A-TEAM GROUP
Smart Thinking for Financial IT Executives

Financial Markets Tune In To Wireless
Applications, Solutions, Technologies, Challenges

October 2014

Intelligent Trading Technology
Low-Latency. Big Data. Analytics.

In Collaboration with and Sponsored by

CFN Services
Optimizing the power of your network

TABLE OF CONTENTS

Business Perspective: Executive Summary

Wireless communications is not new. What is new is that wireless networks have evolved from supporting mainstream communications to specialty high value applications. One emerging application is its use in financial services, specifically electronic trading. Such communication links are connecting high frequency trading firms to exchanges and liquidity providers.

The latest generation of microwave and millimeter wave technologies are being adopted by leading trading firms because they reduce latency by at least 40% when compared with fiber optic networks between the same points. For latency-sensitive trading strategies, wireless transport of both market data and order entry is seen as essential to remain competitive.

Cost, time to market, control and speed all play into a buy-versus-build decision, as a means to create a differential advantage. However, investment in a wireless network does not guarantee leading performance. Others may deploy faster solutions, and the need for ongoing investment to improve or upgrade technology can impact overall profitability if a market advantage is not achieved.

Given the potential to deliver a significant speed advantage over fiber, wireless networks can be a means for the most latency-sensitive trading firms and market participants to achieve an edge in trading performance. Although financial markets deployment is relatively new, firms have more options to integrate these technologies into their trading platforms than they did just a few years ago – whether that be building and operating a proprietary dedicated network, or leasing bandwidth from a commercial shared services provider, or even leasing from another trading firm that invested in its own proprietary build but has excess bandwidth capacity and is seeking to offset its costs.

Regardless of deployment approach, it is important to have an understanding of what's involved in designing and deploying wireless networks – from network route design and site selection, equipment and processing latencies, bandwidth management and optimization, to overall integration and operation of wireless networks with existing market data and trading systems, networks and infrastructure.

This industry briefing will provide an overview of the adoption of low latency wireless in the capital markets for trading strategies and specific financial applications, dynamics of the wireless marketplace and key considerations for integrating next-generation wireless technologies.

Low Latency Wireless in the Capital Markets

Trading Strategies

Market participants are keen to leverage the speed advantage offered by wireless communications in order to increase the efficiency and potential profitability of a number of key trading approaches and strategies. Today's trading strategies fall into two main categories: structural and economic.

Structural strategies are rooted in technology since speed will allow for the strategy to take advantage of asynchronous market structure. Examples of structural strategies are the following:

- Passive Market Making: A market participant is a liquidity provider by layering the book with bids and asks of different prices and sizes. The participant is attempting to capture the spread between the liquidity rebate and buying and selling the bid and offer. This strategy requires the trading algorithm to quickly respond to shifts in the bid and offer price.

- Arbitrage: This strategy is rooted in capturing market inefficiencies between assets, products and markets. This strategy can be executed within the cash equities market due to RegNMS. Today, it is more common to see this strategy being utilized across asset classes. For example, traders are arbitraging interest rate futures being traded in Chicago with the cash treasuries being traded on electronic platforms in New York. The same can be said for equity index futures and exchange traded funds and underlying equities.

- Spread Strategies: Electronic trading strategies have taken traditional spread strategies to the next frontier. Algorithms are now creating spread strategies across asset classes and products to find new alpha generating opportunities. This requires a distributed infrastructure across regions requiring communication linkage. To improve data coherence, market participants are using wireless communication to transmit essential information at the lowest latency.

Economic strategies use data points that influence price movements and pricing models. Compared to the structural strategies, the trader is assuming an unhedged position in the product and anticipating a price movement. Examples of economic strategies are the following:

- Directional: The movement can be a result of earnings release, news announcement, social media or higher volume. The success of this strategy is to receive the information ahead of the competition in order to act on it first. Market data vendors, traditional and value added, are exploring wireless to deliver services.

