

Research Note

March 2025

# Cohere Technologies

## AI, Networks and The Intelligence Plane

Author: Robert Curran



[www.appledoreresearch.com](http://www.appledoreresearch.com)



**Appledore**  
RESEARCH

Published by Appledore Research LLC • 44 Summer Street Dover, NH. 03820

Tel: +1 603 969 2125 • Email: [info@appledorerg.com](mailto:info@appledorerg.com) • [www.appledorerresearch.com](http://www.appledorerresearch.com)

© Appledore Research LLC 2025. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means – electronic, mechanical, photocopying, recording, or otherwise – without the publisher's prior written permission.

Figures and projections contained in this report are based on publicly available information only and are produced by the Research Division of Appledore Research LLC independently of any client-specific work within Appledore Research LLC. The opinions expressed are those of the stated authors only.

Appledore Research LLC recognizes that many terms appearing in this report are proprietary; all such trademarks are acknowledged, and every effort has been made to indicate them by the normal USA publishing standards. However, the presence of a term, in whatever form, does not affect its legal status as a trademark.

Appledore Research LLC maintains that all reasonable care and skill have been used in compiling this publication. However, Appledore Research LLC shall not be under any liability for loss or damage (including consequential loss) whatsoever or howsoever arising because of the use of this publication by the customer, his servants, agents, or any third party.

Publish date: 25 February 2025

Cover Photo by Bing/Robert Curran

## Introduction

In March 2024, Appledore profiled **Cohere Technologies**, a pioneer in the field of mobile network technology, and their Universal Spectrum Multiplier (USM) software. USM addresses the problem of congestion in mobile networks by effectively creating additional available spectrum (as much as double) when cells are at capacity (i.e., all available radio channels are consumed).

USM does this by through an advanced new technique for allocating radio spectrum in real time to the many devices that need service, based not on RF modelling, but considering device location and speed, and other information about the channel between transmitter and receiving device.

The initial application for USM has been as a “supercharged” scheduler within the base station, kicking in when cells are congested, reverting to a conventional scheduler under quieter loads. The business value of that application lies not only in ensuring customers stay connected (especially in busy areas/periods), but also in the avoidance of more costly or intrusive alternatives: purchasing and installing additional network equipment or acquiring more spectrum.

Through 2024, Cohere proved its technology with successful field trials, most notably at Vodafone in a dense urban environment (boosting 5G capacity by 50%, achieved entirely through software), and at Bell Canada. The first operational deployments are expected by late 2025.

But capacity multiplication is just an initial application of Cohere’s technology. Combining the data generated by USM with the latest AI breakthroughs – LLMs in particular – has significant new implications and opportunities for mobile providers.

In this paper we take an analyst’s view of the new value potential for telcos that arises from a combination of a completely new and untapped data source, off-the-shelf AI, and programmable networks.

## Mapping the Channel

The job of estimating how much spectrum to allocate to every active call or session is the beating heart of a mobile network. It is what enables mobile networks to support many concurrent users or devices per cell. Getting it wrong means dropped calls, failed calls, frustration for users. Getting it right means more users (devices) successfully connected for as long as they need, which translates into squeezing more revenue from each and every cell – and an increased return from investment in expensive spectrum.

Finding a way to improve it means that operators can generate more revenue for years longer from their existing networks, delaying the need for expensive and disruptive network upgrades (and potentially spectrum). And since newer spectrum is generally higher spectrum with worse propagation characteristics, squeezing more out of *existing* low- and mid-band spectrum (and equipment) is better for customers, too.

The current state of the art in channel estimation is statistical. Cohere’s innovative modelling of channels, rather than of frequencies, incorporates temporal and spatial dependencies, environmental factors and current network conditions. With better-informed predictions, RAN

parameters can be more finely tuned and optimized – modulation schemes, coding rates, and power allocation.

