# Research Biography

**Mahadev Satyanarayanan**

## A. Distributed File Systems (early 1980s to early 1990s)

### A-1. Andrew File System (AFS)

Key concept: *Scalable, secure, and seamless data caching from servers*

Impact

- Deployed at CMU, MIT, Michigan, and hundreds of universities and companies worldwide for many decades (some continue to present)
- Commercialized by Transarc, IBM, and open-sourced since 2001
- Inspiration for DropBox, Microsoft OneDrive Files-on-Demand, and many other caching-based distributed file systems

Key papers: Satya1985, Howard1988, Spasojevic1996

Awards

- ACM SOSP 1985 and SOSP 1987 Outstanding Paper Awards
- 2008 ACM SIGOPS Hall of Fame Award
- 2016 ACM Software System Award

Starting in 1983, Satya was a principal architect and implementer of the Andrew File System (AFS), the technical forerunner of modern cloud-based storage systems. AFS has been continuously deployed at CMU since 1986, at a scale of many thousands of users. From its conception in 1983 as the unifying campus-wide IT infrastructure for CMU, AFS evolved through versions AFS-1, AFS-2 and AFS-3. In mid-1989, AFS-3 was commercialized by Transarc Corporation and its evolution continued outside CMU. Transarc was acquired by IBM, and AFS became an IBM product for a number of years. In 2000, IBM released the code to the open source community as OpenAFS. Since its release as OpenAFS, the system has continued to be used in many enterprises all over the world. In the academic and research lab community. OpenAFS has been used at more than 30 sites in the United States (including CMU, MIT, and Stanford) and dozens of sites in Europe, New Zealand, and South Korea. Many global companies have used OpenAFS including Morgan Stanley, Goldman Sachs, Qualcomm, IBM, United Airlines, Pfizer, Hitachi, InfoPrint, and Pictage.

Over a 30-year period, AFS has been a seminal influence on academic research and commercial practice in distributed data storage systems for unstructured data. Its approach to native file system emulation, scalable file caching, access-control based security, and scalable system administration have proved to be of enduring value in enterprise-scale information sharing. The design principles that were initially discovered and validated in the creation and evolution of AFS have influenced virtually all modern commercial distributed file systems, including Microsoft DFS, Google File System, Lustre File System, Ceph, and NetApp ONTAP. In addition, AFS inspired the creation of DropBox whose founders used AFS as part of Project Athena at MIT. It also inspired the creation of Maginatics, a startup company advised by Satya that provides cloud-sourced network-attached storage for distributed environments and was purchased by EMC in 2014 (now part of Dell). The NFS v4 network file system protocol standard has been extensively informed by the lessons of AFS.

The 2016 ACM Software Systems Award recognized this work as follows:

> *"AFS was the first distributed file system designed for tens of thousands of machines, and pioneered the use of scalable, secure and ubiquitous access to shared file data. To achieve the goal of providing a common shared file system used by large networks of people, AFS introduced novel approaches to caching, security, management and administration. AFS is still in use today as both an open source system and as the file system in commercial applications. It has also inspired several cloud-based storage applications."*

## A-2. Coda File System

Key concept: *Disconnected and Weakly-Connected Operation*

Impact
- Inspiration for "Offline Access" in Google Docs, Microsoft 365, Microsoft Outlook, and DropBox
- Inspired Lustre File System, used today in virtually all supercomputing systems (HPC systems)

Key papers: Satya1990, Kistler1992, Satya1993, Mummert1995, Satya2002

Awards
- ACM SOSP 1991 Outstanding Paper Award
- 2015 ACM SIGOPS Hall of Fame Award
- 2016 ACM SIGMOBILE Test-of-Time Award
- 1999 LinuxWorld Editor's Choice Award

