QUANTUM BLOCKCHAIN TECHNOLOGIES PLC

ASIC Ultra Boost

A More Efficient Approach to SHA-256 Hashing for Bitcoin Mining

White Paper

Based on US Patent Application 18/696,073, “Message Scheduling for Cryptographic Hashing”

Notice of Allowance received from the USPTO, 23 July 2026

This paper sets out, for a technically literate but non-specialist audience, the problem this innovation solves, how it works, how it compares with the established ASICBoost technique, and the efficiency gains it achieves.

1. Executive Summary

Quantum Blockchain Technologies plc ("the Company") has developed and patented a hardware-level optimisation to the SHA-256 hashing process used in Bitcoin mining, referred to internally as **ASIC Ultra Boost**. The innovation is the subject of US Patent Application 18/696,073, *"Message Scheduling for Cryptographic Hashing,"* for which the United States Patent and Trademark Office ("USPTO") issued a Notice of Allowance on 23 July 2026.

The innovation targets a specific and previously unaddressed source of redundant computation inside every Bitcoin-mining ASIC: the repeated, unnecessary recalculation of arithmetic that does not, in fact, change between successive hashing attempts. By identifying which parts of this calculation can be computed once and safely reused, the invention reduces the number of arithmetic and logic-gate operations required per hash attempt by **approximately 17% to 24%**, depending on the processing stage, without altering the final hash output in any way.

This is a hardware efficiency gain, not a shortcut around Bitcoin's security guarantees: every hash produced remains fully compliant with the SHA-256 standard. Because dynamic power consumption in digital ASICs scales directly with the number of gate-level operations performed, this reduction translates into a proportionate opportunity to lower energy consumption per hash, increase hash-rate within a fixed power budget, or both — a meaningful consideration at the scale of an always-on mining fleet.

The innovation is complementary to ASICBoost, an existing, widely referenced optimisation in the mining industry, and can in principle be deployed alongside it for compounding benefit. This paper explains, in non-specialist technical terms, how SHA-256 hashing works, where the inefficiency the Company has addressed arises, how the invention resolves it, and the quantified savings achieved.

2. Introduction

Bitcoin mining hardware exists for a single purpose: to compute the SHA-256 hash function as many times per second as possible, within the constraints of chip area, cost, and power supply. Because mining ASICs run continuously and at enormous scale — a single chip performs billions of hash computations every second, and mining facilities operate thousands of chips around the clock — the efficiency of the underlying hashing circuitry has a direct and compounding effect on operating economics and energy consumption.

Meaningful improvements to mining hardware efficiency are therefore commercially significant, and have historically come from removing redundancy in the standard hashing calculation without affecting the correctness of the result. ASICBoost, introduced publicly in 2016, is the best-known example. The innovation described in this paper identifies a further, previously unaddressed category of redundancy and secures patent protection for a hardware method that removes it.

3. Technical Background: SHA-256 and Bitcoin's Proof of Work

SHA-256 is a cryptographic hash function that compresses an input message of any length into a fixed 256-bit output, or *digest*. Every SHA-256 computation, regardless of input, is performed by the same two hardware stages, executed in sequence:

- A message scheduler, which expands a 512-bit input block into 64 intermediate words used by the next stage.
- A compression engine, which mixes those 64 words over 64 rounds of arithmetic, using eight working registers, to produce the 256-bit digest.

Bitcoin's proof-of-work protocol requires miners to repeatedly hash an 80-byte **Block Header** describing a candidate block of transactions, searching for a header whose resulting hash is numerically below a network-defined target. Miners search for a valid header primarily by varying a 32-bit field called the **Nonce**, trying billions of candidate values per second per chip. For additional security, Bitcoin applies SHA-256 twice to every candidate header — a process commonly termed “double SHA-256.”

4. Inside a SHA-256 Core

Figure 1 illustrates the two hardware stages described above and the flow of data between them.