Cohere’s technique for calculating the requirements of a radio channel is based on a combination of the range and velocity of a user device and how the radio signal propagates from the cell site to that device. Instead of time and frequency, Cohere’s model uses distance (as measured in signal delay) and speed (as measured in Doppler shift). The resulting map is good for up to 50 milliseconds, thereby reducing the processing load within base stations, which otherwise may have to re-estimate channel requirements much more frequently. This allows allocated channels to be used for longer – effectively slowing down the channel aging process.

With the delay-Doppler channel model, Cohere has a real time view of everything in the wireless channel, which can be used to better optimize performance and the experience for mobile users. This means Cohere maps all energy, interference and reflectors, thereby creating a combined mapping of the physical world with the wireless world.

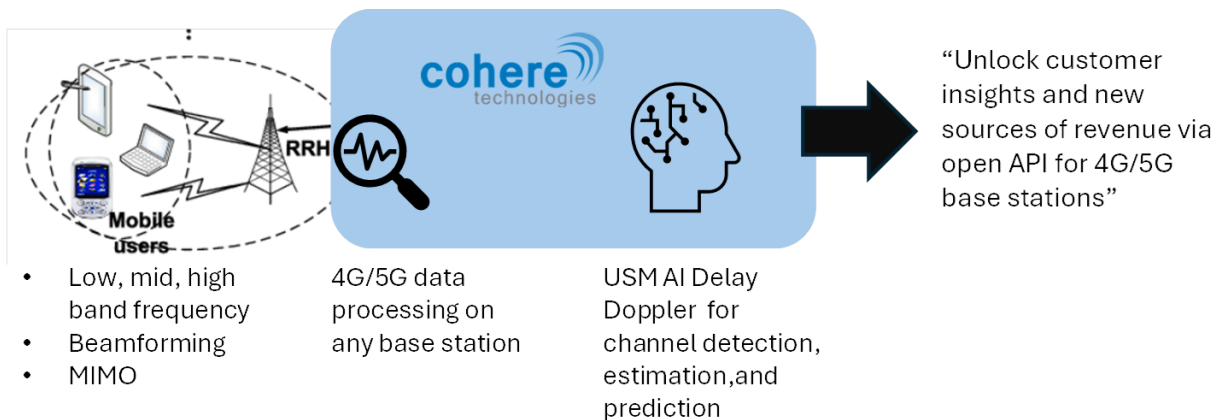
With a more realistic view of how specific signals propagate within a specific environment, beams can be formed more precisely to UEs and spectrum can be used more efficiently. Users can be scheduled in the *same* time and frequency slots, compared to conventional methods that require *separate* time or frequency slots for each user.

## Into the Intelligence Plane

The next step takes advantage of AI and cloud to automate channel map information, enabling “digital twin” style insights of the physical and wireless channel data which will improve user experience and real-time location data of all energy in the wireless domain.

Alongside conventional network control and data planes, this could be considered as a separate “intelligence plane”, providing a distinct resource about current and predicted network status and potential that can be used for a range of applications.

