

Coherent Multiplex: Real-Time Signal Coherence Analysis and Feature Extraction Platform

Noah Shore

August 2025

Abstract

The Coherent Multiplex is a real-time signal analysis and visualization platform designed to identify and monitor coherent relationships among multiple time series signals. By combining fast Fourier analysis, wavelet coherence, and interactive graph representations, the Multiplex enables efficient exploration of dynamic interactions in live signals. This whitepaper introduces the system architecture, core methodology, current prototyping efforts, and potential applications in research, trading, and live monitoring. Possibilities and goals for future directions with the Multiplex are also discussed.

1. Introduction

Understanding and employing wavelet coherence between signals is becoming increasingly popular in fields like neuroscience, finance, geophysics, and communications, where uncovering relationships between signals can reveal underlying mechanisms and patterns. Conventional approaches to coherence analysis are limited to high-level programming languages, run outdated transforms, and are incapable of handling high-volumes or real-time data streams efficiently. The Coherent Multiplex overcomes these challenges by offering a rapid and scalable solution for live coherence tracking, allowing users to explore frequency- and time-dependent interactions with greater speed and precision.

By leveraging advanced algorithms and optimized computing frameworks in the *Fastest Fourier Transform in the West* (FFTW) [1] and the *Fast Continuous Wavelet Transform* (FCWT) [4], the Multiplex significantly reduces the processing time required for coherence analysis, making it suitable for real-time and large-scale applications. This innovation, combined with the living multilayered network architecture, will enable researchers and industry professionals to efficiently monitor and investigate dynamic signal relationships as they evolve, supporting new opportunities in data-driven decision making and responsive system design across diverse scientific and industrial domains.

2. System Overview

The Coherent Multiplex absorbs, analyzes, and displays relationships among signals through several key components described in this section and detailed in the next.

Signal Ingestion

The platform supports flexible signal ingestion through APIs or direct connections to monitoring devices. This allows users to seamlessly integrate diverse data sources, whether signals are generated internally or are continuous real-world data streams, enabling comprehensive and timely analysis.

Real-Time Spectral Transform Compute

To efficiently analyze signals in the frequency domain, Coherent Multiplex employs fast Fourier transforms computed in real-time using optimized libraries. This approach ensures rapid processing of high-throughput data streams, facilitating immediate insight into the spectral characteristics of each signal.

Multilayered Graph Network

A multilayered graph network is constructed to represent the complex relationships among signals. Each node corresponds to an individual signal, and edges denote the strength of coherence or similarity between them. This layered structure supports exploration at multiple levels of granularity and aids in visualizing interdependencies.

Pairwise Cosine Similarity Analysis

Cosine similarity metrics are calculated pairwise between FFT magnitude vectors in the frequency domain. This measurement quantifies the alignment of spectral content between signals, serving as a computationally efficient heuristic to identify pairs exhibiting potential coherence, which can then be further analyzed.

Selective Wavelet Coherence

Additional edges between nodes will grow and decay based on cosine similarity. These edges contain wavelet coherence arrays that are updated at a sub-sampled interval because of prohibitive cost. Utilizing continuous wavelet transforms, this method provides detailed time-frequency resolution, revealing dynamic and transient coherent interactions between signals that evolve over time.

Interactive Dashboard

An interactive dashboard presents a comprehensive visualization suite including time-domain signal plots, FFT spectra, network graphs illustrating signal relationships, and wavelet coherence heatmaps. These elements can be explored dynamically, supported by live updates delivered via server-sent events, to gain actionable insights in real time.

3. Technical Outline

3.1. Incoming Data

A live signal can be represented as a real-valued function of time:

$$x(t) : \mathbb{R} \rightarrow \mathbb{R}$$

where $x(t)$ is defined only up to the present moment t_0 :

$$x(t), \quad \text{for } t \leq t_0$$

A continuous-time signal $x(t)$ sampled at a frequency f_s (in Hz) becomes:

$$x[n] = x\left(\frac{n}{f_s}\right), \quad n \in \mathbb{Z}$$

where f_s is the sampling frequency (samples per second), and n is the discrete time index. For the real-time processing employed by the Multiplex, a finite-length buffer is defined:

$$X[n] = \{x[n - N + 1], x[n - N + 2], \dots, x[n]\}$$

where N is the window size. Multiple signals are indexed by $i \in \{1, \dots, M\}$, each with its own discrete-time sequence and buffer $X_i[n]$. By aligning our signals in an M -dimensional vector, the full datastream is represented in the $N \times M$ matrix D :

$$D = \begin{bmatrix} x_1[n - N + 1] & x_1[n - N + 2] & \cdots & x_1[n] \\ x_2[n - N + 1] & x_2[n - N + 2] & \cdots & x_2[n] \\ \vdots & \vdots & \ddots & \vdots \\ x_M[n - N + 1] & x_M[n - N + 2] & \cdots & x_M[n] \end{bmatrix}$$

3.2. Frequency Domain

To analyze the spectral properties of each signal, a discrete Fourier transform (DFT) is applied row-wise to the data matrix D . For a discrete-time signal $X_i[n]$, its DFT over the window of size N is defined as:

$$f_i[k] = \sum_{n=0}^{N-1} X_i[n] \cdot e^{-i2\pi \frac{kn}{N}}, \quad k = 0, 1, \dots, N-1$$

where $f_i[k]$ represents the complex frequency component at bin k , and i is the imaginary unit. This transform yields a frequency-domain representation of each signal, capturing its harmonic content within the sliding buffer. The entire frequency-domain data across all M signals is represented by the matrix $F \in \mathbb{C}^{M \times N}$, where:

$$F = \begin{bmatrix} f_1[0] & f_1[1] & \cdots & f_1[N-1] \\ f_2[0] & f_2[1] & \cdots & f_2[N-1] \\ \vdots & \vdots & \ddots & \vdots \\ f_M[0] & f_M[1] & \cdots & f_M[N-1] \end{bmatrix}$$

Each row corresponds to a signal's frequency spectrum, and each column represents a particular frequency bin across all signals.

As long as the incoming signal data is real-valued, the Fourier values will exhibit *Hermitian symmetry*, meaning:

$$X[k] = \overline{X[N-k]}, \quad k = 1, 2, \dots, N-1,$$

where $\overline{X[\cdot]}$ denotes the complex conjugate. This symmetry implies that the frequency components for positive frequencies uniquely determine those for negative frequencies, allowing the Fourier transform of a real signal to be fully represented by the first $L = \lfloor \frac{N}{2} \rfloor + 1$ frequency bins (including the zero-frequency and, if N is even, the Nyquist frequency). Exploiting this property, matrix F can be safely reduced to $F_r \in \mathbb{C}^{M \times L}$:

3.3. Cosine Similarity

To identify signal pairs with similar spectral profiles, Coherent Multiplex computes the pairwise cosine similarity between rows of the reduced frequency matrix $F_r \in \mathbb{C}^M$. Given two complex-valued frequency vectors f_i and f_j , their cosine similarity is defined as:

$$s(i, j) = \frac{|\langle f_i, f_j \rangle|}{\|f_i\| \cdot \|f_j\|},$$

where $\langle f_i, f_j \rangle$ denotes the complex inner product:

$$\langle f_i, f_j \rangle = \sum_{l=0}^{L-1} f_i[l] \cdot \overline{f_j[l]},$$

and $\|f_i\|$ is the Euclidean norm:

$$\|f_i\| = \sqrt{\sum_{l=0}^L |f_i[l]|^2}.$$

This metric yields a real-valued similarity score between 0 and 1, where 1 indicates perfect spectral alignment. The resulting vector of similarities is defined as:

$$S = [s_1, s_2, \dots, s_p],$$

where $p = \binom{M}{2}$ is the number of unique unordered signal pairs.

3.4. Exploratory Network

The similarity vector S is used to construct a fully connected graph to represent the relationships among the M signals. Each signal corresponds to a node in the network, and every pair of nodes (i, j) is connected by an edge weighted by the cosine similarity score $s(i, j)$.

Formally, the network $G = (V, E, W)$ is defined as follows:

- $V = \{1, 2, \dots, M\}$ is the set of nodes, each representing a signal.
- $E = \{(i, j) \mid i, j \in V, i \neq j\}$ is the set of edges connecting every distinct pair of nodes.
- $W : E \rightarrow [0, 1]$ assigns weights to edges, where $W(i, j) = s(i, j)$ is the cosine similarity between signals i and j .