In 1987, Satya began work on the Coda File System to address a fundamental shortcoming of AFS-like systems. Extensive first-hand experience with the AFS deployment at CMU showed that users are severely impacted by server and network failures. This vulnerability is not just hypothetical, but indeed a fact of life in real-world deployments. Once users become critically dependent on files cached from servers, a server or network failure renders these files inaccessible and leaves clients crippled for the duration of the failure. In a large enough system, unplanned outages of servers and network segments are practically impossible to avoid. Today's enthusiastic embrace of cloud computing rekindles many of these concerns because of increased dependence on centralized resources. The goal of the Coda project was to preserve the many strengths of AFS, while reducing its vulnerability to failures. Over 30+ years, research on Coda has proved to be highly fruitful in creating new insights and mechanisms for failure-resilient, scalable and secure read-write access to shared information by mobile and static users over wireless and wired networks. Coda was the first system to show how server replication could be combined with client caching to achieve good performance and high availability. Coda invented the concept of "disconnected operation", in which cached state on clients is used to mask network and server failures. Coda also demonstrated bandwidth-adaptive weakly-connected operation over networks with low bandwidth, high latency or frequent failures. Coda's use of optimistic replication, trading consistency for availability, was controversial when introduced. Today, it is a standard practice (under the name "eventual consistency") in all data storage systems for mobile environments. Coda also pioneered the concept of translucent caching, which balances the full transparency of classic caching with the user visibility needed to achieve a good user experience on bandwidth-challenged networks. The Coda concepts of hoarding, reintegration and application-specific conflict resolution are found in the cloud sync capabilities of virtually all mobile devices today. Key ideas from Coda were incorporated by Microsoft into the IntelliMirror component of Windows 2000 and the Cached Exchange Mode of Outlook 2003.

The 2015 ACM SIGOPS Hall of Fame Award recognized this work as follows:

> *"This paper was the first to describe the use of caching to provide availability in addition to improved performance in a distributed setting where clients use files stored at remote file servers, leading to potential loss of service during disconnection. The Coda design provided a thoughtful and elegant approach to supporting continued service during disconnection. Disconnected clients continued to service user requests using locally cached content; however all potential modifications performed while disconnected were logged locally, and when service was restored the system attempted to reconcile the local modifications with the current server state. The Coda design inspired much follow-on research on distributed file systems and its techniques were adopted in other systems."*

The 2016 ACM SIGMOBILE Test-of-Time Award recognized this work as follows:

> *"The Coda file system remains influential today as did in 1990s when it first appeared. Coda first recognized that the future mobile devices and their file systems should be designed to operate in disconnected environments. Coda treated disconnection as a mode of operation rather than failure. Examples of its impact are all around us from Google Drive to Dropbox."*

## B. Mobile Computing (early 1990s to late 2000s)

### B-1. Disconnected and Weakly Connected Operation *(see A-2 above)*

### B-2. Offloading or "Cyber Foraging" by mobile devices

Key concept: *Remote execution to amplify compute power of lightweight, resource-poor, battery-dependent mobile devices*

Impact

- Used by original implementation of Apple Siri speech recognition
- Cloud offload is used by many Android and iOS apps today
- Used extensively for "Edge AI" applications today

Key papers: Noble1997, Satya2001a, Flinn2001, Flinn2002, Balan2002, Balan2007, Satya2021

In 1993, Satya wrote a short thought piece that identified the inherent resource poverty of mobile devices as a key long-term constraint of mobile computing. As the article states: "Mobile elements are resource-poor relative to static elements. Regardless of future technological advances, a mobile unit's weight, power, size and ergonomics will always render it less computationally capable than its static counterpart. While mobile elements will undoubtedly improve in absolute ability, they will always be at a relative disadvantage." In the 30+ years since that article, the prediction has remained consistently true, a phenomenon now referred to as "mobility penalty.". In 1997, Satya and his team proposed a solution that is now widely used for many compute-intensive applications: a mobile device offloads heavy computations over a wireless network to a server that is much more powerful than the mobile device. The approach was first demonstrated in the Odyssey system. An important aspect of this implementation was Odyssey's ability to select the optimal execution mode (local, remote or hybrid) based on runtime factors such as current network bandwidth. Today, the hybrid mode of execution is known as "Split Computing." Odyssey was thus the technical forerunner of today's mobile speech-to-text systems, as well as modern mechanisms for adaptive offload. Since 2001, this approach has also been known by the term *cyber foraging* and has been a key area of mobile computing research. The evolutionary path from its 1997 origin to modern mobile and IoT systems is described in Satya's 2014 retrospective "A Brief History of Cloud Offload: A Personal Journey from Odyssey Through Cyber Foraging to Cloudlets."