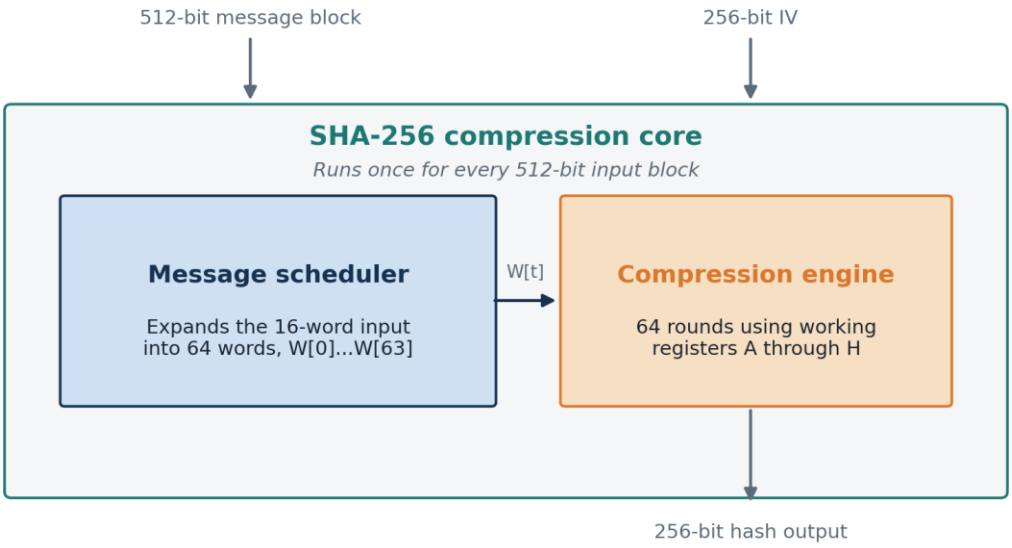

Figure 1. Block diagram of a single SHA-256 compression core.

The message scheduler expands the 16 input words, M[0] through M[15], into 64 output words, W[0] through W[63], using the following recurrence for word indices 16 through 63:

$$W[t] = \sigma 1(W[t-2]) + W[t-7] + \sigma 0(W[t-15]) + W[t-16]$$

where σ1 and σ0 each combine bitwise rotate-right (ROTR), shift-right (SHR), and exclusive-OR (XOR) operations. This recurrence executes 48 times per 512-bit block, contributing a substantial share of the total arithmetic performed on every hash attempt. The Company's analysis of this specific stage — the message scheduler — is the foundation of the innovation described in this paper.

5. Why Bitcoin's Double Hash Requires Three Processing Stages

"Double SHA-256" might be expected to require exactly two executions of the core shown in Figure 1. In practice it requires **three**, because the 80-byte Block Header exceeds the 512-bit (64-byte) block size that a single SHA-256 execution can process.

Figure 2. The three-stage pipeline used to compute Bitcoin's double SHA-256.

Stage	Header fields processed	Commentary
SHA-256₀	Version, HashPrevBlock, first part of HashMerkleRoot	Version is the only field here that changes between attempts, and then only occasionally.
SHA-256₁	Remainder of HashMerkleRoot, Timestamp, Target, Nonce	The Nonce changes on almost every attempt — this is the field the mining search is primarily conducted over.
SHA-256₂	Output digest of SHA-256₁, plus padding	The second hash of "double SHA-256," constant in size regardless of the header content.

6. The Innovation: ASIC Ultra Boost

Standard SHA-256 hardware treats every hashing attempt identically: it re-executes the message-scheduling recurrence from Section 4 in full, on every attempt, with no mechanism for recognising that most of the underlying input data has not changed since the previous attempt. Between one Nonce value and the next, only a small fraction of the 512-bit input actually differs — yet the entire 64-word expansion is recomputed regardless.

The Company's invention addresses this by classifying every input word as either **invariant** (identical across attempts) or **variable** (changing on every attempt), and identifying a mathematical property of the scheduling recurrence that had not previously been exploited: even certain calculations that ultimately feed into a *variable* output word can be partially completed using only *invariant* input data.

