10 Years of Bismuth

Architectural anticipation in a fair-launched blockchain, 2017–2026

Bismuth Foundation — June 2026

Abstract

Bismuth is a from-scratch, fair-launched blockchain whose mainnet began on 1 May 2017 with no initial coin offering and no premine [1, 2]. Across its first decade the project repeatedly shipped designs that the wider industry arrived at independently, and named, years later — and it did so with unusual rigour: its difficulty controller and block-validation changes were analysed in the peer-reviewed control-systems literature as early as 2017 [3, 4, 5]. This article surveys Bismuth's public record — its whitepaper, its peer-reviewed papers, its development blog (2018–2026), and a documented corpus of application protocols [6] — and maps each early design decision to the mainstream pattern it anticipated, with dates. The cases examined include: the separation of consensus over data from consensus over meaning (later: modular blockchains and data-availability layers; Bitcoin inscriptions and off-chain indexers); a plugin "building-blocks" architecture and a documented ecosystem of on-chain protocols (tokens, aliases, decentralized social, naming/DNS, event sourcing, state channels, NFTs); confidential assets (2019); on-chain governance (2019); the recording of real-world sensor data on-chain in 2020, two years before "DePIN" was coined; and cross-chain bridges (2021). We argue that Bismuth functioned, largely unheralded, as a long-running and academically grounded testbed for ideas that later defined market cycles, and we identify the single recurring caveat its own architecture exposes: where interpretation carries value, the semantic layer must re-converge on one canonical engine.

---

1. Introduction

Most narratives of blockchain progress are organised around the projects that captured attention during a given market cycle. This survey takes a different vantage point. It examines a single, continuously developed chain — Bismuth — whose documented design decisions frequently preceded, by months or years, the patterns later popularised under names such as "modular blockchains," "inscriptions," "DePIN," and "DAO governance," and several of which were subjected to formal academic analysis at the time.

The method is deliberately conservative. Each claim of anticipation is anchored to a dated primary source from Bismuth's own record — whitepaper, peer-reviewed paper, blog entry, or protocol specification — and compared against a dated secondary source marking when the corresponding idea entered mainstream practice. Where Bismuth's artifact clearly predates the mainstream terminology,

standardisation, or market emergence of an idea, we say it anticipated that idea; we do not claim invention of primitives that have older, independent lineages.

Two properties frame everything that follows. First, Bismuth launched **fairly** — "no premine, no ICO" [1, 2] — and was, by several contemporary accounts, the **first blockchain written in Python** [7], a from-scratch codebase rather than a fork of Bitcoin or Ethereum. This gave its authors unusual latitude to experiment at the protocol layer, but it also removed the marketing apparatus that typically converts technical priority into public recognition. Second, the project engaged the academic literature directly: its consensus mechanism was the subject of refereed publications in Modeling, Identification and Control [3, 4], a discipline-of-control journal, well before "blockchain research" became a crowded field.

2. A fair-launched, modular base layer (2017–2018)

Bismuth's mainnet launched on 1 May 2017 [1]. Its base layer combines a SHA-224 proof-of-work function (the CPU/GPU-friendly Heavy3), a 60-second target block time, and per-block difficulty retargeting via a PID controller — a control-theoretic approach to difficulty markedly more responsive than the windowed moving averages common at the time [2]. Block rewards decay smoothly from roughly 15 BIS toward a perpetual tail emission of 0.5 BIS per block, deliberately avoiding the discontinuous "halving" cliffs of Bitcoin-derived schedules [2].

Two choices in the transaction model proved especially consequential. Bismuth uses an **account-based model with a signature-derived transaction identifier**, achieving replay protection through the intrinsic uniqueness of each signature rather than Ethereum-style nonces [2]. And every transaction carries an **`operation` field and an arbitrary `openfield` data field**, so that custom semantics — `token:issue`, `token:transfer`, and others — are expressed as data on ordinary transactions rather than as bespoke protocol changes [2, 6]. This "arbitrary data plus custom operations" pattern is, in retrospect, a direct ancestor of the inscription model discussed in §4 and §5.