Figure 1: New uses for low-level network data

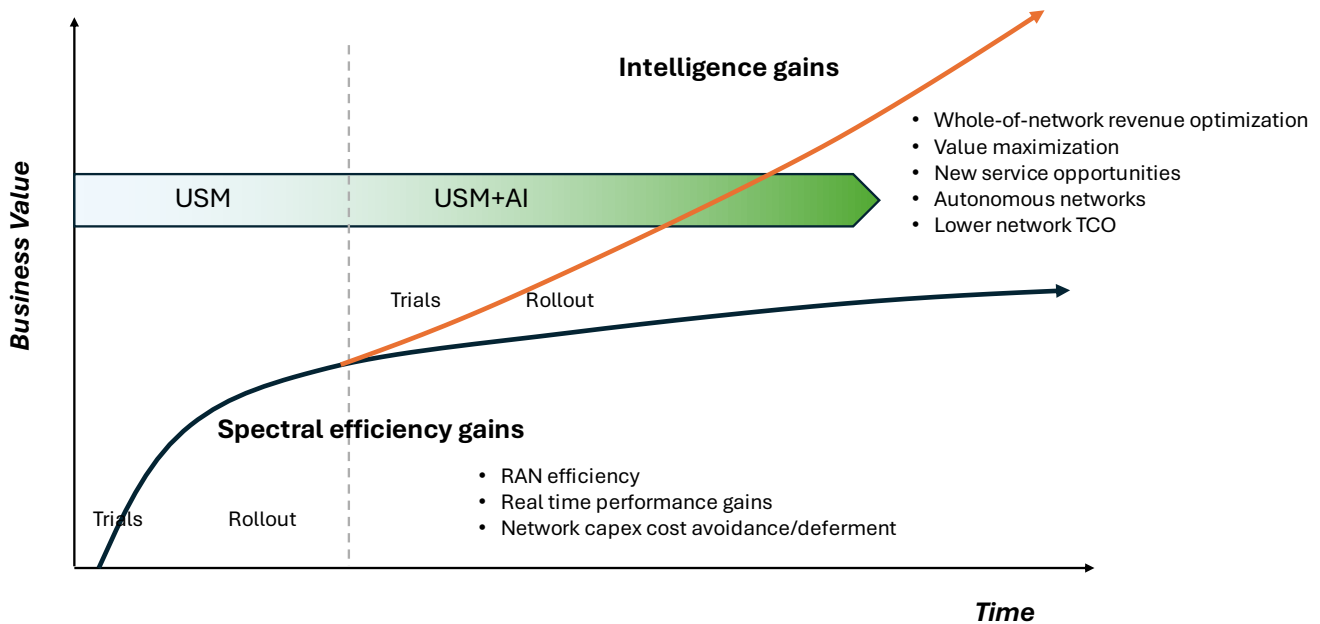


Source: Appledore

## Next generation architecture

Pairing USM with AI is the next step. With a pre-existing basic channel model or one developed from an IQ sample drive test of an area, a Retrieval Augmented Generation LLM architecture pattern is envisioned to develop learning models that explicitly reflect the characteristics of each cell. These future models will enable a new generation of digital twin that fuses the wireless and physical worlds. This twin can be used to model and recommend changes, along with augmenting capacity and financial planning for network improvements leading to more efficient spectrum utilization and improved network performance.

**Figure 2: Benefits of intelligence beyond spectrum multiplication**



Source: Appledore

USM's Machine Learning and sector fingerprinting capabilities today are extremely capable but are real time systems. Mobile networks can be considered as a set of control loops, from fast loops (such as the short cycle time required for channel estimation and scheduling) to slow loops (in the extreme, a network planning program). Slow loops afford an opportunity to learn from data accumulated over time, to spot patterns and anomalies. Cohere's technology uses a vast amount of real-time network data to construct its channel estimations and predictions (for "fast loop" use in scheduling). But the data generated can provide value across "slow loop" processes, not only for network planning but also sales and marketing, and product development.

The continuous learning of the system will include traffic patterns, effects of planned and unplanned events, environmental factors, network load, and geographic areas of challenging network conditions.

Pattern recognition capabilities in such a system allow it to detect subtle correlations in conventional USM data that might be missed by conventional modeling approaches, leading to more

efficient spectrum utilization and improved network performance. This synergy could revolutionize how operators approach wireless network planning, optimization, and management.

Operators could use this to proactively identify revenue opportunities – which cells are capable of supporting additional traffic, and how much. The closed loop that automates the process of changing transmission parameters means this can be a fast, business-led decision, rather than a long-cycle engineering project.

In a FWA context, this could allow operators to generate more revenue, through a better-informed analysis of cells available capacity. Since device mobility is low, the business impact could be significant.