- Fundamental: Combining traditional fundamental valuation and technology, this strategy is continuously evaluating the fundamental price of a product by reshaping its assumption in real-time. The model is using both native and inferred data to make trading decisions, such as liquidating or increasing a position, in coordination with the strategies risk threshold.

Applications

In order to achieve success for any one of these trading strategies, firms must think carefully about how they will leverage the use of high performance low latency microwave communications in the right way, specifically when applied to market data and order execution. With both applications, design or selection of a low latency wireless solution begins with understanding the necessary latency, capacity and reliability requirements to meet business objectives.

To put microwave speeds in a practical context, wireless signals through air cover one mile in 5.4 microseconds, compared to 8.1 microseconds via fiber optic cables. In practice that translates to 8.2 – 9.0 milliseconds round trip latency between the CME data center in Aurora, IL to the equities markets in the New York area, compared to 13-16 milliseconds via fiber – a 33% improvement.

While microwave is faster than fiber, trading firms need to intelligently decide how to best utilize the available bandwidth capable with wireless transport. Compared to fiber, which has capacity for terabits of data, the capacity of wireless networks is limited. Today, long haul microwave channels of 100-150 Mbps are common with metro millimeter waves providing bandwidth capacity up to 1-2 Gbps. Increased capacity is in the development roadmap of some vendors (up to 10 Gbps for millimeter wave in particular), however the current capability requires careful consideration of what data should be selected for transmission via wireless, and which protocols to use.

Market Data

When trying to accelerate the delivery of market data over wireless, capacity is usually the greatest challenge. Delivering data in its entirety from major markets is not feasible as the required volume of bandwidth is too large, so typically firms just a send a subset of the data – usually by message type and by symbol – for wireless transmission. While this requires filtering technology to be deployed and configured (which introduces latency), it is often the case that latency-sensitive strategies focus on a small symbol set, and so this data subset approach is entirely workable.

After working through capacity requirements, network reliability needs to be considered, especially for applications that require every price tick to be received. Apart from engineering a network for highest reliability (also introducing latency), firms can decide to deploy architectures that seek to achieve both lowest latency and high reliability by harnessing dual packet transmission over fiber or a secondary wireless network. This approach calls for technology to fill in sequence gaps on the primary network with messages sent via the secondary link and also requires the implementation of a message-sequencing scheme to ensure the correct order and handle duplicity of the data.

The simultaneous transmission approach requires a purpose-built, dedicated, very low-latency wireless network to be augmented with purchased bandwidth on a shared commercial offering that offers higher reliability at the expense of latency. Using a fiber connection to carry a simultaneous transmission may also be employed for even higher reliability. The bottom line, though, is that delivering 100% of market data messages via wireless is challenging, and requires additional technology for filtering and multi-feed arbitration.

Order Execution

Today, there is a variety of trading platform deployment schemes. Trading firms use a bifurcated model that assumes that each server operates independently and assumes that the other server is acting as expected. With an arbitrage strategy, one server initiates a long position in its local market and it is expecting the other server located in another market's data center to execute the short position. This type of model requires that each server communicate trading signals to notify the overall platform of positions and average prices to have an accurate view of the portfolio. While each server is executing in its local market, trade data is being transmitted. The platform will use a binary proprietary protocol to communicate this data and requires guaranteed transmission. Latency is not the main focus but rapid updates will allow for the strategy to respond appropriately in today's fast paced markets.

Centralized trading firms are still commonplace in today's market, especially for economic models where technology is not the competitive advantage. However, latency is still very critical and the trading firm still requires fast order execution. In this case, the strategy is using the wireless network for direct connectivity into the market. Again, this requires guaranteed transmission since the traffic represents actual orders. Relatively compact binary protocols, e.g. OUCH, will be used, instead of standard, but less efficient FIX messaging. The wireless network is protocol agnostic but it needs to ensure reception. This can be accomplished by increasing latency by enabling technologies to reduce packet loss or leveraging synchronized routes to send the packets multiple times.