Concretely, where some of the terms in the recurrence $W[t] = \sigma 1(W[t-2]) + W[t-7] + \sigma 0(W[t-15]) + W[t-16]$ are invariant across attempts, that portion of the sum can be calculated once, stored as a **reusable partially-calculated value**, and reused on every subsequent attempt. Only the genuinely variable portion of the calculation is performed fresh each time. The patent specification catalogues eight distinct "computation variants," each corresponding to a different combination of invariant and variable terms, together with the precise operation savings each one delivers.

7. Comparative Analysis: ASICBoost and ASIC Ultra Boost

ASICBoost, referenced in the patent's own background discussion, is a well-established prior technique in Bitcoin mining hardware. It is useful to be precise about how it differs from the Company's invention, as the two are frequently, and incorrectly, treated as interchangeable.

ASICBoost

ASICBoost avoids re-executing an entire hash stage by reusing a previous **whole intermediate digest** (H0, the output of SHA-256_0) whenever the same underlying input recurs. This is achieved either by identifying multiple Merkle roots that happen to collide in their last 32 bits ("Covert" ASICBoost), or by deliberately incrementing the Version field so the same H0 can be reused across several attempts ("Overt" ASICBoost).

Critically, ASICBoost provides no benefit whatsoever inside SHA-256_1 or SHA-256_2 — the stages where the Nonce search itself is conducted — and offers no optimisation for SHA-256_0 on the occasions it must still run in full.

ASIC Ultra Boost

The Company's invention operates at a structurally different level: it reduces the **cost of the arithmetic inside every hashing stage**, rather than attempting to avoid re-running stages altogether. Its savings therefore apply on every single attempt, including the Nonce increments inside SHA-256_1 where ASICBoost offers no assistance.

The two techniques are complementary rather than competing: ASICBoost reduces *how often* a hashing stage must be executed in full; ASIC Ultra Boost reduces *how expensive each execution is*. A mining ASIC could, in principle, implement both simultaneously for compounding efficiency gains.

8. Quantified Efficiency Gains

The patent specification quantifies, for each of the three hashing stages, the reduction in Add, ROTR, SHR, and XOR operations achieved once the reusable partially-calculated values have been derived from an initial input and reused for subsequent attempts. These figures are summarised in Figure 3 and Table 2 below.

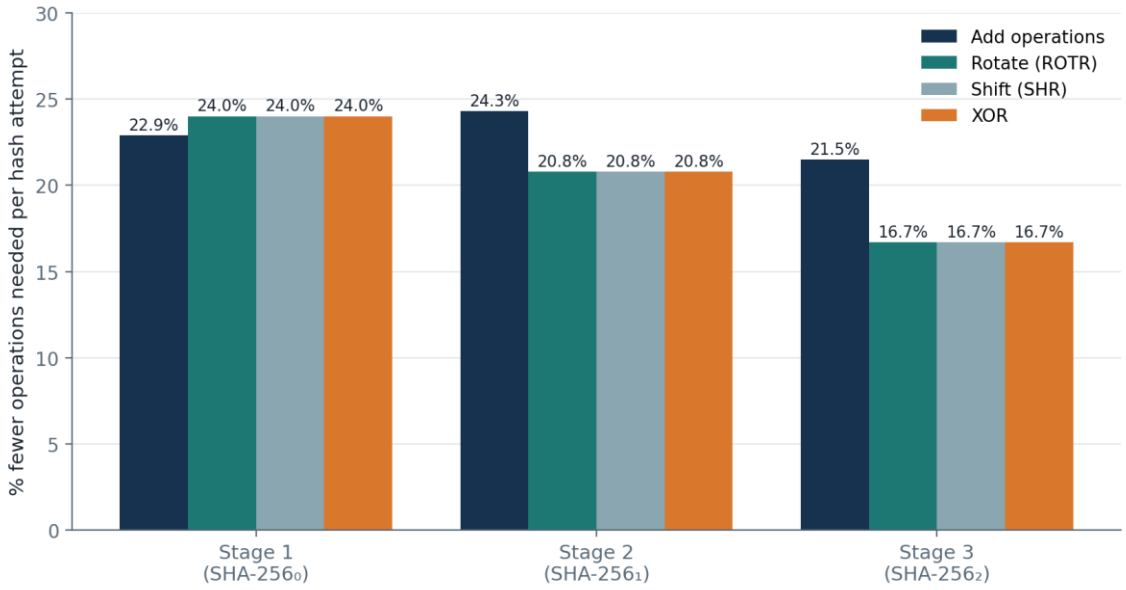

Figure 3. Percentage reduction in operations required per subsequent hash attempt, by stage and operation type.