The architecture was described from the outset as a "building-blocks" system with optional, plugin-based extensibility [2], and was paired with a second consensus tier — **HyperNodes**, a proof-of-stake layer that recorded network key-performance metrics on its own loosely coupled chain rather than holding currency [2]. The dual-tier design operated from block 800,000 and was later retired at the HF4 hard fork (block 4,380,000) — an unusually candid example of a project removing a major subsystem once its purpose had passed [2].

3. Peer-reviewed foundations: difficulty control as a control-systems problem (2017–2018)

Where most chains treat difficulty adjustment as folklore, Bismuth's was formally modelled. In Nonlinear Feedback Control and Stability Analysis of a Proof-of-Work

Blockchain (Hovland & Kučera, Modeling, Identification and Control 38(4):157–168, 2017) [3], the difficulty mechanism is cast as a feedback control loop in which "the controller output equals the difficulty adjustment in the mining process while the feedback variable is the average block time," with hash power modelled as a disturbance. The paper demonstrates "stability and a fast response" across constant, step, ramp, and high-frequency sinusoidal disturbances, validated in both simulation and on the live Bismuth testnet and mainnet [3]. Casting PoW difficulty as a control problem, and analysing its stability formally, in 2017 anticipated a subsequent academic thread on control-oriented modelling of proof-of-work — for example Leva, Strada & Tanelli's IEEE work on control-oriented PoW modelling [9] and Bissias, Thibodeau & Levine's Bonded Mining [10].

The same authors followed with Tail Removal Block Validation: Implementation and Analysis (Kučera & Hovland, MIC 39(3):151–156, 2018) [4], which proposes, implements, and measures a fix for long-tail block times; on Bismuth mainnet "the variances in the key variables, difficulty level and blocktime, were approximately halved after the tail removal code was enabled" [4]. A companion analysis treated the security of the chain [5].

Bismuth also drew third-party academic attention beyond consensus. It appears in the literature on blockchains as adaptive systems (Liaskos et al., SEAMS 2019) [11] and on Industry-4.0 knowledge representation via blockchain (Pinheiro, Santos & Barbosa, Springer 2018) [12] — the latter directly resonant with the on-chain machine-data work of §7 — and is cited among prior art in an IBM blockchain-performance patent [13]. For a project with no marketing budget, this is a notable density of refereed engagement, and it grounds the "ahead of its time" claim in something firmer than retrospective analogy.

4. Semantic interpretation: consensus on data versus consensus on meaning (July 2018)

The project's most prescient document is the essay Semantic Interpretation (19 July 2018), also issued as a paper [8, 14]. It proposes a three-layer architecture: a base layer holding **passive data** under consensus; a layer of **interpretation engines** that read that data and compute meaning, modifiable or replaceable independently; and a conflict-resolution layer in which users choose which interpretation to trust. Crucially — and contrary to a common misreading — the essay also specifies that where interpretation must be authoritative, it should collapse back to a single canonical engine. Its threat model is equally notable: it identifies the **code repository and its social consensus**, not the chain, as the true single point of failure, citing the 2017 Ethereum hard fork [8].

The wider industry reconstructed this architecture from several directions over the following five years:

- **Modular blockchains / data availability.** Al-Bassam's LazyLedger (May 2019), later **Celestia** (mainnet October 2023), reduced the base layer to consensus-plus-data-availability with execution elsewhere [15]. Ethereum's **rollup-centric roadmap** (Buterin, October 2020) made the same division [16].
- **Inscriptions and off-chain indexers.** Bitcoin **Ordinals** (Rodarmor, January 2023) and **BRC-20** (Domo, March 2023) put arbitrary data on L1 and let off-chain **indexers** interpret it into balances [17] — passive data plus interpretation engines, exactly. The ecosystem promptly hit the problem the 2018 essay had named: with no on-chain logic enforcing ownership, "the canonical interpretation is whatever the dominant indexer says it is," now discussed as **"indexer consensus"** [17]. That is precisely the case in which, per [8], the semantic layer must re-converge on one canonical engine.

5. The plugin system and a documented protocol ecosystem (2018–)

Three days before the semantic essay, Exploring the Bismuth Plugin System (16 July 2018) described how wallets, explorers, monitoring, token logic, and custom features attach as plugins rather than core modifications [18, 2] — the engineering counterpart to the semantic thesis: if meaning lives above consensus, the software computing it should be modular and replaceable.

