BERNSTEIN
SOCIETE GENERALE GROUP

21 September 2026

China Semiconductors

China Semis: Huawei's first LogicFolding chip Kirin 9050 Pro reaching 3nm-class performance


Qingyuan Lin, Ph.D.
+852 2123 2654
qingyuan.lin@bernsteinsg.com


Kai Zhang
+852 2123 2665
kai.zhang@bernsteinsg.com


Francis Ma
+852 2123 2626
francis.ma@bernsteinsg.com

Huawei (private) launched their 2nd gen tri-fold smartphone Mate XT2 on September 7th, 2026, paired with the Kirin 9050 Pro chip, the first commercially shipped LogicFolding chip. **We continue to highlight LogicFolding is the DeepSeek moment for China semis and is an underestimated breakthrough, benefiting the local supply chain on advanced logic/packaging, especially SMIC, NAURA, Piotech (all OP).**

A breakthrough that was thought to be impossible. If anyone had told us a 7nm process can be used to make chips with performance equivalent to 3nm-class, we would have dismissed the claim as utterly absurd—if not an outright joke. It's widely perceived that 7nm represented a hard ceiling for DUV only fabrication. And yet, against every established precedent and technical prediction, Huawei did exactly that. This latest flagship mobile SoC Kirin 9050 Pro exceed the performance of Apple A17 Pro, a 3nm-class SoC released in 2023, narrowing the gap between Huawei and Apple chips to ~3 years (was ~4 years for the previous gen). This is no marketing trick or benchmark manipulation or magic: it's the result of LogicFolding, a groundbreaking innovation that is drastically underestimated. What's even more staggering is that it was accomplished not by the world's leading foundries with unlimited access to every advanced tool, but by a local foundry working under severe equipment constraints.

Deep hardware/software co-optimization across single core, multicore and OS. The chip is deployed on Huawei's new tri-fold smartphone Mate XT2, so we chose the third-party test (from Geekerwan) that compares it with Samsung's tri-fold phone as a close proxy. Using the Geekbench multicore score as a benchmark, Kirin 9050 Pro beats Apple A17 Pro (N3) or Snapdragon 8 Gen3 (N4P). Single-core remains further behind, Kirin 9050 Pro is only slightly better than Snapdragon 8 Gen2 (N4) or Apple A14 (N5). This indicates Huawei can optimize multicore design better than peers so extra cores and threads are able work together seamlessly to deliver better performance than just single core. However, the actual performance on the device was even better than shown by the multicore score on native HarmonyOS, Mate XT 2 matches Samsung Galaxy Z TriFold (Snapdragon 8 Elite, 2024, N3E) on FPS/power to run the latest games that requires high SoC performance (e.g. Genshin Impact). The results can be attributed to the hardware/software co-optimization between the chip and HarmonyOS, a unique advantage for Huawei.

LogicFolding is the semis innovation behind the achievement. According to Huawei's lab tests, Kirin 9050 Pro uses 20-40% less power to deliver the same performance as Kirin 9030 Pro for CPU P-Core, or 12-18% better performance at the peak frequency (based on different benchmark). LogicFolding achieve this through reducing the critical path for communication between transistors after stacking two dies o
Performance improvement on GPU and NPU is even bigger (60%+ les
performance) as there's more room for design optimization. Huawei al
multi-year journey for further performance improvements while remai
manufacturing. Manufacturing costs manufacturing and available man
are the only remaining uncertainties for wide adoptions of the technol
not stopping Huawei to resume the #1 China mobile phone brand pos



See the Disclosure Appendix of this report for required disclosures, analyst certifications and other important information. Alternatively, visit our Global Research Disclosure Website.

First Published: 20 Sep 2026 23:05 UTC Completion Date: 20 Sep 2026 23:05 UTC

BERNSTEIN TICKER TABLE

Ticker	Rating	Cur	18 Sep 2026 Closing Price	Price Target	TTM Rel. Perf.	Cur	Reported EPS 2025A	Reported EPS 2026E	Reported EPS 2027E	Reported P/E (x) 2025A	Reported P/E (x) 2026E	Reported P/E (x) 2027E
688012.CH (AMEC)	O	CNY	354.92	660.00	84.4%	CNY	2.28	5.30	7.26	155.5	67.0	48.9
688256.CH (Cambricon)	O	CNY	1,113.14	1,340.00	(0.8)%	CNY	3.28	7.50	14.52	339.9	148.5	76.7
1347.HK (Hua Hong-H)	O	HKD	115.50	220.00	67.8%	USD	0.03	0.14	0.20	N/M	718.8	481.1
688347.CH (Hua Hong-A)	O	CNY	234.00	370.00	158.2%	USD	0.03	0.14	0.20	N/M	N/M	N/M
688041.CH (Hygon)	O	CNY	241.38	450.00	(21.5)%	CNY	1.10	2.04	3.59	219.4	118.1	67.2
002371.CH (NAURA)	O	CNY	686.51	1,380.00	47.0%	CNY	5.66	10.11	19.92	121.2	67.9	34.5
688072.CH (Piotech)	O	CNY	720.06	1,200.00	218.9%	CNY	3.32	10.18	16.61	216.9	70.7	43.3
981.HK (SMIC-H)	O	HKD	65.15	120.00	(30.4)%	USD	0.09	0.13	0.17	95.4	65.3	48.3
688981.CH (SMIC-A)	O	CNY	122.00	220.00	(23.2)%	USD	0.09	0.13	0.17	209.3	143.3	105.9
ASIAX			1,953.49									

O - Outperform, M - Market-Perform, U - Underperform, NR - Not Rated, CS - Coverage Suspended

Source: Bloomberg, Bernstein estimates and analysis.