Buy versus Build Debate

Simply put, trading firms looking to leverage wireless communications today face a *buy versus build* choice. The main issues that typically comprise the decision to build or buy are upfront and operating cost, flexibility and latency, and time to market.

But it's not always been this way. A couple of years ago, such a choice was not available, as shared wireless services did not exist. Early adopter trading firms were thus required to design, build and deploy their own dedicated networks.

The build route – whether by necessity or choice – does have a number of tangible benefits. Key ones are being able to focus on engineering the lowest latency and having complete control over the use of bandwidth. Against those benefits, the drawbacks include the complexity, cost, and time taken to deploy such networks.

Even though a wireless build represents a significant undertaking, microwave networks can typically be deployed much more rapidly and at a fraction of the cost of fiber networks. Installing long haul fiber is costly, but installing dedicated "shortest path" fiber networks over hundreds of miles can be incredibly expensive with installations running $25,000 to $100,000 per mile or more. Comparatively, wireless networks can be installed at less than $20,000 per mile with a much shorter installation time.

However, once a network is deployed, ongoing management and maintenance by a team of experts to drive consistent optimal operation, as well as ongoing technology upgrades to hold latency competitive or ahead of the competition is a major consideration when evaluating the *buy-versus-build* decision.

The challenges of the build route have led a number of trading firms to pursue emerging shared services. In some cases, firms that built their own networks through necessity have since sold them to shared service providers, and become customers in the process.

The benefits of the buy method or shared services model are reduced upfront and ongoing costs, a faster time to market, and overall reduced business risk since the services are underwritten by a third party and supported by a number of customers. Technical upgrades to reduce latency are driven by a common need, and costs are distributed fairly across the customer base.

The top reason not to follow the buy option is that firms lose the latency and control aspects of network ownership that can in some cases offer a differentiated advantage over trading competition. Shared services will sometimes compromise on latency due to network routing or equipment deployed. In addition, bandwidth when shared can be even more limited, and auditing actual availability and service level agreements against measured wireless network performance can be challenging.

Even if a firm chooses the buy route, it is important to maintain some in-house or contracted expertise to conduct due diligence on shared services and to provide independent advice. This expertise should be familiar with various practical and technical issues.

Findings & Industry Next Steps

As more shared services continue to become available in line with current trends, the existing trade-offs between achieving wireless speeds and having to give up available capacity and service reliability will subside. By correctly deploying intelligent latency and reliability schemes across multiple routes, firms will get closer to replicating the consistent performance from fiber at faster and faster wireless speeds.

In order to secure latency improvements across long haul routes, some trading firms and shared service providers will build hybrid wireless/fiber links that utilize pre-existing lowest latency fiber to deliver faster speeds across geographic areas unfit for radio towers. Today, a few providers have already built such networks using wireless to carry data from New York and London financial centers to the Atlantic coastline and fiber to transport the data across the ocean. More hybrid networks will arise as firms continue to seek opportunities connecting other intercontinental markets.

Technology Perspective: Executive Summary

Rolling out a wireless network is a challenging, complex and costly undertaking. Even established equipment providers and service operators continue to face challenges implementing low latency wireless networks that meet the strict capacity, reliability, and bandwidth management.

In addition, market participants must consider whether a hybrid wireless and fiber network approach is suitable to address both latency and capacity needs. The specific network approach will drive contingency planning in the event of signal degradation leading to packet loss, frequency interference, or outages due to various weather extremes. Resiliency methods can vary from auto-failover to a secondary wireless network or a (slower) fiber connection, to an automatic and orderly withdrawal from the market.

Regardless of deployment approach, it is important to have an understanding of the steps involved in designing and deploying wireless networks – from network route design and site selection, equipment specification and processing latencies, frequency licensing and regulatory compliance, to overall integration and operation of wireless networks with existing market data and trading systems, networks and infrastructure.