That model is not merely theoretical; it is documented as a working developer corpus, Hack-with-BIS [6], which specifies application protocols built entirely as interpretations of the `operation`/`openfield` data model — a **token** protocol; an **alias** (human-readable naming) protocol; **Bisnet-DNS**; **event-sourcing** patterns and on-chain-event plugins; **off-chain secure storage** ("Scatter"); a **social** protocol; **state channels**; and worked dApps including a "simple dApp vs Solidity" comparison demonstrating contract-like applications without a Solidity-style VM [6].

Two of these specifications reward direct quotation, because their framing is strikingly early. The **state-channels** research note (2018) describes keeping consensus over results rather than content:

> "The concept of Bismuth state channels comes from an idea that it is possible to keep an objective consensus around behavior and results of on-chain contracts, without the need of the contract to be hosted on the Bismuth network itself. … we only need to keep consensus around contract behavior, exploiting Vitalik Buterin's data unavailability phenomenon for contract content hashes." [6]

It generalises the mechanism beyond contracts — "Basically a chainless consensus mechanism that can be applied to any situation where state/result coherence is

important” — and proposes it as an oracle substitute: “This concept can be also used instead of oracles for connecting contracts to the real world, as a checksum mechanism” [6]. Maintaining consensus over execution results while treating contract content as data-unavailable is precisely the division later formalised by optimistic rollups and modular data-availability designs (§4); advancing it, and an oracle alternative, in 2018 is markedly forward-looking.

The **social** protocol applies the same `operation`/`openfield` model to social data: a single operation registers a post on-chain — “operation ‘twitter’; openfield ‘tweet_id’ where tweet_id is the twitter id of the tweet to register” [6] — consumed by a rewards application (“Used by the twitterizer reward app” [6]). Recording social attestations on-chain as a first-class, interpretable operation in 2018 anticipated, in primitive form, the later decentralized-social movement (Nostr, Farcaster, Lens; 2022–2023). The corpus’s **event-sourcing** patterns — on-chain events consumed as a queryable index — likewise mirror the indexer/subgraph approach later standardised by services such as The Graph (mainnet 2020) [19], the same data-then-interpretation shape as inscriptions (§4).

The flagship example is **Dragginator**, a collectible-creature application (a “complex app” in the corpus [6]) active by mid-2018 [20] — i.e. non-fungible digital collectibles on a non-Ethereum chain contemporaneous with the standardisation of NFTs on Ethereum (EIP-721 was introduced in January 2018, following CryptoKitties in late 2017 [21]) and roughly three years before the 2021 NFT market cycle. That a small, fair-launched chain supported user-created collectibles at the moment the NFT concept itself was being formalised is a concrete instance of the pattern this article documents.

6. Confidential assets: shielded tokens (2019)

Introduction to Bismuth Shielded Tokens (20 September 2019), also issued as a paper [22], introduced privacy-preserving tokens, extending the token system of §2 and §5 with confidentiality. Privacy as an asset-layer property — rather than a whole-chain property as in dedicated privacy coins — anticipated the broader interest in confidential assets and selective on-chain privacy. The theme persists: the project’s current post-fork roadmap brings stealth addresses, ring signatures, and confidential amounts into the core.

7. On-chain governance (2019)

Through the second half of 2019 Bismuth executed a deliberate **governance shift**, documented in two parts (August 2019), and held **BGV-01**, its first formal governance vote (announced 10 September 2019; outcome 6 November 2019) [23]. Chain-anchored governance of protocol decisions predated the wave of DAO-governance tooling — off-chain voting systems such as Snapshot emerged around 2020 — and the broader normalisation of token-holder governance. The exercise is consistent with the 2018 threat model [8]: formalising governance removes the

repository-and-social-consensus single point of failure by making decisions explicit and accountable.

8. Physical-world data: DePIN before the name (2020)

In spring 2020 Bismuth published a sequence on recording real-world sensor data on-chain: Decentralized Condition Monitoring (30 April 2020), Android Battery Monitoring (9 May 2020), and Tesla Battery Monitoring (25 May 2020) [24]. These demonstrated industrial and consumer telemetry written to, and interpreted from, the chain — machine data as a first-class citizen, and a practical instantiation of the blockchain-for-industrial-knowledge thesis examined academically in [12].

