Title: Sample, Don't Optimize: Quantum computing's new playbook for molecular electronic structure

Abstract: Quantum computers were long expected to attack quantum chemistry's hardest problems by directly optimizing molecular wavefunctions on noisy hardware — the Variational Quantum Eigensolver (VQE) approach. But in practice, this strategy stalled: measurement overhead, barren plateaus, and compounding hardware noise kept VQE demonstrations confined to a handful of atoms for years. This talk traces the field's subsequent pivot: rather than trusting a noisy quantum device to *optimize*, recent methods use it only to *sample* — proposing candidate electron configurations that a classical computer then diagonalizes, in the spirit of selected configuration interaction. This shift, embodied in Sample-based Quantum Diagonalization (SQD) and its self-consistent recovery loop, proved far more robust to noise and rapidly scaled from small molecules to iron-sulfur clusters relevant to bioinorganic chemistry. The talk will highlight recent refinements — including physics-informed generative screening (PIGen-SQD) and density-matrix embedding combined with SQD, both developed at IIT Bombay — that sharpen this approach using chemical intuition rather than brute computational methods. The talk closes with the field's current frontier: a fragment-embedded quantum-classical workflow that recently simulated protein-ligand complexes exceeding 12,000 atoms, a scale simply not reachable on a quantum computer just two years ago. That trajectory — from a few qubits to industrially relevant biomolecules in under three years — captures how rapidly quantum computing and chemistry have converged since SQD's introduction, and where this fast-moving intersection may be headed next.