Additionally, shared service providers need to be cognizant of providing consistent service to customers expecting reliable high performance service. They need to ensure that their clients receive fair and equal access to the network to dispel any gaming of the radio link queue. As reliability can be an issue for firms choosing to trade at wireless speeds, providers must ensure the operation and management of their networks meets availability parameters and service level agreements.

As network reliability is a critical issue for firms choosing to trade at wireless speeds, service providers need to intelligently mange for tight bandwidth constraints inherent in wireless technology today, ensuring adequate bandwidth for data transmission or order signaling. Some providers have chosen to use feed handlers that strip native feeds to popular symbols, forcing adoption of specific APIs and add latency; while others provide more bandwidth-intensive full-depth native market data feeds at the expense of additional bandwidth capacity.

Microwave Network Implementation

Network Route Design and Site Selection

One of the most important aspects of wireless network design is the network route, which needs to be as short as possible to reduce propagation latency. Compared to fiber, microwave can deliver lower end-to-end latency because it can take a more direct line between two ends of the network. Deviating from the most direct route adds latency to transmission.

Underground fiber routes tend to trace winding surface level transportation infrastructure including streets, electrical lines, and railroads that avoid traversing through difficult geographical terrain and private property. Microwave, operating hundreds of feet in the air via radio towers can achieve more direct lines of sight and thus a shorter and faster transmission path.

Even at wireless speeds, every extra mile adds 5.4 microseconds to latency in one direction, or 11 microseconds of additional latency for a round-trip trading transaction, which is significant for low-latency financial applications. In practice, absolute direct routes are hard to achieve, though deviations of just a few percent are often possible. As the mileage between the two ends increases, more radios need to be installed to link the two points together, thereby contributing additional latency due to technology processing.

Wireless networks must be engineered to minimize the overall latency for a particular route by using the shortest path possible and minimizing the number of towers. Although there is a goal to make longer links and have fewer hops, microwave radio power and signals dissipate over longer distances as links are stretched out, which has the potential to make the technology less reliable for trading applications.

While it is possible for microwave signals to be broadcast across dozens of miles, it is not always possible to obtain a clear line of sight between endpoints, due to terrain, buildings, other immovable obstacles. Transmission across water can also be an issue due to potential air disturbance from water evaporation or reflection from the water surface. While minimizing distance across water is desirable to boost reliability, it may require a non-optimal route to achieve it.

Another consideration when constructing a wireless network is whether to leverage existing wireless telecommunications towers for reduced cost and time to market. This depends mostly on whether the towers are already optimally located for shortest distance. Leveraging existing tower infrastructure can be desirable if space on them is available. However, without ownership it is not easy to ensure acceptable tower condition and location. Clearance height and ruggedness of the tower foundation to accept additional antennas without the need for costly structural reinforcements must be verified in advance as microwave antennas are large and heavy.

Alternatively, tower construction is complex, costly and time consuming. Apart from actual construction costs, it is often necessary to obtain planning permission from authorities and building rights from landowners as well as accept towers as subject to regulatory oversight. Curvature of earth issues can also require tall towers to be built high enough to ‘see’ over the horizon, which create additional exposure to wind and lead to reliability issues.

In metro areas, where existing building walls or roofs are used for mounting, it will be necessary to secure appropriate rights from building owners (which can be costly if said owners determine the business value of the network), including access rights to perform on-going maintenance. For connectivity over longer distances, such as Chicago to New York City, or across the English Channel, route optimization can be costly but also result in significantly lower latency.

Equipment and Processing Latency

Wireless transmission not only achieves a latency advantage due to route efficiency; it is also faster than fiber because wavelengths move faster through air than through glass. Glass adds additional resistance that slows down light by 30-40%. However, since radio signals attenuate and degrade as they travel across distance, and are impacted by phenomenon such as weather and other radio interference, reliable transmission requires the deployment of latency adding repeater devices along a route, in order to “clean up” and regenerate, or simply boost those radio signals. The number of repeaters required is a complex determination based on many factors, including the wireless waveband used, the power of transmission and the gain/amplification characteristics of the radio/antenna equipment.