The category this prefigures, **Decentralized Physical Infrastructure Networks (DePIN)**, did not acquire its name until late 2022, when the analytics firm Messari coined it; the sector's expansion to hundreds of projects came in 2022–2023 [25]. While individual antecedents exist (Filecoin 2014, Render 2017, Helium 2019 [25]), Bismuth's 2020 monitoring work placed verifiable physical-world data on a general-purpose chain more than two years before the surrounding category had a vocabulary.

9. Interoperability and programmability (2021)

In 2021 the project shipped **cross-chain bridges** — the "Crystal" connectors BIS↔ETH (21 April 2021) and BIS↔BSC (12 May 2021) [26] — during the same window in which cross-chain bridging became a central, and contested, theme of the multi-chain era. In parallel, Python and State Machines (24 March 2021) [27] articulated an on-chain programmability model built on explicit state machines, complementing the state-channel and dApp work of §5 [6] and the "Beyond DeFi" application framing of August 2020 [28].

10. Discussion

A consistent pattern emerges. A small team, working without ICO capital or a promotional apparatus, used a from-scratch (and first-in-Python [7]) codebase to ship working primitives — semantic data layering, modular interpretation, a documented protocol ecosystem, digital collectibles, confidential assets, on-chain governance, physical-world data, and bridges — ahead of the cycles that later popularised each, and did so with enough rigour to be analysed in refereed control-systems venues [3, 4]. The fair launch that lends the project its integrity is also, plausibly, why the work went largely uncited in popular accounts: there was no marketing engine to convert priority into recognition.

The retrospective also vindicates the project's own central caveat. The 2018 semantic model is correct that data and meaning can be decoupled, but it is equally clear — in the essay and now in practice — that **value-bearing meaning cannot remain pluralistic**. BRC-20's "indexer consensus" problem [17] is the lesson the

essay anticipated: when interpretation determines ownership, the network must converge on one canonical interpreter, which quietly re-imports a coordination point. Bismuth's current direction reflects exactly this. The original signature-derived transaction identifier [2] is being replaced, at the project's next hard fork, by a content-addressed identifier — a hash of the transaction's canonical contents — of the kind standard elsewhere; and a deterministic on-chain virtual machine promotes a single canonical interpreter back into consensus for the cases that require authoritative state. The arc runs from "interpretation above consensus" toward "one canonical interpreter, folded into consensus where value demands it" — precisely the collapse the 2018 essay prescribed.

11. Conclusion

Across its first decade Bismuth repeatedly occupied design space that the broader field would later treat as novel, and it did so with both a working developer ecosystem and a peer-reviewed foundation. The value of this record is not a claim of priority for its own sake; it is evidence that durable architectural ideas — separating data from meaning, modular interpretation, asset-level privacy, physical-world data, explicit governance — can appear early, in unfashionable places, be formally analysed, and be validated by the market only in hindsight. The project's continued development, including a hard-fork modernisation of its transaction identity and execution layers, suggests the disposition that produced the 2017–2018 papers remains intact: build and analyse the primitive first, and let the vocabulary catch up.

References

Primary — Bismuth

1. Bismuth Foundation. Mainnet Overview (launch 1 May 2017; no ICO/premine; consensus; HyperNodes). bismuthcoin.org/docs/mainnet/overview/
2. Bismuth Foundation. Bismuth Whitepaper. bismuthcoin.org/docs/papers/whitepaper/
3. G. Hovland and J. Kučera. "Nonlinear Feedback Control and Stability Analysis of a Proof-of-Work Blockchain." Modeling, Identification and Control, 38(4):157–168, 2017. DOI: 10.4173/mic.2017.4.1.
4. J. Kučera and G. Hovland. "Tail Removal Block Validation: Implementation and Analysis." Modeling, Identification and Control, 39(3):151–156, 2018. DOI: 10.4173/mic.2018.3.1.
5. G. Hovland and J. Kučera. Security of the Bismuth Blockchain. 5 May 2018. bismuthcoin.org/pdf/bis-security-20180505.pdf
6. Bismuth Foundation. Hack-with-BIS — concepts and application-protocol corpus (token, alias, decentralized social, Bisnet-DNS, event sourcing, off-chain storage, state channels, Dragginator, dApp-vs-Solidity). github.com/bismuthfoundation/Hack-with-BIS
7. C. Blackbeard, "Bismuth — the first Python Blockchain" (16 Jan 2019); Kingscrown, "Bismuth — No ICO, No Premine and First Python Coin" (5 Oct 2017).
8. J. Kučera. Semantic Interpretation. 19 July 2018. bismuthcoin.org/blog/2018-07-19-semantic/ ; bismuthcoin.org/pdf/2018-07-19-semantic.pdf

