-type connectors. While we considered FSO linkage, we determined that the cost-benefit was not reasonable due to the relatively static nature of consist linkage. We source power from the existing, in-car, DC-DC converters on a per-car basis and condition it by way of passing it through a space efficient power conditioner and battery backup that feeds a small PoE++ distribution network to each LTE and WiFI access point. We pay the cost of standardizing to PoE++ to increase bidding sources and decrease customization for access point providers.

As mentioned prior, we utilize short-range FSO (Free Space Optical) transceivers on front and rear-faces of train assemblies to account for unplanned, extremely tight (<25m), congestion spacing. These FSO links allow an NNI-isolated neighboring train, devoid of NNI access, to tunnel through the front and rear blocking trains' NNIs. Although the sharing compromises bandwidth, it is designed only as a backup for extreme congestion.

In order to remove the requirement to deploy WiFi "concentrators", we utilize low-latency, hardware-based VXLAN tunneling from all trains to the core of the network so that we are able to provide the highest quality mobility experience possible via WiFi. In this arrangement, sometimes referred to as local switching, control traffic may traverse "CAPWAP" style tunnels, but all data plane traffic is switched (and subsequently routed) directly. While this is not a conventional deployment concept today, it yields the highest quality of service, lowest-latency AP handoff and lowest latency data plane by allowing for an access point deployment that mimics many APs having visibility to the same layer 2 network. We are unaware of any seamless WiFi mobility deployment extending across the distances involved here and caution the MTA against any vendor proposing to use concentrators and/or home-registration tunneling/forwarding. With this as a base infrastructure we are able to avoid vendor lock-in and potentially pursue a multi-vendor deployment as a price and technical risk mitigation strategy.

Key features of the user access network:

1. 802.11ax WiFi Access Points
    - Backwards compatible with legacy 2.4Ghz and 5Ghz 802.11 access protocols
    - 3 APs per car, localized MU-MIMO
    - Hotspot 2.0 Release 2 (WPA2-EAP / 802.11X) for regular riders
    - Open Network also available for enrollment and ease of connectivity, but will be subtly discouraged for repeat riders due to security risks
    - We will pressure Transit Wireless to upgrade their "concentrator" technology to support the Hotspot specification concept so that users can seamlessly roam when embarking/disembarking and enjoy better WiFi security.
2. LTE-A/5G-NR Multi-Carrier (Multi-Operator) Small Cells
    - In-car multi-tenant small cell platform (supporting multiple carriers)
    - In order to minimize O&M for the carriers, we will present each entire train as a single base station (eNB) and each car as a sector. Careful attention must be paid to frequency overlap in open gangway cars.
    - Operating as virtual cells within a "moving-cell" with the important distinction that the backhaul of the traffic from the small cell in the train to the carrier network is via IP and the fact that it is in a moving platform is transparent to the carriers

- Provide IEEE 1588v2 (PTP) timing source from the NNI

MTA Net. The MTA Net is a private, partitioned, network space for low-latency backhaul of high bit rate data. While not a core objective of this challenge, it is available in the margins of the design. The key features of the MTA Net are:

1. In-Tunnel Fixed Point Camera Backhaul Bandwidth (NBase-T with PoE++ terminals spaced at 25m throughout tunnels)
2. In-Tunnel WiFi connectivity for worker connectivity, autonomous light rail vehicles, sensors, etc.

Internet. This is not an easy network to engineer at any level and will require carrier class Internet connectivity. Fortunately, NYC is one of the easiest places to implement such connectivity.

Assuming the population of NYC is 8.5m, and the MTA has a daily subway ridership of 6m, and assuming 40% of riders being round-trip, the MTA will be close to, if not, the largest ISP in NYC with 3.6m potential unique Internet users/day. With a high quality service, more users will be active and, during rush hour, the bandwidth requirements to/from the Internet will be large. The network will interconnect (peer) with the largest sources of traffic (Google, Facebook, YouTube, etc.) at 111 8th Ave. and 60 Hudson St. and for all other sources, purchase transit at these locations from major carriers. It is important to note that at this scale of Internet operations, peering not only reduces operational costs significantly it also increases user experience as there are no 3rd party networks in the path.