Wireless networks may incorporate both analog and digital repeater technologies. Analog repeaters can typically only amplify and re-broadcast the signal while digital repeaters, or regenerators, can theoretically reconstruct a signal to near its original quality. While analog repeaters can eliminate digital processing latency, they do not reduce signal noise and can result in an increased error rate.

The more repeaters/less distance between repeaters in a network increases reliability (against factors such as weather or the signal-to-noise ratio which can impact the accuracy of received data), but each repeater typically adds hundreds of nanoseconds of latency.

Depending on their deployment and configuration along the network path, regenerators can add significant end-to-end latency. One vendor estimate puts added latency at 4% to 14% for each additional hop, as they convert analog signals to digital, apply Forward Error Correction (FEC) processing, and then convert digital back to analog for onward transmission. Digital signal regeneration has the advantage that it will clean up any noise (errors) in a signal along the route, but at the cost of increased latency.

Often, a mix of analog and digital repeaters are deployed, or hybrid adaptive repeaters are used. Configuring a network with both regenerators and repeaters is one method to reduce end-to-end latency; and creating a smart network using hybrid adaptive repeaters, where regenerators and repeaters are alternately activated based on data errors or environmental conditions, is beneficial but more costly.

Some networks are specifically over-engineered in terms of number of repeaters to build in reliability beyond what might reasonably be required. Such an approach obviously wins at times when extreme weather might cause other networks to go down, but at the expense of increased latency all of the time.

Today’s Wireless Network Landscape

There has been significant investment in wireless networks. Figure 1 illustrates a consolidated view of wireless networks in operation today including a fiber optic route as well. While the fiber route looks to be traveling a similar route, the divergence from the direct path adds distance translating into significant latency.

Looking more closely at the wireless networks, the market is crowded with many networks. Subsequently, bandwidth capacity is not the limiting factor but all networks are not equal. As a result, the fastest or shortest path route demands a cost premium is protected by the monopolistic nature of regulation. Other networks can be more cost effective with minimal impact to latency and overall performance, opening up the excess supply to the market.