This weighted network layer encodes the spectral similarity structure among signals, enabling subsequent layer construction.

3.5. Wavelet Coherence

To uncover time-localized, frequency-specific relationships between signals, the Coherent Multiplex contains a second, sparse layer where the edges represent *wavelet coherence*.

To compute coherence, wavelet transform arrays must be constructed for participating nodes. Given a signal x , the wavelet transform [3] is computed as:

$$W_x(t, s) = \sum_{n=0}^{N-1} x[n] \cdot \psi^* \left(\frac{n-t}{s} \right),$$

where ψ is a complex-valued mother wavelet, s is the scale parameter (related to frequency), t is the time index, and ψ^* denotes the complex conjugate.

The wavelet coherence [2] between signals x_i and x_j is then defined as:

$$C_{ij}(t, s) = \frac{\left| S \left(W_i(t, s) \cdot \overline{W_j(t, s)} \right) \right|^2}{S(|W_i(t, s)|^2) \cdot S(|W_j(t, s)|^2)},$$

where $S(\cdot)$ is a smoothing operator in both time and scale, and the result satisfies $0 \leq C_{ij}(t, s) \leq 1$. Edges of this layer are constructed and decay based on the corresponding edge's power from the previous exploratory layer.

4. Prototyping

To evaluate the Coherent Multiplex pipeline, a full simulation and analysis environment has been implemented. The following subsections describe each stage of the prototype, along with representative figures.

4.1. Signal Simulation

Incoming data is procedurally generated to show signals that come in and out of covariance. 8 input signals are synthesized as linear combinations of sine waves with randomized amplitudes, frequencies, and phases. This emulates the complex, non-stationary behavior of real-world signals. Coherence is artificially introduced between selected pairs by sharing frequency components or phase structure at randomized intervals.

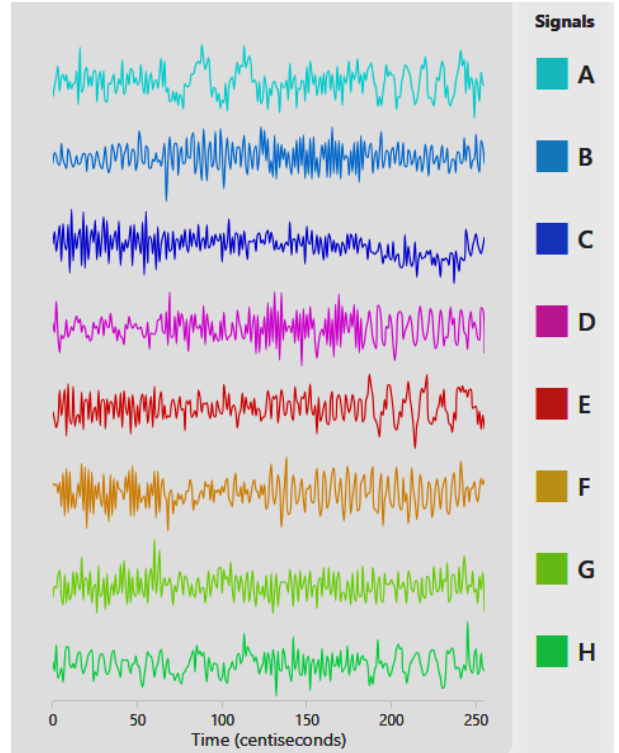


Figure 1: Signals A-H are color coded for easy tracking.

Signals are processed in real-time using a rolling buffer, and maintain a window of 256 samples. The model is set to produce and kick one sample each second.

4.2. Buffering and FFT Analysis

For each update, the FFT is computed using a fast backend with pre-aligned reusable memory buffers. This implementation imports the FFTW C sub-routine, which will enable scaling up. The magnitude spectra of the FFT results are extracted and used to compute pairwise cosine similarity. This provides a first-pass screening of potentially coherent signal pairs based on their frequency-domain peaks.




Figure 2: FFT magnitude spectra of each signal at a given time step. Pairwise cosine similarities are computed between rows.

4.3. Graph Construction

A fully connected network is made to represent the 8 signals as nodes, and the edges between nodes are weighted by the cosine similarity between the corresponding signals' frequency-domain vectors. This network serves as the first layer of analysis, allowing visualization of potential coherence relationships and identification of candidate pairs for further analysis.