INVESTMENT IMPLICATIONS

We rate SMIC, NAURA, and Piotech Outperform.

This launch demonstrates significant upside on China's advanced logic and advanced packaging demand. The technology right now is still proprietary to one company, but we believe this would be gradually adopted by other local fabless as well after the ecosystem becomes more matured. Within our coverage, we believe **SMIC** is the foundry behind the LogicFolding chip manufacturing on 7nm process using DUV, which should be the major beneficiary of increasing demand. Compared with a single wafer in a planar Kirin 9030 Pro, each Kirin 9050 Pro doubles logic wafer consumption. **NAURA** is the major vendor that supplies the etch and deposition equipments for SMIC's 7nm process. **Piotech** is the domestic leader on W2W bonding tools for advanced packaging, which could also benefit from the further adoption of the technology.

DETAILS

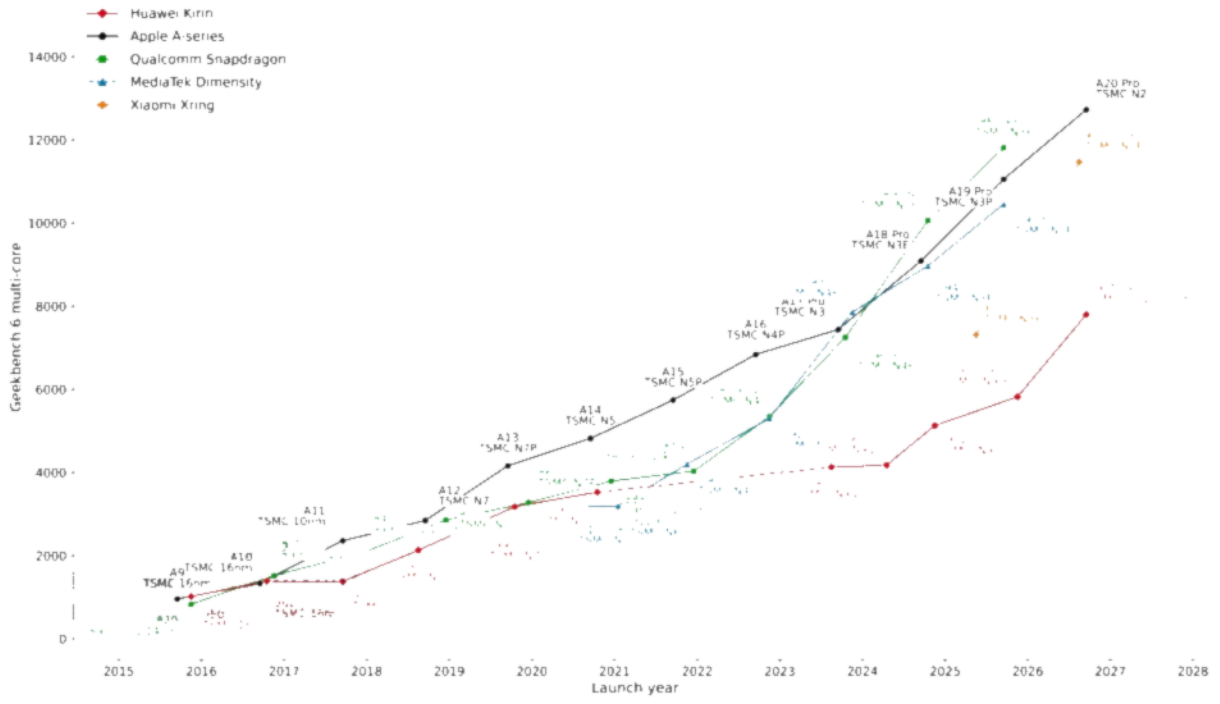
A 7NM-CLASS CHIP REACHING 3NM-CLASS PERFORMANCE

ACHIEVING A BREAKTHROUGH THAT WAS THOUGHT TO BE IMPOSSIBLE

If anyone had told us just a few years ago that a company could take a 7nm-class manufacturing process and push it to deliver performance equivalent to a 3nm-class chip, every industry expert, analyst, and engineer would have dismissed the claim as utterly absurd—if not an outright joke. At the time, the entire semiconductor consensus held that advancing beyond 7nm was physically impossible without extreme ultraviolet (EUV) lithography, the extraordinarily complex and tightly controlled technology that had become the exclusive gateway to denser transistor geometries and higher performance. Conventional wisdom dictated that 7nm represented a hard ceiling for DUV-only fabrication; any attempt to squeeze 3nm-level performance out of it would violate everything known about transistor scaling, power efficiency, and thermal limits.