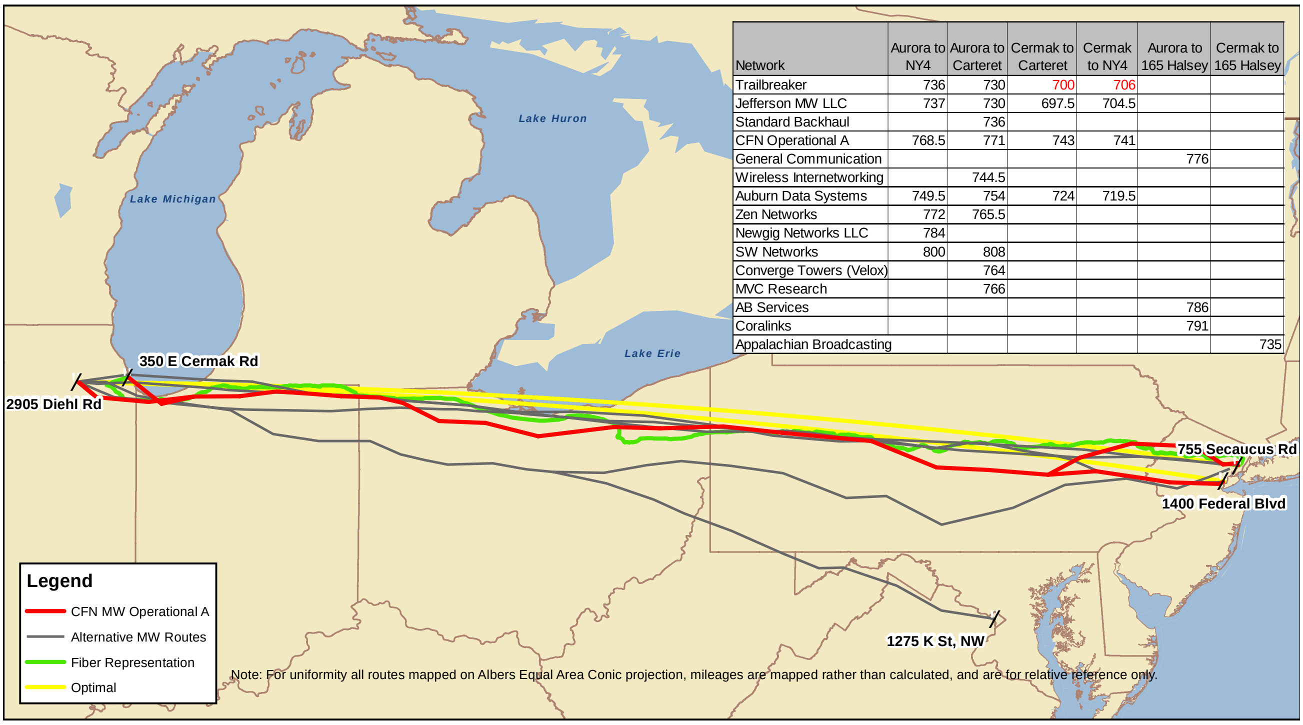

Network	Aurora to NY4	Aurora to Carteret	Cermak to Carteret	Cermak to NY4	Aurora to 165 Halsey	Cermak to 165 Halsey
Trailbreaker	736	730	700	706		
Jefferson MW LLC	737	730	697.5	704.5		
Standard Backhaul		736				
CFN Operational A	768.5	771	743	741		
General Communication					776	
Wireless Internetworking		744.5				
Auburn Data Systems	749.5	754	724	719.5		
Zen Networks	772	765.5				
Newgig Networks LLC	784					
SW Networks	800	808				
Converge Towers (Velox)		764				
MVC Research		766				
AB Services					786	
Coralinks					791	
Appalachian Broadcasting						735

Figure 1 - Chicago to NJ Wireless Routes

Bandwidth Management

A microwave network has limited bandwidth compared to a fiber circuit. The operator of the microwave network needs to allocate bandwidth amongst the users and ensure that all users have fair and equal access. Electronic trading platforms are real-time applications requiring fair and equal access to the network. Fair and equal access is a set of stringent requirements, guaranteed throughput based on bandwidth allocation and identical priority for all users.

Generally, all users demand the same resource at the same time, such as market open and close and events. It is important that access to the network is deterministic to handle traffic bursts but prevent network saturation from bursts. Otherwise the network will not be available to the majority of users at critical times during trading hours.

Bandwidth management is designed to provide a deterministic service to all users while policing bandwidth allocation. It is important to deploy the correct scheduling algorithm that will meet those demands of the technology. Round robin scheduler is a circular approach for assigning user data packets to data frames on the network. It ensures that each user has equal priority and prevents any one user from starving other users.

To manage fair access, each user has a specific allocation of bandwidth. The challenge is the mismatch between frame size and IP packet size. APIs primarily drive IP packet size, hence it is unreasonable to assume that each IP packet equals the frame size. This results in orphaned bandwidth when the IP packet is less than the frame size, or the user will experience jitter when an IP packet is fragmented across multiple frames due to two or more scheduling cycles.

There are several established algorithms used to overcome this challenge. Token bucket is just one example that maximizes bandwidth utilization while accommodating IP packet sizing differences and handling bursts. To improve the quality of service, the user’s packet might be split into smaller sizes to avoid orphaning capacity.