Bandwidth requirements estimation for this solution are somewhat different than conventional estimation techniques because of the unique circumstances that the users are in: they are riding on the subway and are therefore a captive audience. Therefore it is likely that we will experience something close to, if not exceeding, a 90% rider to user ratio. What this means is that when engineering for worst case scenarios, i.e. rush hour, we need to assume an active user base of approximately 1m. Now, statistically, not everyone is streaming the morning news broadcast or their favorite YouTube star, many are reading and responding to emails which is a low-bandwidth service, but regardless, we're still on the terabits/sec scale.

Key features:

1. Primary Site: 111 8th Avenue
2. Secondary Site: 60 Hudson Street
3. Intelligent peering with large traffic sources/destinations (Google, Netflix, etc.)
4. Large capacity transit from "tier-1" service providers
5. CG-NAT: Due to the exhaustion of IPv4 address space, it is not likely that we will be able to obtain an IPv4 allocation large enough to satisfy the concurrent user base. Therefore we will have to employ a carrier grade NAT (CG-NAT) solution for IPv4 destination access. This should not be a concern, it is a well established practice that when done properly will not degrade services in any way. As the world moves closer to IPv6 support everywhere, this will become less of a concern.

D. Implementation Detail

1. Our solution in whole has never been deployed commercially because it is highly specific to the topology and constraints of the NYC subway system. However, close to the entirety of the technology behind our solution has been in widespread use for years and is well proven. As stated earlier, our solution is highly modular, with proprietary bits isolated to very specific areas. For example, our solution to packet reordering, without blocking during handoff, may seem proprietary, but can be thoroughly verified in a lab environment. However, there are many areas where significant verification is required. As an example, when selecting antenna systems, we must rigorously verify beamsteering latency under real world conditions. We generally see pure R&D as a separate activity from verification,

although verification at the unit or system level is integral to system performance and often requires R&D support.

2. The resources that we require to fully verify the entire solution before proceeding to ordering/manufacturing are divided amongst two areas:
    1. Lab Resources (Unit Testing)
        - We will open a dedicated R&D space for this project with sufficient space for radio component testing and cable assembly prototyping.
    2. Proving Ground Resources (Train, Track & Tunnel)
        - We will conduct first-phase post-lab testing of the train mounted radios on modified vans in a controlled race track setting. We will initially lease track time from nearby tracks while we construct a long-term evolution testing facility tailored to the MTA subway tunnel topology.
        - For in-car mounted equipment, we will request the MTA to provide a train car, with enough operational capacity to operate the doors in order to design and verify the required WiFI and LTE antenna signal propagation attenuation during door opening/closing events and to model congestion conditions.
        - For interior ceiling and front/rear-face assembly testing and installer training, we will require fully operational test cars throughout the project.
        - We will ask for the usage of decommissioned tunnel sections to verify the results of point cloud and algorithmically designed parameters.
3. Throughout the design of this solution, we have been tracking implementation requirements to be in compliance with FCC regulations with regards to: licensing of spectrum, emission limitations and E-911 (both handset-based and network-based) requirements. Further, for the portions of general railway safety standards and additional code regarding fire/smoke safety, we will be sure that all aspects are in compliance.
4. The general key risks center around two areas: **underspecified conditions** and **cost slippage** due to vendors becoming aware of the project.
    1. **Underspecified Conditions:** are areas where limited pre-project information can prevent a full-scale model and simulation from being built. These can arise from out of date or inaccurate schematics (CAD/GIS data), misrepresented performance parameters from vendors and precision of implementation (e.g., poor tradesman quality or supervision). So while we cannot guarantee that there will be no issues, we always strive to ensure that nothing is missed. It is important to remember that part of the design is that we only plan to make one, 2mph, construction pass through each tunnel section, therefore it is paramount in our design that we have accounted for conditions with the utmost precision as the costs to revisit tunnel sections will be unreasonable.
    2. **Cost Slippage:** Careful diligence with multi-sourcing agreements, labor planning and efficiency and the choice of standards based technology where possible will help to mitigate the risk of cost slippage, but there will be cases where a vendor may be in a position to squeeze the project. Our best mitigation for these cases is to utilize a shadow-project for all of our sourcing and discovery agreements. For example, we would approach vendors under the premise that we are building a municipal network for the city/state but that budget has not yet been committed. Once contracts and pricing are favorable, we will pivot.