And yet, against every established precedent and technical prediction, Huawei did exactly that. In its latest flagship mobile chip Huawei Kirin 9050 pro, our analysis show that the company has achieved what the broader industry had long written off as unachievable: harnessing a 7nm-class production line entirely without EUV equipment to deliver overall performance on par with leading 3nm-class processors. This is no trick of marketing or benchmark manipulation, and it is certainly not magic. It is the result of LogicFolding, a groundbreaking architectural and manufacturing innovation that reinvents how transistors are arranged and optimized at the circuit level—and one that the global semiconductor community has, in my personal opinion, drastically underestimated in both its technical sophistication and its disruptive potential. What makes the feat all the more staggering is that it was accomplished not by the world's leading foundries with unlimited access to every advanced tool, but by a company working under severe supply constraints on equipments, turning a supposedly insurmountable technical barrier into an opportunity for fundamental reinvention.

EXHIBIT 1: **Same-test Geekbench6 multi-core scores for the leading flagship mobile phone SoCs**



Krin 9050 Pro has no public Geekbench 6 ranking; we scale 9030 Pro's Geekbench 6 score by the same-test Geekbench 7 step versus Apple A17 Pro to get the implied Geekbench 6 score.
Source: TOPCPU, Geekerwan, Bernstein analysis

GEEKBENCH SCORE COMPARISON FOR THE LEADING MOBILE PHONE CHIPS

Let's take a closer, more rigorous look at how this chip actually performs in practice. To make a relatively fair and apples to apples comparison, we use Geekbench 6 test scores as our primary benchmark, running both single-core and multicore workloads to compare the Kirin 9050 Pro against other leading flagship chips on the market (Apple A-series, Qualcomm Snapdragon, MediaTek Dimensity and Xiaomi Xring).

We acknowledge that this benchmark is not perfect, and critics will rightly point out several limitations: there are long-running debates in the tech community about the potential for benchmark score manipulation or targeted optimizations, and synthetic tests like these do not fully capture the software level and system level tuning that can dramatically elevate real world performance beyond what the raw silicon alone would predict (for example the Apple chips always quoted to deliver better user experience at the same score compared to Qualcomm chips). Factors like thermal management, memory scheduling, OS level optimizations, and app specific co-optimization all play a role in daily use that a standardized benchmark cannot fully measure. That said, Geekbench 6 remains one of the most widely adopted, consistently applied benchmarks in the industry for evaluating cross-platform mobile chip performance, and crucially, it applies the exact same test suite and scoring methodology to every chip it runs on — meaning at the very least, the testing conditions are held constant across all devices being compared. Most importantly, for this analysis we have only selected scores submitted from commercially available retail mobile phones running normal consumer firmware, rather than controlled lab test scores collected under unrealistic conditions such as extreme cryogenic cooling, unlimited power delivery, or pre-production silicon with cherry-picked bins. By using only real-world scores from devices as consumers actually use them, we can ensure the comparison aligns as closely as possible with the performance people will actually experience in their daily lives.

On Geekbench6 multicore scores, Huawei Kirin 9050 Pro is already better than Apple A17 Pro, which is the world's first 3nm-class mobile phone chips (Exhibit 1).

The Apple A17 Pro is Apple's flagship mobile system-on-chip, officially unveiled on September 12, 2023, as the core silicon powering the iPhone 15 Pro and iPhone 15 Pro Max. It earned industry-wide attention as the first mass-produced 3nm smartphone chip to reach consumers. It was fabricated by TSMC on the foundry's first-generation N3B 3nm process node — a custom, launch-phase 3nm platform developed for Apple. The chip integrates ~19 billion transistors (a ~19% increase over the preceding A16 Bionic), leveraging the 3nm node's gains in transistor density and power efficiency to push performance within strict mobile thermal and power limits. The A17 Pro uses a 6-core CPU design: 2 high-performance "Everest" cores clocked up to 3.78 GHz, paired with 4 energy-efficient "Sawtooth" cores. Its 6-core GPU was Apple's first mobile graphics processor with hardware-accelerated ray tracing and mesh shading. Apple officially claimed ~10% faster CPU performance, ~20% faster GPU performance, and a 2x faster 16-core Neural Engine (35 TOPS INT8) versus the A16 Bionic. As the first volume product built on TSMC's 3nm node, the A17 Pro defined the baseline for premium 3nm-class mobile performance and efficiency in its generation.

On the other hand, the Huawei Kirin 9050 Pro is still using SMIC 7nm-class process (N+2 or N +3) plus LogicFolding, using DUV rather than EUV. Yes, the chip can only match the performance of the Apple chip that was released 3 years ago, and is still about 30% lower in performance compared to the Apple's latest flagship chip A20 pro that is releases around the same time, but is still very impressive to see the improvement, as the industry generally believe that China will be stuck at 7nm. Indeed, if we look at the performance improvement of Huawei chips since export control (Huawei was banned from using TSMC process since Sept 2020), the first gen Huawei chip manufactured in SMIC back in 2023 was only able to match the performance of Apple's chips made in 2019, with 4 years of gap, and the chips from Huawei in 2024/25 was keeping the gap with Apple at about the same 4-years of gap. Just when the industry believes the gap would further widen in the next few generations, here comes the LogicFolding technology that brings the gap back to about 3 years, and from the roadmap that Huawei disclosed, there is still plenty of room to further improve the chip performance beyond this gen 1 LogicFolding. Importantly, even though this process still relies on tools from global vendors, the fact that Huawei is willing to public talk about the chip in mobile phone release indicate that they don't have any concern that US government can future put restriction on the company to stop it from manufacturing/improving, a significant step further towards self-sufficiency in advanced logic development.