In simple terms, the bandwidth is divided into a fixed number of tokens and each user is given a token on a set rate based on its bandwidth allocation. When it is the user’s turn, the scheduler validates 1) the user has enough tokens to send the data, and 2) the size of data divided by the token size is greater than zero. If yes, the tokens are “cashed in” and the data passes. Otherwise, the user does not conform because of insufficient tokens and must wait until sufficient tokens have been accumulated.

Another common method to improve capacity and maximize data is to remove Ethernet and IP headers at transmission and reconstruct the header data with static data on the receiver.

Line Rate Arbitration

The industry relates data arbitration with market data feeds. The feeds use multicast to transmit order book updates to the market. Multicast doesn't guarantee the delivery of a packet so the exchanges transmit two identical feeds, A and B feeds. Data arbitration technology creates a single stream for a trading application by receiving and arbitrating both feeds.

Using a similar approach to market data systems, network equipment replicates the traffic across two routes, microwave and fiber, at the source. At the destination, the network arbitrates both routes into a single stream. This combats the higher likelihood of packet loss over the microwave link without having to build retransmission capabilities. Possibly, the traffic may be replicated over two microwave networks to improve the overall reliability of transmission.

Line rate arbitration enables the user to optimize system performance based on his or her operational risk associated with packet loss. The following table outlines the data reliability threshold settings and use cases.

Threshold	Criteria	Performance
High	Data reliability > Latency	For a packet loss, the output stream will be delayed for a set period of time while waiting for the missing packet on the secondary route.
Medium	Data reliability = Latency	The system will have a waiting period based on the propagation delay delta between both routes.
Low	Data reliability < Latency	The system will have no or very small waiting period before skipping the missing packet.

Network arbitration reduces the complexity of the overall network. Specifically, it reduces the burden on data retransmissions. Limited bandwidth capacity precludes the service provider the ability to offer market data retransmission capabilities on the microwave network, reverting all retransmissions on the higher latency route. More importantly, line rate arbitration is application agnostic. For raw bandwidth usage, this enables the trading firm to improve its reliability without having to implement retransmission capabilities directly into its application. It is handled at the network level and appears to be a single stream of data. As a result, the user needs to set threshold tolerance based on packet delivery reliability.

Improving Reliability

Reliability is often cited as a major obstacle to a trading firm's adoption of wireless, specifically the weather conditions such as rain, fog and snow. The financial markets require "5 nines" reliability for its mission critical trading operations. Not being able to trade for any reason means missed opportunities. While wireless networks can be engineered to this standard, it is at the expense of increased latency. This includes more radio towers, signal regeneration or enabling forward error correction (FEC) features.

For this reason, trading firms adopt contingency plans to handle with wireless disruptions and outages. Using two wireless networks and a fiber route improve your reliability without sacrificing latency. The primary route is the fastest wireless network configured to maximize latency reduction between the nodes. A secondary wireless route with higher latency is more reliable because all sites are regenerating the signal and using FEC. The secondary route is still 33% faster than the fastest fiber route. This enables the firm to remain active in the market instead of making defensive trades, such as closing out positions. At last resort, the trading firm will revert to a fiber route to maintain connectivity between both nodes.

This contingency model requires route management. The simplest approach is to replicate the network traffic across all routes and use smart device or software to arbitrate the traffic, passing the first received packet to the application. However, this introduces non-deterministic behavior in the system because the packet can either arrive on the fastest route or the slowest route, a potential 40% latency difference.

Simple is good, but it comes at a cost and impact to the overall system. Another approach is to enable the trading firm to actively manage its own network traffic across the routes. This will allow the firm to maximize bandwidth utilization by not duplicating data. The trading firm will monitor the health of the wireless linkage via a telemetry feed. This real-time feed will provide weather data at each tower enabling the user to predict pending network degradation. As a result, the system will begin to route over the secondary wireless or fiber route before the wireless network packet loss increases beyond the trading firms own thresholds.