    Additionally there are several key **technology risks** that are important to point out:
    1. **Mechanical vibrations** that may cause equipment fatigue or failure. In order to mitigate this, we plan to conduct hyper-time mechanical simulation of all parts meaning that we will stress components under excessive vibrational forces 24/7 for 6 actual months, but with stresses that

translate to several years of real world exposure. This allows us to identify simple fabrication issues such as weak solder joints all the way to long-duration effects such as gross alignment issues.

2. **Track abnormalities** that cause uncorrectable beam alignment jumps. This is relatively easy to correct for, but it is an important risk that we must account for. In the early stages of the project implementation, we will be passively collecting these statistics from actual train-mounted sensors operating under real world conditions. Upon review of the data sets, we can either widen our beams, decrease our beamforming tracking latency, include mechanical dampers in the train mounted antennas or provide active tracking corrections based on the learned track conditions. Again, a risk, but an easy one to address.
3. **Wireless spectrum availability vs cost of components.** Our design is modular and can operate in various mmWave frequencies and we will negotiate with suppliers until the "last minute" of lead time, but there is always a risk that the most cost effective components are not well aligned with FCC spectrum licensing exemptions. Our recourse is threefold:
    1. Utilize unlicensed mmWave spectrum only (60Ghz), which while currently penalized for poor attenuation may soon become the front-runner as component costs come down in preparation for 802.11ay.
    2. Allow our costs to increase.
    3. Reduce our bitrates/energy levels to work around imposed restrictions. In the event that we must reduce our bitrates, we still forecast that in a worst case scenario we will have no problems achieving a minimum of 10Gbps per-train assembly, although we really do wish to achieve the target minimum of 40Gbps.

5. In order to accomplish a rollout of this magnitude and complexity in a 2-year window requires a very well planned strategy that separates the workload into parallel efforts with well defined, friction-free, interfaces. We will need the MTAs unwavering support in several areas to meet this aggressive deadline. The implementation schedule consists of us forming the following design and engineering divisions and moving them along in parallel (see project plan for details):
    1. System Simulation Division
        - Tunnel and station modeling. We are operating under the assumption that 3D CAD and GIS data is current and readily available to us for both manual work and so we may algorithmically run simulations against the as-built tunnels and stations.
        - Point cloud collection: High resolution LIDAR mapping of tunnels utilizing a service vehicle such as the geometry train. A LIDAR unit can easily be retrofitted here.
        - Network behavior simulation and BER/PER estimates.
    2. Station Plant Division
        - DC Power Plant
        - Optical Cable Management
    3. NNI/Wayside Division
        - Hybrid Cable Design Group
        - Radio Design Group
    4. MAN/Backhaul Division
        - Optical Transport Group
        - Fiber Leasing Group (for end-station loop closure)
    5. In-Car Network Division
        - WiFi Group
        - Carrier Group

6. Internet Division
    - Access Group
    - Backbone Group
    - Peering & Transit Group
7. Operations and Management Division
    - Network Monitoring and Management Software
    - Inventory and Repair Dispatch Software

6. We require the MTA to be fully engaged with us for the duration of this project as the engineering of this solution requires intimate knowledge of the subway system. Specifically we need the following supporting data elements as well as dedicated liaisons enumerated below:
    - Station, Tunnel and Train 3D CAD models
    - Tunnel GIS data
    - Station survey site access liaison
    - Tunnel survey access liaison