Kirin 9050 Pro still trails on single-core performance. If we look at the scores closely, we have to admit that process limitation is still a bottleneck for Huawei, there is no magic that it can do to overcome the physical constrain at the transistor level. The single core score for Huawei (~2000) is around TSMC N4 level, similar to Qualcomm Snapdragon 8 Gen2 (2,004, TSMC N4), MediaTek Dimensity 9200 (1,949, TSMC N4), and Apple A14 (2,087, TSMC N5). The single-core score reflects more of the intrinsic performance of the transistors, but what the multicore performance is what the consumers can really feel, and that reflects more on system level optimization capabilities for the fabless, so we believe is a more fair comparison. For

multicore score, HarmonyOS also plays an important role, by deciding which threads sit on which cores also lifts the multicore score more than the single-core score. That matches the laboratory paper: LogicFolding plus extra mid-cores help when work can be split; a single thread still walks one pipeline, which is why P-Core turbo was only increased by 18%.

EXHIBIT 2: **Geekbench 6 single / multi-core score by launch year**

Lunch Year	Huawei Kirin	Apple	Qualcomm Snapdragon	MediaTek Dimensity	Xiaomi Xring
2015	**Kirin 950** (TSMC 16nm)	**Apple A9** (TSMC 16nm)	**Snapdragon 820** (Samsung 14nm)	-	-
	327 / 1,016	623 / 957	323 / 827		
2016	**Kirin 960** (TSMC 16nm)	**Apple A10** (TSMC 16nm)	**Snapdragon 835** (Samsung 10nm)	-	-
	408 / 1385	879 / 1,332	415 / 1,519		
2017	**Kirin 970** (TSMC 10nm)	**Apple A11** (TSMC 10nm)	**Snapdragon 845** (Samsung 10nm)	-	-
	386 / 1,377	1,097 / 2,357	566 / 2,075		
2018	**Kirin 980** (TSMC N7)	**Apple A12** (TSMC N7)	**Snapdragon 855** (TSMC N7)	-	-
	834 / 2,128	1,301 / 2,843	939 / 2,855		
2019	**Kirin 990 5G** (TSMC N7)	**Apple A13** (TSMC N7P)	**Snapdragon 865** (TSMC N7P)	-	-
	969 / 3,176	1,748 / 4,161	1,128 / 3,277		
2020	**Kirin 9000** (TSMC N5)	**Apple A14** (TSMC N5)	**Snapdragon 888** (Samsung 5nm)	**Dimensity 1000+** (TSMC N7)	-
	1,266 / 3,529	2,087 / 4,820	1,481 / 3,794	1,040 / 3,152	
2021	-	**Apple A15** (TSMC N5P)	**Snapdragon 8 Gen1** (Samsung 4nm)	**Dimensity 1200** (TSMC N6)	-
				1,118 / 3,198	
		2,333 / 5,741	1,664 / 4,030	**Dimensity 9000** (TSMC N4)	
				1,599 / 4,199	
2022	-	**Apple A16** (TSMC N4P)	**Snapdragon 8 Gen2** (TSMC N4)	**Dimensity 9200** (TSMC N4)	-
		2,627 / 6,838	2,004 / 5,344	1,949 / 5,301	
2023	**Kirin 9000S** (SMIC N+2)	**Apple A17 Pro** (TSMC N3)	**Snapdragon 8 Gen3** (TSMC N4P)	**Dimensity 9300** (TSMC N4P)	-
	1,334 / 4,128	2,953 / 7,441	2,181 / 7,250	2,225 / 7,857	
2024	**Kirin 9010** (SMIC N+2)	**Apple A18 Pro** (TSMC N3E)	**Snapdragon 8 Elite** (TSMC N3E)	**Dimensity 9400** (TSMC N3E)	-
	1,410 / 4,176				
	Kirin 9020 (SMIC N+2)	3,582 / 9,089	3,234 / 10,059	2,874 / 8,969	
	1,603 / 5,129				
2025	**Kirin 9030 Pro** (SMIC N+3)	**Apple A19 Pro** (TSMC N3P)	**Snapdragon 8 Elite Gen5** (TSMC N3P)	**Dimensity 9500** (TSMC N3P)	**Xring O1** (TSMC N3E)
	1,683 / 5,817	4,019 / 11,054	3,718 / 11,810	3,518 / 10,459	2,362 / 7,313
2026	**Kirin 9050 Pro** (SMIC LogicFolding)	**Apple A20 Pro** (TSMC N2)	-	-	**Xring O3** (TSMC N3P)
	1,965 / 7,799	4,732 / 12,727			2,977 / 11,472

Krin 9050 Pro has no public Geekbench 6 ranking; we scale 9030 Pro's Geekbench 6 score by the same-test Geekbench 7 step versus Apple A17 Pro to get the implied Geekbench 6 score.
Source: TOPCPU, Geekerwan, Bernstein analysis