To improve reliability, trading firms and network operators can enable Forward Error Correction (FEC). This feature adds data to each message that allows errors to first be detected, and then located and corrected. Logic to implement FEC is complex and adds to latency but can be minimized by using Field Programmable Gate Array (FPGA) technology. However, an issue with FEC is that the data added to each message can be sizeable, around 25% of the original payload for some FEC algorithms. Thus, improved reliability comes at the expense of increased processing latency and longer messages, and thus reduces capacity for application data.

Trading firms building their own networks might opt to turn off FEC to reduce latency, understanding the risk of lower reliability as well, and perhaps opting to employ contingency approaches when glitches do occur. A shared service network provider is able to run the fastest route with little overhead and a second route with FEC enabled in order to meet service levels for reliability without impacting latency and appealing to a broader set of customers.

Industry Next Steps and Looking Ahead

Since 2012, the financial markets have been using wireless networks to reduce latency between strategic data centers and markets. Similar to other financial technology, the market continues to learn from experience to reduce latency and improve reliability. Being that the fastest routes, analogous to the shortest path, are deployed, latency improvements will come from optimizing the technology. These improvements include faster and better radios and new error correction processing to eliminate signal regeneration.

More importantly, providers are best positioned to improve wireless networks because they are able to couple together multiple routes to improve reliability and continually invest in route improvements. The provider is able to pass on aggressive economics to all of the users while delivering a better service due to its route diversity and deep technical experience.

The industry will most likely see a consolidation of networks and a transition of the networks from trading firms to technology providers. The most latency-sensitive high frequency trading firms were the first movers but many have realized their capacity needs are not as great as originally forecasted. Instead of carrying heavy operating expenses, some trading firms are looking to divest ownership but still use the service. Technology providers can step in to run and manage the networks while opening up the market to the unused bandwidth at competitive prices. The providers will also leverage the different networks to create product offerings that will meet the different requirements of market participants.

Trading firms must weigh best-fit approaches for integrating next-generation wireless technologies into their trading infrastructure. Each firm needs to consider the implications of the network provider's deployment model against its specific business and technology requirements. Specific infrastructure providers are now focusing on centralizing a portfolio of wireless networks into a single platform. This enables trading firms 1) to maximize bandwidth utilization without sacrificing on data such as full depth, native feeds; 2) to prioritize data into latency profiles to optimize bandwidth utilization for specific applications; and 3) to improve reliability across the entire network. The provider with a diverse network portfolio uses the scales of economies to deliver these capabilities at very competitive economics.

CFN Services
Optimizing the power of your network

About CFN Services

CFN Services is a leading provider of managed services that improve the trading performance of many of the world's most sophisticated financial markets participants.

CFN developed and operates the Alpha Platform™, a high-performance global platform for ultra-low latency market data delivery and trade execution across key liquidity venues, major asset classes and emerging geographies worldwide. Alpha Platform integrates best-of-breed trading technologies and services into a single platform, connecting more than 70 leading financial data centers, trading venues and proximity locations worldwide. For more information, please visit www.cfnservices.com.

A-TEAMGROUP

www.a-teamgroup.com

A-Team Group 'gets' the business of financial information technology. It's what we do and what we are passionate about.

Since 2001 we have delivered news and analysis, in-depth research, and events for financial institutions, technology and data suppliers, regulatory bodies and more. So whether you are into reference data, looking for low latency, want to manage your risk technology, or thinking big data, we've got the information you need!

Find out more about what we do by visiting our website: www.a-teamgroup.com

And if you are a technology or data vendor on a quest to demonstrate your thought leadership, generate quality sales leads and measure the success of your marketing campaign, we have creative ways to help you achieve your goals. Find out about Smart Marketing, The A-Team Way by visiting www.a-teamgroup.com/smart-marketing.

Or contact us:
Call **+44 (0)20 8090 2055** Email **info@a-teamgroup.com**

For sales, contact Caroline Statman at **caroline@a-teamgroup.com**

For editorial, contact Andrew Delaney **andrew@a-teamgroup.com**