E. Cost Detail

1. We estimate that the complete costs will be between $120m and $185m, however, there are many factors beyond our control (e.g., the price of copper at any given point in time, requirements to use specific labor pools, etc.). We ask that you consider a reasonable margin of variability either in terms of materials cost or in terms of time to completion (or a combination of both) so that we can either time the supply and demand pricing or keep our timelines tight. Please note that we will be able to narrow the margin of error on the estimate as soon as we have access to tunnel, station and train schematics.
2. The major cost components are:
    - Hybrid Inter-Station Cable: ~2,100km of hybrid cable @ $15k/km => $32m
    - Remote Radio Units: ~45,000 @ $1,000/ea => $45m
    - Station Switching/Routing/BBU/ROADMs: $15m
    - DC Power System: $12m
    - In-Train Systems:
        - WiFi APs: $14m
        - Multi-tenant Carrier Small Cells: $18m
    - Primary & Backup Datacenters
        - Internet Edge Switches/Routers: $6m
3. The Budget Summary and Project Plan are included below with key milestones indicated by red diamonds.

Budget Summary

	2018					2019					
	Q1	Q2	Q3	Q4	CY Total	Q1	Q2	Q3	Q4	CY Total	Grand Total
SYSTEM SIMULATION	**16%**	**12%**	**12%**	**12%**	**53%**	**13%**	**11%**	**11%**	**11%**	**47%**	
Labor	250	250	250	250	1,000	250	250	250	250	1,000	2,000
Materials & Equipment	150	50	50	50	300	75	25	25	25	150	450
System Simulation (Total)	**400**	**300**	**300**	**300**	**1,300**	**325**	**275**	**275**	**275**	**1,150**	**2,450**
STATION CONSTRUCTION	**2%**	**2%**	**19%**	**19%**	**42%**	**19%**	**19%**	**19%**	**1%**	**58%**	
Labor	350	500	750	750	2,350	750	750	750	250	2,500	4,850
Materials & Equipment	100	150	4,500	4,500	9,250	4,500	4,500	4,500	50	13,550	22,800
Station Construction (Total)	**450**	**650**	**5,250**	**5,250**	**11,600**	**5,250**	**5,250**	**5,250**	**300**	**16,050**	**27,650**
INTERNET	**4%**	**4%**	**5%**	**5%**	**17%**	**34%**	**34%**	**8%**	**7%**	**83%**	
Labor	250	250	250	250	1,000	250	250	500	500	1,500	2,500
Materials & Equipment	50	50	150	150	400	2,500	2,500	150	50	5,200	5,600
Internet (Total)	**300**	**300**	**400**	**400**	**1,400**	**2,750**	**2,750**	**650**	**550**	**6,700**	**8,100**
STATION-TO-STATION NETWORK	**1%**	**1%**	**1%**	**1%**	**3%**	**79%**	**12%**	**4%**	**2%**	**97%**	
Labor	100	100	100	100	400	250	250	250	250	1,000	1,400
Materials & Equipment	50	50	50	50	200	15,000	2,000	500	100	17,600	17,800