A wider, darker edge indicates a small angle between the two nodes that it connects. At the timestep that Figure 1 was captured, the highest similarities were between pairs A-E, A-H, and E-H, which is reflected by the darkest edges linking nodes A, E and H in Figure 3.




Figure 3: Cosine similarity graph. Nodes represent signals; edge weights represent similarity. All possible pairs are connected.

4.4. Wavelet Coherence

For deeper analysis in time and frequency relationships, wavelet coherence is selectively computed between the pair exhibiting the highest similarity at each update. Transform arrays are efficiently computed via the FCWT algorithm for each signal, and coherence is evaluated in the time-frequency domain. Visualizations in this implementation include phase arrows indicating lead/lag relationships but ignore the cone of influence, which masks regions that are unreliable.

5. Applications

Though originally developed for examining multi-decade correlations between the El Niño-Southern Oscillation and monsoon activity in South Asia [3], wavelet coherence has been applied broadly across numerous domains. Applications with interest in these time-frequency relationships include economics, neuroscience, mechanical and electrical engineering and bio-medical signals.

In economics, wavelet coherence analysis is widely used to investigate the dynamic relationships between financial time series, such as stock indices, exchange rates, and macroeconomic indicators. By allowing researchers to examine correlations at different time scales and periods, wavelet coherence has enabled the identification of transient co-movements and structural changes in markets that may not be visible through traditional meth-




Figure 4: Wavelet coherence analysis between selected signal pairs. High-coherence regions and phase relationships are visible.

ods.

In neuroscience, wavelet coherence provides a powerful tool for exploring the synchronization and functional connectivity between different regions of the brain. For example, it has been applied to electroencephalogram (EEG) and magnetoencephalogram (MEG) data to detect time-varying patterns of coherence between neural signals, helping to identify underlying mechanisms for cognitive processes, brain disorders, and responses to stimuli. Coherence is also used to examine correlations across physiological processes such as heart rate and respiration, enabling detection of abnormal patterns and assessment of disease impacts.

In all of these fields, coherence is studied by first selecting signals of interest from a large pool. EEG for research will regularly employ 128 - 512 nodes in high-density mappings. The NYSE National platform continuously updates the prices of over 8000 listed securities. Traditional coherence implementations would never be considered as methods of extraction for the entire corpus. The Coherent Multiplex is capable of executing these computations at scale, and is highly optimized for performance. The model can be tuned to accommodate hardware, and the most recent prototype is hosted smoothly on a basic dual-core, 3.5 GB RAM cloud platform.

6. Future Work

The Multiplex architecture described in this whitepaper can be summarized as heuristics for distinguishing signal pairings worthy of deeper analysis at a given time. The "heuristics" and "deeper analysis" in this case refer specifically to FFT, cosine similarity, and wavelet coherence, thought this can be appended or substituted depending on the goal. If scale-independency is not necessary for the

object of study, the *cross-wavelet* transform can be substituted in for coherence. A different metric could be used to determine the edge powers in the graph, supposing the spectral content is not the object of interest in the incoming signals.

Additional layers, both sequential and non sequential will be appended to the multiplex. Wavelet coherence can be calculated alongside traditional metrics like macro trends for high-powered edges in the graph, giving a more complete extraction of possible features. Additionally, subsequent layers with trained convolution neural networks will be added for automated analysis and description of the coherence readings.

References

- [1] M. Frigo and S. G. Johnson, *The Design and Implementation of FFTW3*, Proceedings of the IEEE, vol. 93, no. 2, pp. 216–231, 2005.
- [2] A. Grinsted, J. C. Moore, and S. Jevrejeva, *Application of the cross wavelet transform and wavelet coherence to geophysical time series*, Nonlinear Processes in Geophysics, vol. 11, no. 5/6, pp. 561–566, 2004.
- [3] C. Torrence and G. P. Compo, *A Practical Guide to Wavelet Analysis*, Bulletin of the American Meteorological Society, vol. 79, no. 1, pp. 61–78, 1998.
- [4] Arts, L.P.A., van den Broek, E.L. *The fast continuous wavelet transformation (fCWT) for real-time, high-quality, noise-resistant time-frequency analysis*. Nat Comput Sci 2, 47–58, 2022.