KIRIN 9050 PRO: HUAWEI'S FIRST COMMERCIAL LOGICFOLDING CHIP

THE LAUNCH OF HUAWEI'S GEN2 TRI-FOLD PHONE

Huawei officially unveiled the Mate XT2 at its fall flagship product launch event held in Guangzhou on September 7, 2026 (Exhibit 3), with general sales of the device set to kick off on September 12, 2026 at a starting price of RMB 19,999 (or roughly USD 3k). Positioned at the top end of Huawei's ultra-premium mobile lineup, the Mate XT2 marks the company's second-generation tri-fold smartphone, building on the groundbreaking three-panel form factor introduced with the original Mate XT in 2024 to deliver a more refined, durable, and versatile foldable experience that transitions seamlessly from a compact handset to a near-tablet-sized display (10.2-inch screen after unfolded). The device ships preloaded with HarmonyOS 7, Huawei's latest native operating system, which features a full suite of form-factor-specific optimizations—including adaptive UI scaling across all three display states, streamlined multi-window multitasking, and seamless app continuity when folding or unfolding the screen—to fully leverage the tri-fold design's unique versatility.

Most notably, the Mate XT2 is the first commercial smartphone to launch with Huawei's new Kirin 9050 Pro system-on-chip. During the launch keynote, Huawei executives framed the new silicon as the most powerful Kirin processor ever created, and highlighted that it represents the first all-new high-performance Kirin flagship chip to debut at a full-scale Huawei flagship launch event in six years (the last time the company talked about chip performance in mobile phone launch was Sept 2020 on Kirin 9000). That six-year milestone carries particular industry significance, as it marks the company's return to unveiling top-tier in-house silicon at a major consumer launch following a multi-year hiatus triggered by U.S. export controls that cut off Huawei's access to leading-edge global foundry production in late 2020.

As the headline performance claim for the device (Exhibit 4), Huawei stated that overall real-world system performance of the Mate XT 2 lands 42% above that of the previous-generation Mate XTs (using Kirin 9020 pro chip). The bulk of this generational leap is attributed to the Kirin 9050 Pro's upgraded CPU and GPU architecture, alongside HarmonyOS 7's targeted efficiency tuning, improved memory management, and optimizations for the heavier graphical and multitasking workloads that come with driving the tri-fold's large, high-resolution multi-display setup.

EXHIBIT 3: **Huawei launched Mate XT 2 with with Kirin 9050 Pro mobile processor**



Source: Company news, Bernstein analysis

EXHIBIT 4: **Overall performance for Mate XT 2 is 42% above that of previous tri-phone phone Mate XTs.**



Source: Company news, Bernstein analysis

OPTIMIZED NEW MULTICORE DESIGN

Kirin 9050 Pro is the first commercial LogicFolding SoC - Huawei's circuit-level implementation of Tau (τ) Law. Compared with Kirin 9030 Pro, the process and CPU micro-architecture are essentially unchanged on 9050 Pro. Therefore, the incremental performance on Kirin 9050 Pro mainly come from the application of LogicFolding. Performance details:

- **Transistor density**: Transistor density raises from 155 to 238 MTr/mm^2 based on their roadmap (Exhibit 13), increasing 55%, and Huawei attributes the density improvement to the stacking.
- **CPU**: The innovated structure includes nine-core 1+2+4+2 with SMT (Exhibit 5). And P-Core frequency peak increases from 2.75GHz to 3.1GHz. Overall SRAM capacity also slightly increases.
- **GPU**: Kirin 9050 Pro integrates Maleoon 955 GPU, adding hardware ray tracing up to 50 million rays/s, and a claimed 142% lift in real-time rendering. The max frequency of GPU core is also improved from 933 MHz to 1056 MHz
- **NPU**: Da Vinci is rated to run a 30-billion-parameter on-device MoE.