Station-to-Station Network (Total)	**150**	**150**	**150**	**150**	**600**	**15,250**	**2,250**	**750**	**350**	**18,600**	**19,200**
TUNNEL-TO-TRAIN NETWORK	**1%**	**4%**	**30%**	**30%**	**66%**	**30%**	**1%**	**1%**	**1%**	**34%**	
Labor	100	250	250	250	850	250	100	100	100	550	1,400
Materials & Equipment	50	150	3,000	3,000	6,200	3,000	50	50	50	3,150	9,350
Tunnel-to-Train Network (Total)	**150**	**400**	**3,250**	**3,250**	**7,050**	**3,250**	**150**	**150**	**150**	**3,700**	**10,750**
REMOTE RADIO UNITS	**2%**	**2%**	**2%**	**2%**	**8%**	**29%**	**29%**	**29%**	**5%**	**92%**	
Labor	600	600	600	600	2,400	300	300	300	300	1,200	3,600
Materials & Equipment	250	250	250	500	1,250	12,500	12,500	12,500	2,000	39,500	40,750
Remote Radio Units (Total)	**850**	**850**	**850**	**1,100**	**3,650**	**12,800**	**12,800**	**12,800**	**2,300**	**40,700**	**44,350**
HYBRID CABLE	**1%**	**1%**	**32%**	**32%**	**64%**	**32%**	**3%**	**0%**	**0%**	**36%**	
Labor	100	100	100	100	400	100	750	100	100	1,050	1,450
Materials & Equipment	100	100	10,500	10,500	21,200	10,500	250	50	50	10,850	32,050
Hybrid Cable (Total)	**200**	**200**	**10,600**	**10,600**	**21,600**	**10,600**	**1,000**	**150**	**150**	**11,900**	**33,500**
IN-CAR NETWORK	**1%**	**1%**	**2%**	**2%**	**6%**	**52%**	**40%**	**1%**	**1%**	**94%**	
Labor	350	350	350	350	1,400	700	700	350	350	2,100	3,500
Materials & Equipment	50	50	250	250	600	18,000	14,000	100	100	32,200	32,800
In-Car Network (Total)	**400**	**400**	**600**	**600**	**2,000**	**18,700**	**14,700**	**450**	**450**	**34,300**	**36,300**
O&M SYSTEMS	**13%**	**13%**	**13%**	**13%**	**50%**	**13%**	**13%**	**13%**	**13%**	**50%**	
Labor	150	150	150	150	600	150	150	150	150	600	1,200
Materials & Equipment	50	50	50	50	200	50	50	50	50	200	400
O&M Systems (Total)	**200**	**200**	**200**	**200**	**800**	**200**	**200**	**200**	**200**	**800**	**1,600**
RESOURCE ONBOARDING	**13%**	**13%**	**13%**	**13%**	**50%**	**13%**	**13%**	**13%**	**13%**	**50%**	
Labor	100	100	100	100	400	100	100	100	100	400	800
Materials & Equipment	20	20	20	20	80	20	20	20	20	80	160
Resource Onboarding (Total)	**120**	**120**	**120**	**120**	**480**	**120**	**120**	**120**	**120**	**480**	**960**
QUARTERLY TOTAL (%)	**2%**	**2%**	**12%**	**12%**	**27%**	**37%**	**21%**	**11%**	**3%**	**73%**	
QUARTERLY TOTAL ($ in 000s)	**3,220**	**3,570**	**21,720**	**21,970**	**50,480**	**69,245**	**39,495**	**20,795**	**4,845**	**134,380**	**184,860**

Fig. 8. Budget Summary by Category and Time.

