SBC Takeaways

Quite frankly, this year, SBC was a bit underwhelming, with papers and research being focused more on technical niche feature upgrades rather than foundational progress that demands venture capital attention. The era of raising $50M+ in a seed round for a new L1 by way of a new consensus mechanism paper is long gone. I enjoyed some papers and research areas attached below:

- **Elliptic Curve Cryptocurrencies and Quantum Vulnerabilities (Paper):** This paper provides resource estimates for breaking the 256-bit Elliptic Curve Discrete Logarithm Problem (used in Bitcoin and Ethereum) via Shor's algorithm, demonstrating feasibility with <1200 logical qubits and <90 million Toffoli gates. The authors analyze quantum threat timelines to blockchain systems and propose mitigation strategies for the eventual emergence of cryptographically relevant quantum computers.
- **Weighted Batched Threshold Encryption with Applications to Mempool Resilience (Paper):** This paper develops weighted batched threshold encryption schemes with direct applications to mempool privacy in blockchain systems. The work addresses how to protect transaction ordering and hide pending transactions while maintaining protocol security.
    - There was a big research focus on Monad and Category Labs, and their protocol design. They focused on describing Monad's mempool, their new mechanism for consensus design for a higher-performance L1, and their initial ideas for an inference market.
- **The Economics of Proof-of-Useful-Work (Paper):** This paper analyzes the economic viability of proof-of-useful-work (PoUW) blockchains, where computational work secures consensus while simultaneously generating economic value (e.g., ML inference). The author develops a competitive-equilibrium model characterizing how duplex overhead and economic parameters affect compute allocation between pure mining, useful work, and hybrid activity, addressing critiques that useful work creates "paid-to-attack" vulnerabilities.

As is evident here, people have realized that the L1 market is an oligopoly, where you differentiate on the spectrum of permissionlessness / extent of decentralization (ties into speed of execution and control), app customizability (goes back to sequencing and governance), and enshrined features that enable natural distribution. Hence, what they are trying to do is pitch research that can be commercialized as off-the-shelf features for these L1s. A prime example is a quantum strategy or an off-the-shelf privacy solution, all of which can ideally be sold in a "SaaS"-type business model. Not sure if this is venture-backable.

I was also able to gain exposure to some cool frontier AI products and builders. While there was a lot of noise around data sets and environments for training, local models and harnesses, and other AI-equipped consumer tools, there were a select-few that I thoroughly enjoyed using. I have linked these below:

- AGI Inc.: Neutral on-device AI model (optimized by test-time compute methods and computer-use agents) for OEMs and businesses looking to deploy an AI strategy.
- Crater: Predictive desktop assistant.
- Osaurus AI: Gamified local agentic harness which develops your memory and provides access to an inclusive set of inference partners, starting with Venice AI.