EXHIBIT 5: **Comparison of core desing for Kirin 9030 Pro vs. 9050 Pro**



Source: Geekerwan, Bernstein analysis

PERFORMANCE ANALYSIS FROM THIRD-PARTY TESTING

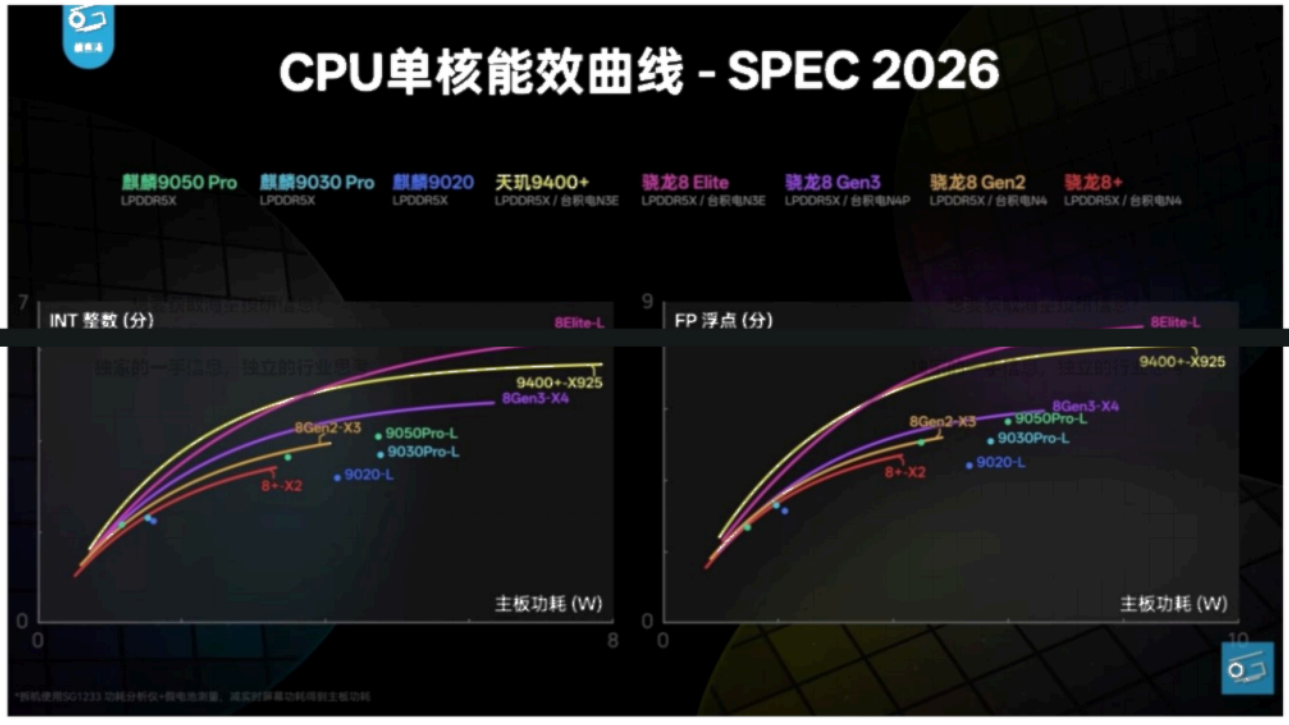
There are various third party tests measures Kirin 9050 Pro performance in a Mate XT2, we use the one from Geekerwan as illustration. Specific test is designed for each block. For CPU, SPEC CPU 2026 measures prime-core work at a given clock, and Geekbench 7 measures single- and multicore peak scores; both ran as HarmonyOS apps. Steel Nomad Light is tested for the GPU, and an INT8 throughput test is conducted for the NPU. Compared with Kirin 9030 Pro, Kirin 9050 Pro does not change much on the design on CPU cores, but both GPU and NPU are larger.

CPU Test

On SPEC CPU 2026, when locked at 2.75 GHz, Kirin 9030 Pro's peak, the Kirin 9050 Pro prime core does the same integer work as the Kirin 9030 pro but uses more than 30% less power (Exhibit 6). That 2.75 GHz match is the clean comparison. Since Kirin 9050 Pro's CPU design and process are essentially unchanged from that of Kirin 9030 Pro, the power cut is achieved from LogicFolding. At its new 3.1 GHz peak, the Kirin 9050 Pro prime then does more work than the Kirin 9030 Pro prime did at 2.75 GHz, without a matching step-up in power. The efficiency chart showed is still sloping at 3.1 GHz: raising the Kirin 9050 Pro clock still buys extra performance rather than mostly extra watts. Therefore, the 3.1 GHz should still be a conservative peak, not the performance limit of the LogicFolding chip.

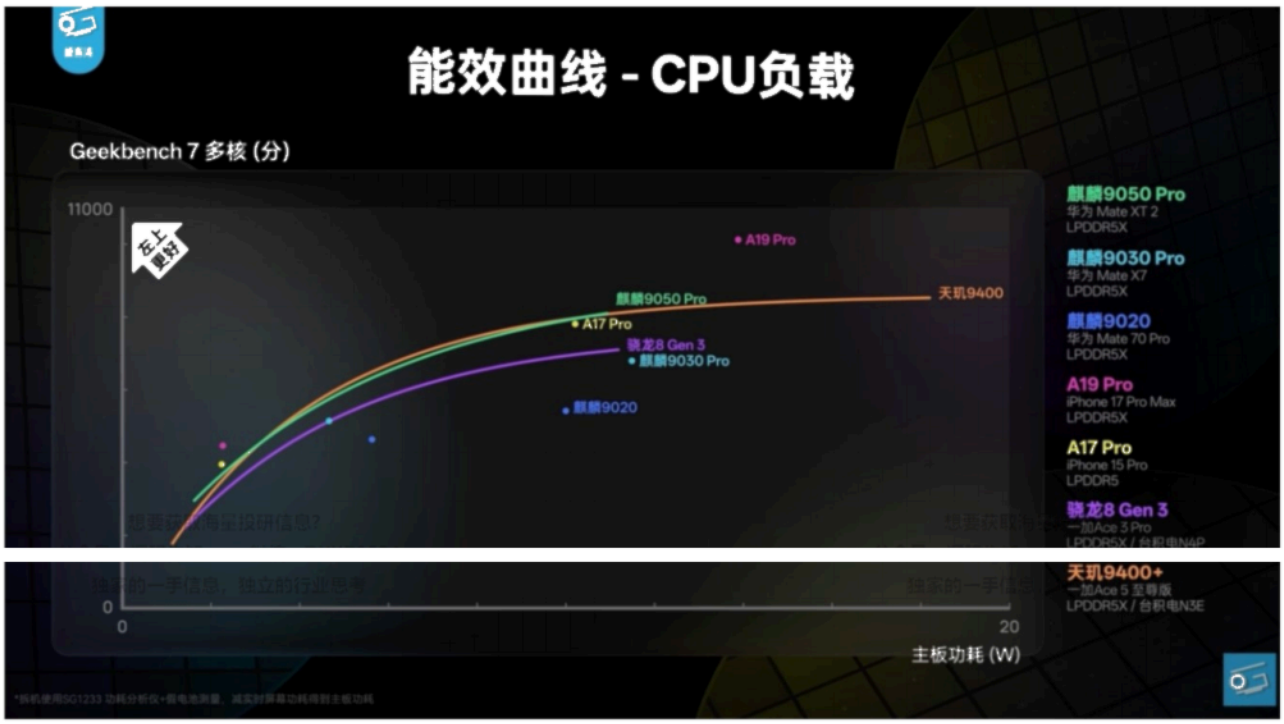
On Geekbench 7 test (Exhibit 7, Exhibit 8), Kirin 9050 Pro scores 1,813 single core and 8,159 multi-core, versus 1,596/6,802 on Kirin 9030 Pro, increasing 14% and 20% separately. The single-core score is from the Kirin 9050 Pro prime core. And the multi-core score runs work across the whole CPU: nine cores and 16 threads (seven larger cores each run two threads). The score is clearly better than Qualcomm Snapdragon 8gen3 (TSMC N4P) and comes in-line with MediaTek Dimensity 9400 (TSMC N3E), also proves that the performance of Kirin 9050pro is equivalent to 3nm-class.