Project Plan

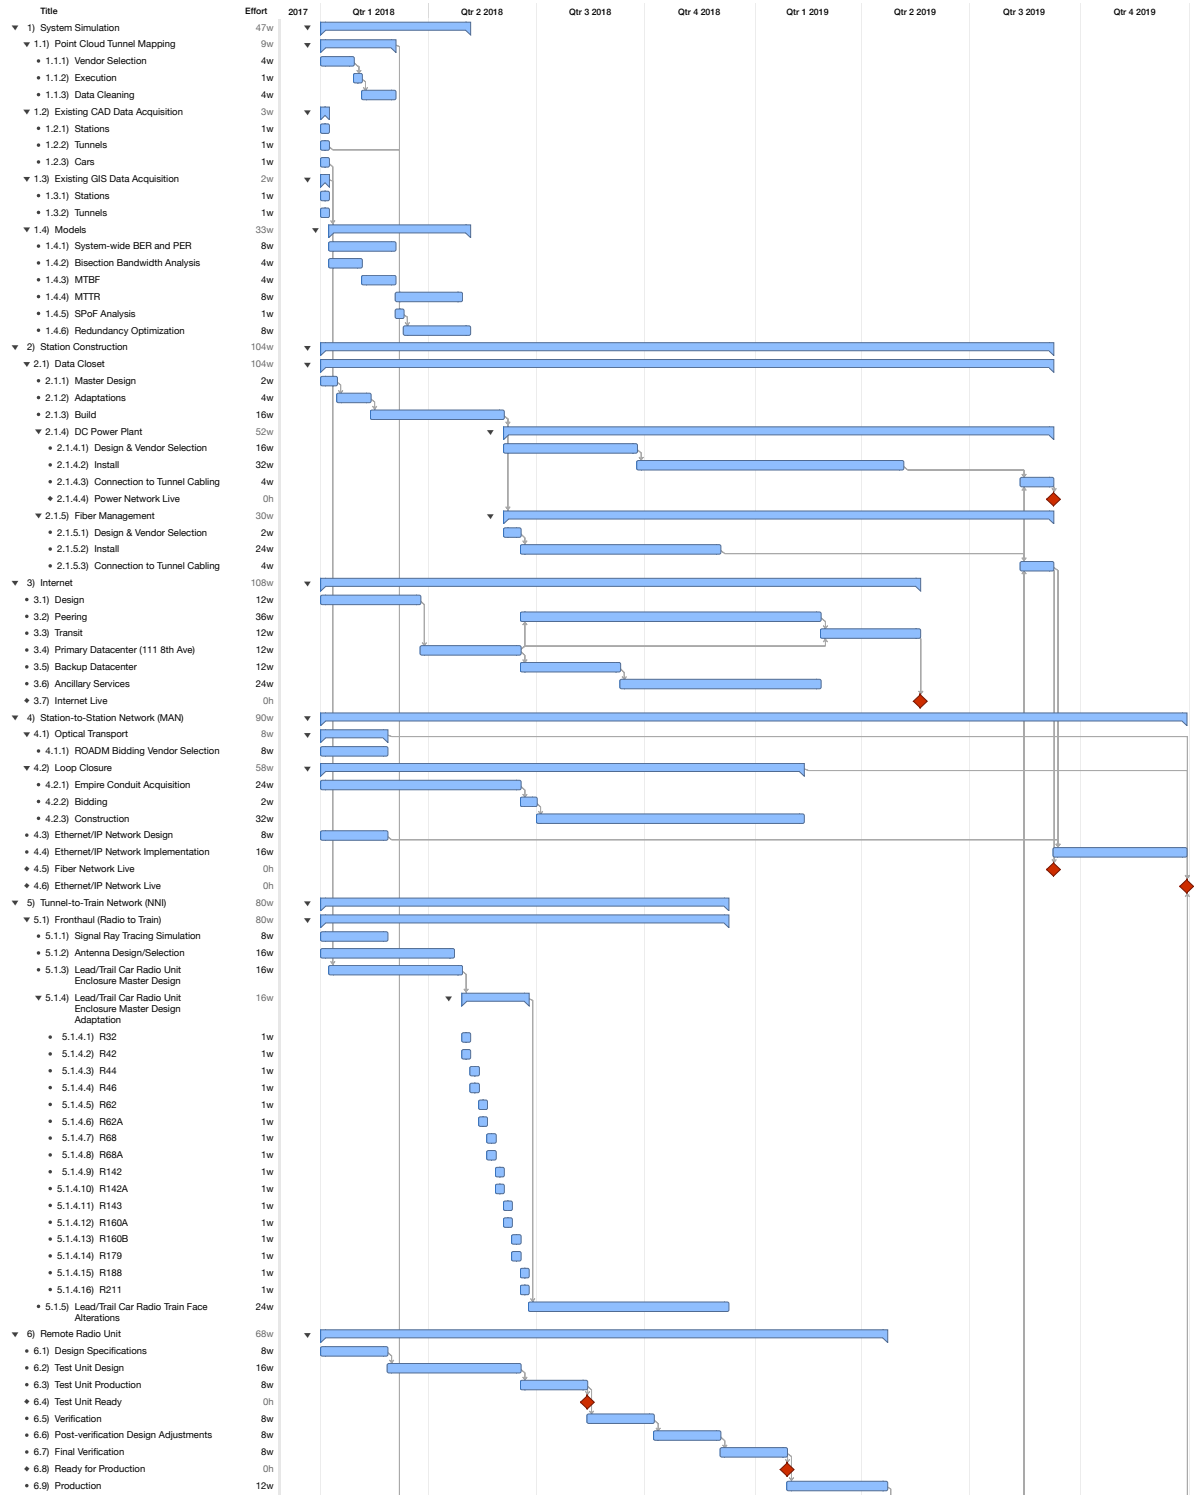


Fig. 9. Project Plan 1/3.