### B-3. Multi-fidelity Computation

Key concept: *Match application quality to resource availability*

Impact

- used by YouTube, Netflix, and many other video streaming services for variable bitrate video delivery

- used in Intermittent Computing for energy-harvesting devices

Key papers: Noble1997, Satya2001b, Narayanan2003

Awards

- 2024 ACM SIGMOBILE Test-of-Time Award

In the mid-1990s, Satya initiated the Odyssey project to explore how operating systems should be extended to support future mobile applications. While Coda supported mobility in an application-transparent manner, Odyssey explored the space of application-aware approaches to mobility. Wireless network bandwidth and energy (i.e., battery life) were two of the key resource challenges faced by mobile applications. Odyssey invented the concept of application-aware adaptation and showed how the system call interface to the Unix operating system could be extended to support this new class of mobile applications such as video delivery and speech recognition. Odyssey envisioned a collaborative partnership between the operating system and individual applications. In this partnership, the operating system monitors, controls and allocates scarce resources such as wireless network bandwidth and energy, while the individual applications negotiate with the operating system on their resource requirements and modify application behavior to offer the best user experience achievable under current resource conditions. The concepts of multi-fidelity algorithms and predictive resource management that emerged from this work have proved to be highly influential. In particular, adapting video delivery quality to available network bandwidth is standard practice in all streaming video services such as Netflix and YouTube today.

The 2024 ACM SIGMOBILE Test-of-Time Award recognized this work as follows:

> *"The Odyssey paper introduced fidelity—the degree to which data presented on a client matches the server's reference copy—as a critical factor for mobile applications. It pioneered the concept of application adaptation for mobile environments, extending operating systems to support this adaptation across a wide range of mobile information access applications. Odyssey's notion of application-aware adaptation has had a profound and lasting impact on the evolution of mobile operating systems."*

## B-4. Energy-Aware Adaptation

Key concept: *Adapt application behavior to reduce energy demand*

Impact

- widely used technique for mobile devices today
- "zoned backlighting" concept used by OLED displays today
- "PowerScope" measurement technique widely used today

Key papers: Flinn1999a, Flinn1999b

Awards

- 2020 ACM SIGMOBILE Test-of-Time Award

Energy is a vital resource for mobile computing. Today, there is widespread consensus that advances in battery technology and low-power circuit design cannot, by themselves, meet the energy needs of future mobile computing systems — the operating system and higher levels of the software stack must also be involved. In 1999, Satya and his team demonstrated that a collaborative relationship between the operating system and applications can be used to meet user-specified goals for battery duration.

The 2020 SIGMOBILE Test-of-Time Award recognized this work as follows::

> *"This paper describes how an OS and applications can collaborate to adapt energy use on a mobile device. While at the time this was new ground, the approaches it presents are now simply standard practise, in the devices and apps we use every day. It is remarkable how prescient this 1999 paper was, and how well its ideas have stood the test of time. For example, the hardware and software approaches to energy adaptation it proposes are now commonplace, and the applications it analyses (browser, voice recognition, video player and maps) are still ubiquitous. This paper presaged a huge amount of work in energy adaptation by both academia and industry, and the area remains of core interest to SIGMOBILE."*

## C. Internet of Things (early 2000s to present, aka "Pervasive" or "Ubiquitous" Computing)

### C-1. IoT Framework

Key concept: *IoT as derivative of Mobile Computing and Distributed Systems*

Impact

- Second most widely cited paper in the IoT literature (close to 4000 citations), exceeded only by Mark Weiser's original 1991 Scientific American paper.

Key paper: Satya2001

Awards

- 2018 ACM SIGMOBILE Test-of-Time Award

In the late 1990s, Satya initiated the Aura Project in collaboration with CMU faculty colleagues David Garlan, Raj Reddy, Peter Steenkiste, Dan Siewiorek and Asim Smailagic. The challenge addressed by this effort was to reduce *human distraction* in mobile and pervasive computing environments, recognizing that human attention does not benefit from Moore's Law, while computing resources do. This leads directly to the notion of invisible computing, which parallels Mark Weiser's characterization of an ideal technology as one that disappears. The Aura vision proved to be an excellent driver of research in mobile and pervasive computing in areas such as cyber foraging, location-aware computing, energy-awareness, and task-level adaptation. Based on insights from Aura, Satya wrote an invited paper in 2001 entitled "Pervasive Computing: Vision and Challenges." The concepts discussed in this paper have directly inspired today's drive towards an "Internet of Things (IoT)."

In 2018, this visionary paper was recognized by the ACM SIGMOBILE Test-of-Time Award. The citation for the award stated:

> *"The paper connects the vision of pervasive computing to distributed systems and mobile computing as we knew them then, then draws fundamental observations of what system components still needed to be developed and how. The paper is a travel in time. What we call today the "œInternet of Things" was already described here, along with many other fundamental concepts such as edge computing, energy-driven adaptation, thick and thin clients. As an eminent example of abstract thinking, the author revealed the essence of each and every research challenge independent of the technology available back then."*