EXHIBIT 6: **SPEC 2026 prime-core at the same clock**



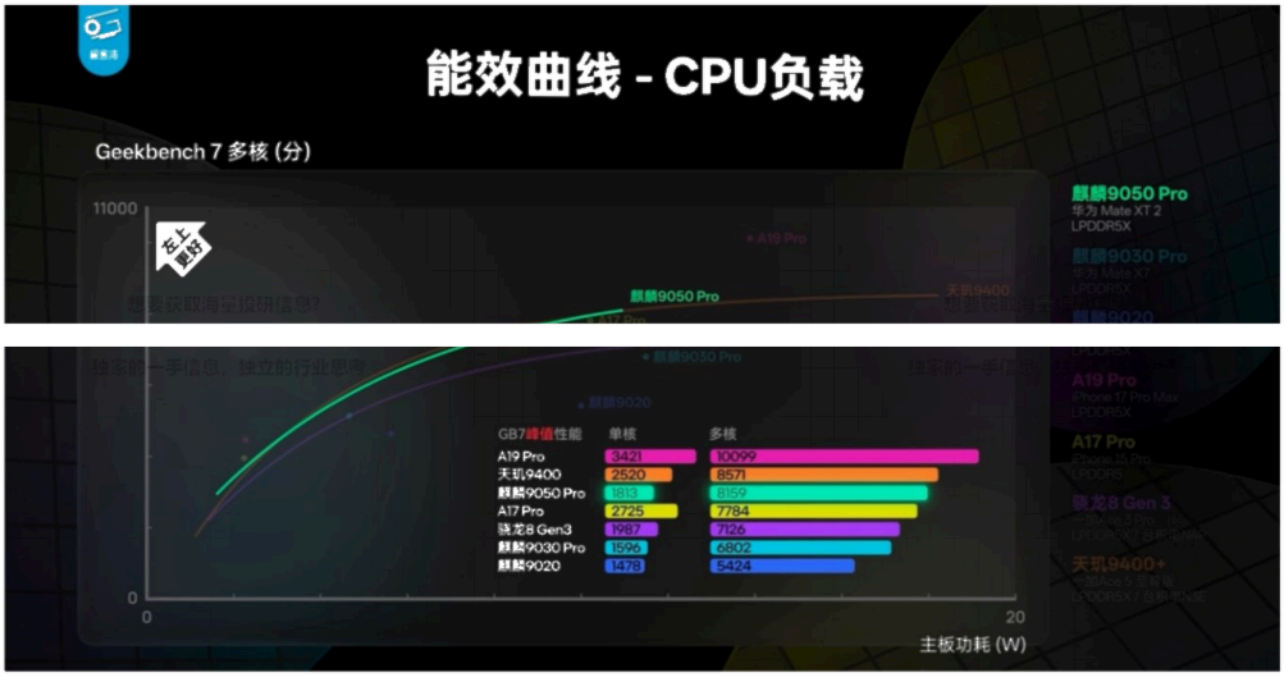
Source: Geekerwan, Bernstein analysis

EXHIBIT 7: **Geekbench 7 multicore efficiency and peak table**



Source: Geekerwan, Bernstein analysis

EXHIBIT 8: **Multicore score of Kirin 9050 Pro is above Apple A17 Pro for GB7 peak performance**