Stage	Add	ROTR	SHR	XOR
SHA-256₀	22.9%	24.0%	24.0%	24.0%
SHA-256₁	24.3%	20.8%	20.8%	20.8%
SHA-256₂	21.5%	16.7%	16.7%	16.7%

From operations to energy

Each Add, ROTR, SHR, and XOR operation corresponds to a discrete block of digital logic — adder circuitry, fixed-wiring shifts and rotations, and banks of XOR gates — that switches state every time it executes. This switching activity is the principal source of **dynamic power consumption** in digital ASICs performing repetitive arithmetic such as hashing. A reduction of approximately one-fifth to one-quarter in the number of these operations, sustained across every attempt after the first, corresponds to a proportionate reduction in switching activity, and therefore energy consumption, within the message-scheduling portion of the chip.

This freed capacity can be applied either to reduce energy consumption at a fixed hash rate, or to increase hash rate within a fixed power budget — the underlying resource made available is the same either way. The specification further notes a secondary benefit: reusable values remove certain sequential dependencies between words, allowing parts of the schedule for a subsequent input to be computed in parallel in a fully unrolled ASIC implementation.

9. Intellectual Property Position

The invention is the subject of US Patent Application 18/696,073, *"Message Scheduling for Cryptographic Hashing,"* filed 27 March 2024 as the US national stage of PCT/GB2022/052458, claiming priority from UK application 2113962.1, filed 29 September 2021, and assigned to Quantum Blockchain Technologies plc.

The USPTO issued a Notice of Allowance on 23 July 2026, confirming that claims 1–14 and 18–20 have been found allowable following examination. In allowing the application, the examiner concluded that the cited prior art — including an earlier academic thesis by the named inventor and a related prior patent application — did not, singly or in combination, teach or suggest a hashing module whose message scheduler derives reusable partially-calculated values from invariant input portions and reuses them, alongside newly calculated variable-portion values, across subsequent hashing attempts.

A Notice of Allowance confirms that prosecution on the merits is closed but is not, of itself, a grant of patent rights. Formal issuance remains subject to payment of the applicable issue fee and completion of the USPTO's standard post-allowance process. The Company will provide further updates on formal grant in due course, in accordance with its ongoing disclosure obligations.

10. Conclusion and Outlook

ASIC Ultra Boost represents a focused, hardware-level efficiency gain targeted at one of the most computationally intensive components of Bitcoin-mining ASICs: the SHA-256 message scheduler. By recognising that portions of the calculation feeding into variable output words can be precomputed from invariant input data alone, the Company has secured patent protection for an approach that reduces the arithmetic and logic-gate operations required per hash attempt by approximately 17% to 24%, depending on the processing stage — without any change to the correctness or security of the resulting hash.

Because the innovation is complementary to established techniques such as ASICBoost, it represents an additive efficiency opportunity for mining hardware designers rather than a replacement for existing approaches. The Company will provide further updates as the patent proceeds to formal grant and as commercial and licensing opportunities are assessed.

Forward-looking statements

This document contains forward-looking statements, including statements regarding the anticipated efficiency, energy, and performance benefits of the technology described. Such statements are based on the Company's current expectations and the technical specification of the patent application referenced herein, and are subject to inherent risks and uncertainties, including the outcome of ongoing patent prosecution, variation in real-world implementation and hardware design choices, and broader market and regulatory conditions affecting Bitcoin mining. Actual results may differ materially from those anticipated. This document is provided for general informational purposes only and does not constitute investment advice or a recommendation to buy or sell any security.