## From Predictive to Adaptive Channels

The integration of AI with USM represents a significant advancement in wireless channel modeling. By leveraging the massive amounts of data generated by USM (including uplink and downlink channel measurements, multipath components, delay spreads, and interference patterns) the combination of USM and AI can create highly sophisticated and adaptive channel models that more accurately represent real-world wireless environments. These models go beyond traditional statistical approaches by incorporating temporal and spatial dependencies, environmental factors, and network conditions in real-time. The USM AI-based system can continuously learn from actual network performance data, adapting its predictions to account for changes in user density, weather conditions, and network load.

This dynamic approach enables more precise optimization of transmission parameters, including modulation schemes, coding rates, and power allocation. The model's ability to process multiple data types simultaneously and identify complex patterns makes it particularly valuable for scenarios involving dense urban environments, high-mobility users, or challenging propagation conditions.

## Beyond Capacity

Being able to accurately predict the channel characteristics experienced by devices moving at different speeds, there is an opportunity to create differentiated, premium-priced mobile service with higher QoS commitments.

Knowing how channels behave reveals information that could be of value in other contexts: for example, the behavior (over time and space) of users/devices. Real-world factors that can affect RF propagation include:

- Urban “canyons” caused by tall buildings
- Indoor environments, where cellular signals are attenuated by building materials
- Old buildings with thick walls
- Rugged terrain with steep valleys or dense vegetation
- Areas subject to extremes of climate – rain, fog
- Highly mobile users, frequently moving between serving cells

- Industrial plant, with many reflective surfaces
- Higher frequency (5G) signals are especially susceptible to propagation challenges.

Some of these factors are temporary (transient) or seasonal (repeated and predictable). And since devices typically move, it is not possible to create a permanent map.

The combination of complex environments and a wide range of user behaviors and network conditions makes for an RF challenge – but also illustrates the opportunity in using AI and geo/temporal/channel data.

For example, **airports**, with their mix of dense and sparse areas, complicated reflection and interference, and wide range of moving UEs – from people with cellphones to baggage and catering trucks, and moving aircraft, plus security and safety monitoring. The sort of analysis available from AI-processed network-sourced data could reveal (and be used to resolve):

- Safety “blind spots” caused by loss of network coverage due to weather conditions or vehicle positions – these could compromise staff or passenger safety.
- People congestion – especially where airports connect with metro transportation systems which may suffer delays.

The same benefits could apply just as well to any complex environment which mixes people, vehicles – such as ports, transportation hubs. Telcos providing wireless coverage could also monetize this extra analysis that is separately available.

This data could be useful in an **environmental** context: using 5G-enabled devices to monitor the condition of remote areas such as forests. Processed data could reveal anomalous changes resulting from poor weather conditions, fire or flood. This could be used to provide an early warning to the appropriate authorities – or just to hikers and tourists.

There is value too for **architects and builders**. Ensuring mobile connectivity within buildings is essential – and telcos could be in a great position to share aggregated, anonymized real-world experience to building owners or neutral host providers to identify problem areas (higher floors?) that need attention.

We should not forget the prosaic: the **virtual drive-test**. Once built, a network must be continually tuned and optimized: power levels, beamforming and handovers. Load balancing and traffic engineering. Constant, routine drive testing is required to assess network quality or investigate “not spots” and “grey spots”.

A process that is necessarily limited in its scope, the aggregated information created by this approach would provide a much more efficient and faster way for telcos to assess coverage problems – before customers simply try an alternative provider.

**Venue operators** are also increasingly aware of the desire of stadium-goers to create and enjoy digital experiences while they attend events. Yet in-stadium coverage is generally problematic, and the environment is even more complex today with mobile retail payment terminals. Armed with

better information on the user experience, venue operators could direct customers in real time to alternative parts of the stadium. The same challenges apply with large scale outdoor events such as music festivals.

There are also applications in a **military context**, where 5G networks are gaining traction for battlefield communications. These environments are likely to involve a mix of vehicles and on-foot personnel, as well as autonomous devices and airborne surveillance/relay.

Spectral *capacity* is a static resource that is only consumed and released on demand. But its *efficacy* depends on a range of other factors which can dramatically improve the utility of mobile networks for specialized applications and environments. And knowing the efficacy allows other decisions to be made more effectively: how to keep people safe, where and when to promote products.