Secondary — academic engagement with Bismuth / related

1. A. Leva, S. Strada, M. Tanelli. “Control-oriented modelling of proof-of-work blockchains.” IEEE, doc. 8795749.
2. G. Bissias, D. Thibodeau, B. N. Levine. “Bonded Mining: Difficulty Adjustment by Miner Commitment.” Springer, 2019.
3. S. Liaskos, B. Wang, N. Alimohammadi. “Blockchain Networks as Adaptive Systems.” SEAMS 2019. DOI: 10.1109/SEAMS.2019.00025.
4. P. Pinheiro, R. Santos, R. Barbosa. “Industry 4.0 Multi-agent System Based Knowledge Representation Through Blockchain.” Springer, 2018. DOI: 10.1007/978-3-030-01746-0_39.
5. Hwang et al. Optimizing Performance of a Blockchain. IBM, US Patent 10,880,073.

Secondary — mainstream-equivalent timeline anchors

1. (Semantic essay PDF; see [8].)
2. M. Al-Bassam, LazyLedger (2019); Celestia mainnet, 31 Oct 2023. theblock.co/post/260167/modular-network-celestia-goes-live-on-mainnet
3. V. Buterin. A rollup-centric ethereum roadmap. Ethereum Magicians, Oct 2020. ethereum-magicians.org/t/a-rollup-centric-ethereum-roadmap/4698
4. C. Rodarmor, Ordinals (Jan 2023); Domo, BRC-20 (Mar 2023); “indexer consensus.” certik.com/resources/blog/ordinals-and-the-brc-20-standard-overview-and-risk-analysis ; chain.link/education-hub/brc-20-token
5. Bismuth Foundation. Exploring the Bismuth Plugin System. 16 July 2018. bismuthcoin.org/blog/2018-07-16-plugin-system/
6. The Graph — decentralized indexing protocol; mainnet 2020. thegraph.com
7. Dragginator — collectibles on the Bismuth blockchain. dragginator.com; coverage cryptooico.com (developer interview, July 2018).
8. W. Entriken et al. EIP-721: Non-Fungible Token Standard (Jan 2018); CryptoKitties (late 2017).
9. J. Kučera. Shielded Tokens. 20 Sep 2019. bismuthcoin.org/blog/2019-09-20-shielded-tokens/ ; bismuthcoin.org/pdf/2019-09-20-shielded.pdf
10. Bismuth Foundation. Governance Shift, Parts 1–2 (Aug 2019); BGV-01 (10 Sep 2019); Outcome of the First Governance Vote (6 Nov 2019). bismuthcoin.org/blog/
11. Bismuth Foundation. Decentralized Condition Monitoring (30 Apr 2020); Android Battery Monitoring (9 May 2020); Tesla Battery Monitoring (25 May 2020). bismuthcoin.org/blog/
12. Decentralized Physical Infrastructure Network (DePIN) — term coined by Messari, late 2022; sector overview. en.wikipedia.org/wiki/Decentralized_physical_infrastructure_network
13. Bismuth Foundation. BIS↔ETH Bridge Crystal (21 Apr 2021); BIS↔BSC Bridge Crystal (12 May 2021). bismuthcoin.org/blog/
14. Bismuth Foundation. Python and State Machines. 24 March 2021. bismuthcoin.org/blog/2021-03-24-state-machines/
15. Bismuth Foundation. Beyond DeFi with Bismuth. 12 August 2020. bismuthcoin.org/blog/2020-08-12-defi/

Dates verified against the cited sources. Primary sources are Bismuth’s whitepaper, peer-reviewed papers (Modeling, Identification and Control), dated blog entries, and

the Hack-with-BIS protocol corpus; secondary sources mark each idea’s entry into mainstream practice or its academic treatment.