## C-2. Privacy and Denaturing of IoT Data

Key concept: *Redaction of data close to capture, under user control*

Key papers: Simoens2013, Davies2016, Wang2017, Satya2022, Satya2022b

Awards

- 2017 ACM MMSys Best Paper Award

Video data is extremely effective in capturing both the spatial and temporal aspects of real-world events. If its bandwidth and storage demands can be met, there is no better way to capture and preserve events, especially if accompanied by audio and other less resource-hungry sensor data. Unfortunately, there are legitimate public concerns about the privacy compromises implicit in widespread use of always-on video cameras. Since 2013, Satya has been exploring privacy-preserving transformation of video data. *Denaturing* refers to the process of making video data "safe" from a privacy viewpoint. A denatured video stream is one whose content has been analyzed to identify potential privacy leaks under the guidance of a privacy policy, and has been modified to eliminate those leaks. Once denatured, a video stream may be safely released to untrusted video analytics software. Real-time denaturing of video streams, shown to be feasible by the 2017 work of Satya's team (Wang2017), is a "killer app" for edge computing.

## C-3. Data Discovery from Uncurated and Unlabeled Sources

Key concept: *Human-in-the-loop context-aware search of raw data*

Key papers: Huston2004, Satya2010, Satya2017a, Feng2018, Feng2023, George2023a, George2023b

Awards

- 2017 IEEE GIOTS Best Paper Award

For over two decades, Satya has been working on the problem of human-in-the-loop search of raw image and video data (prior to curation and indexing) for rare "needles in the haystack." In today's machine learning (ML) parlance, this is the problem of discovering a sufficient number of true positives (TPs) in an extreme class imbalance setting to construct the training set for creation of an accurate ML model. The earliest part of this work was done in the Diamond project, using explicit features such as color, texture, SIFT and HOG. To support interactive discovery of TPs, the OpenDiamond platform provided a storage architecture for *discard-based search* that pipelines user control, feature extraction, and per-object indexing computation and result caching. Diamond attracted interest from the medical and pharmaceutical research communities, resulting in many successful collaborations and leading to the creation of OpenSlide, a widely-used vendor-neutral open source library for digital pathology. Diamond has evolved into Eureka (2018), and then into Hawk (2023), which seamlessly pipelines semi-supervised learning, active learning, and transfer learning, with asynchronous bandwidth-sensitive data transmission for labeling.

## D. Edge Computing (2009 to present)

Key concept: *Bringing the cloud closer*

Impact

- Wearable Cognitive Assistance ("angel on your shoulder" Edge AI)
- Autonomy for ultralight drones
- Archival software reproducibility
- new market valued at $16 billion in 2023, with expected annual growth rate of ~35% from 2024 to 2030.

Key papers: Satya2009, Ha2014, Satya2017b, Wang2019, Satya2019, Satya2023, Bala2023, Bala2024

Awards

- 2022 ACM SIGMOBILE Test-of-Time Award

In 2009, together with Victor Bahl, Ramon Caceres and Nigel Davies, Satya published a paper entitled 'The Case for VM-based Cloudlets in Mobile Computing'. Today, the work is viewed as the founding manifesto of edge computing. At a time when there was widespread euphoria about the limitless possibilities of cloud computing, the paper put forward a contrarian viewpoint. It argued that the extreme consolidation (the concentration of computing resources into a few large data centers) implicit in cloud computing would fundamentally limit its ability to sustain latency-sensitive and bandwidth-hungry applications that would emerge in the future. It also identified bandwidth scalability of cloud-based applications based on video data from sensors as a key concern. To support these future applications, the paper argued for a dispersed infrastructure of micro-data centers called *cloudlets*, which avoids extreme consolidation while preserving cloud computing attributes such as multi-tenancy with strong isolation. With the emergence of edge computing, this vision has become mainstream. Edge computing is still in its infancy. Market estimates place the 2023 global market size of edge computing at $16 billion, growing to $150 billion by 2030 at a compound annual growth rate of more than 36%. In the coming years, edge computing will be the enabler of new IoT, cyber-human and cyber-physical edge-native applications that are transformative. Wearable Cognitive Assistance, in which AI algorithms such as object recognition are executed in the inner loop of latency-sensitive Augmented Reality applications on wearable devices, is one such class of cyber-human applications. Fully autonomous flight by ultralight drones that use edge offload is an example of a cyber-physical edge-native application.

The 2022 ACM SIGMOBILE Test-of-Time Award recognized the significance of the 2009 paper as follows:

> *“What is known as edge computing today can be traced back to this Cloudlets paper, which proposed the need to support localized data and computation, thus avoiding the high latency and bandwidth limitations associated with traditional cloud computing.”*