### Channel Estimation

In mobile networks, uplink channels are typically estimated more frequently than downlink, often on a per-transmission basis using reference signals (in LTE and 5G these are referred to as Sounding Reference Signals, SRS). Downlink channels are estimated less frequently. The frequency with which channels are estimated can range from a few milliseconds to several hundred milliseconds.

For users moving quickly, channels must be estimated more frequently. Different mobile generations have different requirements and ways of estimating channels. This creates the need for considerable processing power at each and every cell site.

Technology	Channel Estimation Characteristics
<b>TDD</b>	<ul style="list-style-type: none"><li>• Can leverage channel reciprocity to infer downlink channel from uplink estimation</li><li>• Less frequent explicit downlink channel estimation</li></ul>
<b>FDD</b>	<ul style="list-style-type: none"><li>• Require separate estimation for uplink and downlink due to different frequency bands</li><li>• May need more frequent downlink estimation compared to TDD systems.</li></ul>

The exact frequency of channel estimation can vary significantly based on the specific implementation, network conditions, and user requirements. Mobile network operators often optimize these parameters to balance performance, overhead, and system efficiency.



## Impact of AI on Network Architecture

There is an active debate in telecom about the pros and cons of centralized vs decentralized AI compute. The extra longevity of USM's channel estimates means that user/data and control planes can be separated, with the bulk of control plane operations located in a cloud data center with clusters of USMs with AI on-board. By using a centralized (i.e. within a data center) USM and AI, it is possible then to have an overall view of traffic and to exercise network control across multiple sites or regions – as a single pool of available capacity. This sort of vision was one of the original ideas for self-optimizing networks (SON), but until now that has been an impracticable option, with SON limited operating on a much more localized set of cells.

In addition, making the USM/AI control plane a sharable, scalable, modular resource in the data center addresses both economic and physical concerns associated with trying to distribute AI capacity to every base station location: capex and opex (for servers in every site, leading to stranded capacity, and repeat engineering) and the additional demands on power, space and cooling.

## If data is the new oil, we just found a whole new well...

Data has always been valuable. What has changed in the last few years is our ability to process very, very large amounts of data. The grocery transactions of millions of shoppers, nationwide – every month. The movement of urban populations through a metropolitan transit system – across a year. The driving habits of drivers – across thousands of miles. By processing data in the right way, new value is unlocked: not only trends and predictions, but affinities, cause-and-effect relationships, likelihoods, risks. The response of companies to that intelligence has shaped the design of new products, the timing of their launch, price points, distribution channels, marketing targeting, things that a generation ago would have been – in effect – something not far from guesswork. So telcos already know that the right data, used in the right way, can be immensely valuable – when it is deployed appropriately, at the right time, in a context where the recipient can take a high-value action based on it.

Tapping transaction and consumer data is only one part – and by far the much smaller part – of the data that a telecom business has access to. What a telco's *network* is capable of revealing can not only enable huge efficiency gains and customer intelligence for a telco, but also as a source of untapped value for a range of third-party users of that data.

With Cohere's approach, operators will have:

1. A technology that allows more mobile users (devices) to be supported within the existing available spectrum.
2. A map of how RF signals behave within a specific spatial environment i.e. considering the specific topography around a cell, and the experience of users/UEs as they move through it.
3. A history of how that model changes over time.

Mobile operators can break free from spectrum, capacity and large amounts of capital being dominant constraints on revenue growth. The ability to extract or derive new data from existing networks provides new commercial opportunities.



**Insight and analysis for telecom transformation.**

 @AppledoreVision

 Appledore Research

[www.appledorerresearch.com](http://www.appledorerresearch.com)

[info@appledorerg.com](mailto:info@appledorerg.com)

+1 603 969 2125

44 Summer Street Dover, NH. 03820, USA

© Appledore Research LLC 2025.