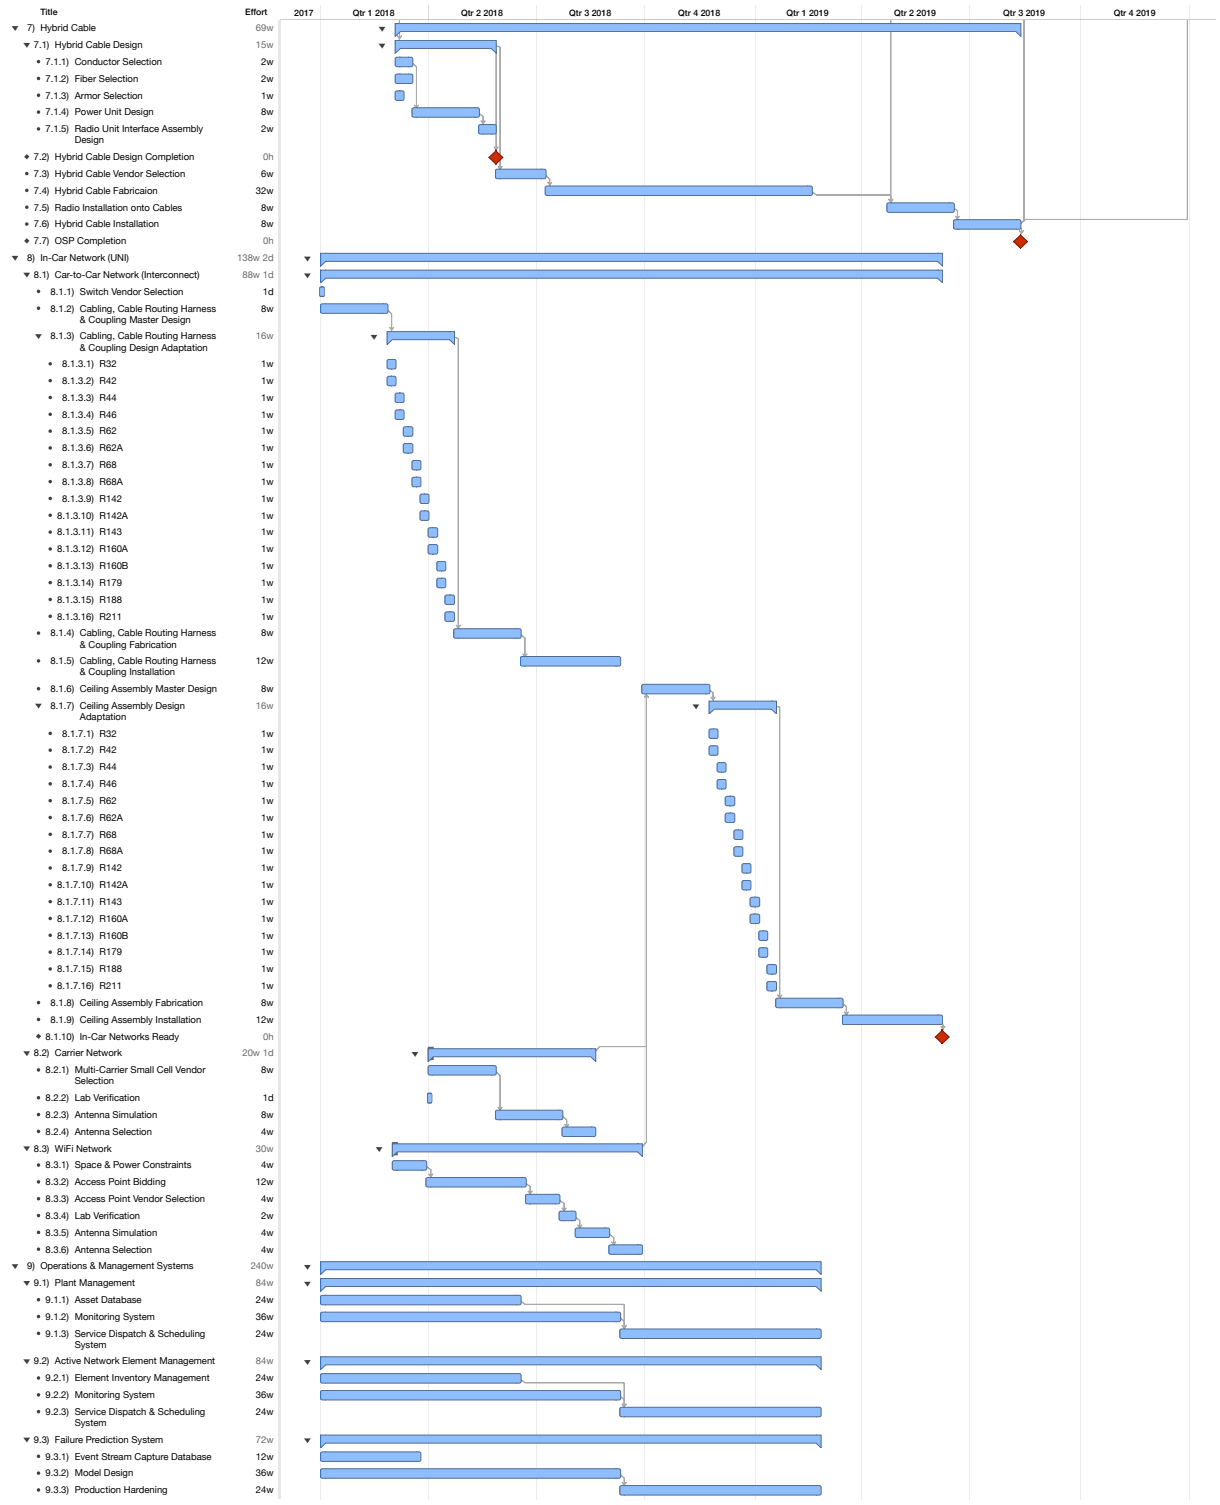


Fig. 10. Project Plan 2/3.

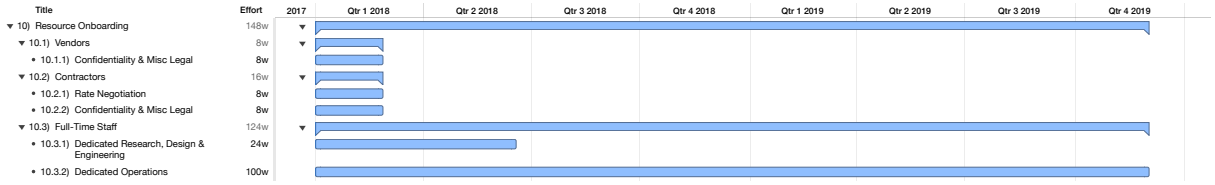


Fig. 11. Project Plan 3/3.