Source: Geekerwan, Bernstein analysis

GPU Test

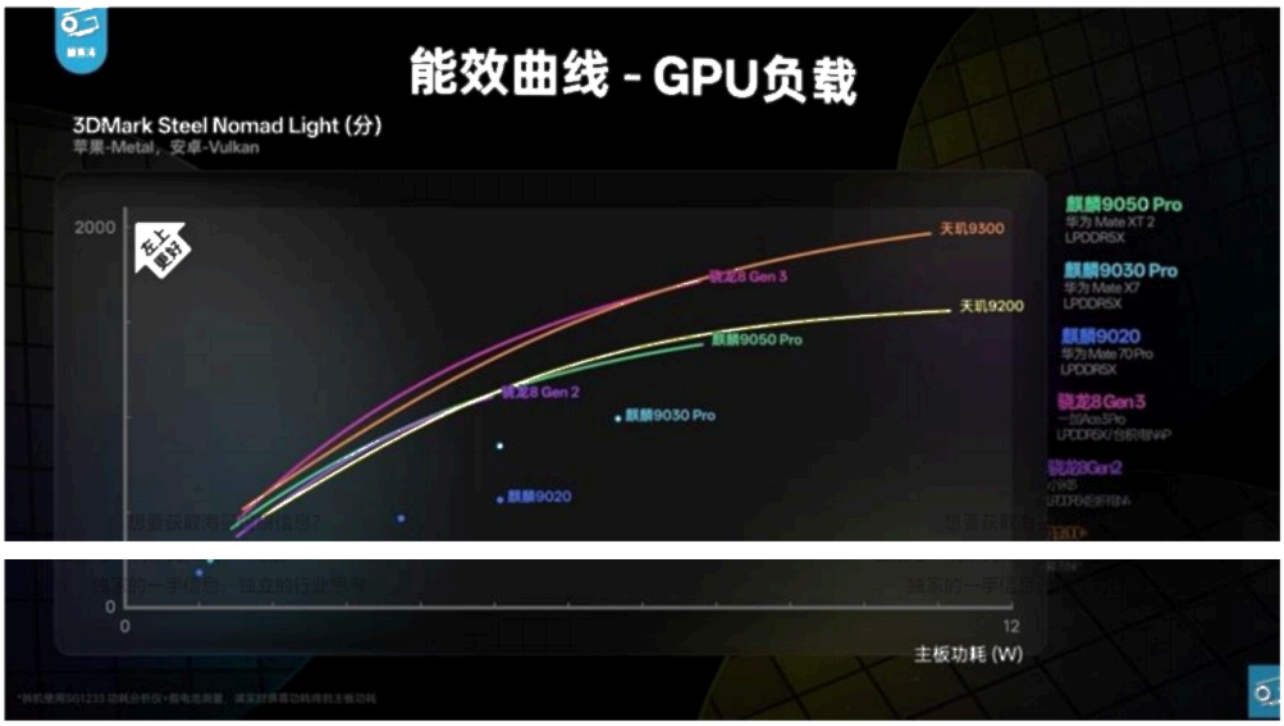
On Steel Nomad Light, a GPU graphics test, Kirin 9050 Pro scores 1,378 versus 993 on Kirin 9030 Pro, increasing 39%. The Kirin 9050 Pro GPU is Maleoon 955. It still has six compute clusters, the same count as the GPU on Kirin 9030 Pro, but each cluster has more arithmetic units and the clock is 13% higher at 1,056 MHz. Peak single-precision throughput therefore rises from 1.4 trillion operations per second on Kirin 9030 Pro to about 2.2 trillion on Kirin 9050 Pro, increasing 57%.

EXHIBIT 9: **Steel Normad Light**



Source: Geekerwan, Bernstein analysis

EXHIBIT 10: **GPU efficiency curve**



Source: Geekerwan, Bernstein analysis

NPU Test

On INT8 neural-network test, the Kirin 9050 Pro NPU delivers 67.7 TFLOPS versus 27 TFLOPS on Kirin 9030 Pro, increasing 150%. This is the largest on-device step versus Kirin 9030 Pro of the three blocks. The NPU in Kirin 9050 Pro is the first in a phone that can run a 30-million-parameter mixture- of-experts model locally, with no cloud round-trip. Mixture-of-experts stores many specialist subnetworks and only turns a few on at a time: Huawei's Pangu 30B-A2B has 30 billion parameters in total, but only about 2 billion are active in any one step, so the phone does not compute all 30 billion at once. HarmonyOS 7's Xiaoyi assistant uses that local model. Examples of Xiaoyi's features in practice include: Xiaoyi can sortthe photo album with the network off, and it can also run on-device vision for accessibility.

EXHIBIT 11: **NPU INT8 throughput**



Source: Geekerwan, Bernstein analysis

REAL PERFORMANCE ON GAME TESTING SHOW EVEN BETTER PERFORMANCE

The Geekerwan tests also run three 30-minute native HarmonyOS games on three trifold phones: Mate XT 2 (with Kirin 9050 Pro chip and HarmonyOS 7), Mate XTs (with Kirin 9020 chip and HarmonyOS 5.1/6), and Samsung Galaxy Z Trifold (Snapdragon 8 Elite chip). The scores are average frames per second, the slowest 1% of frames (stutter), and board power. All three phones are trifolds, so they shed heat in a similar body and, therefore, make the tests results more reliable to reflect the actual chip performance.

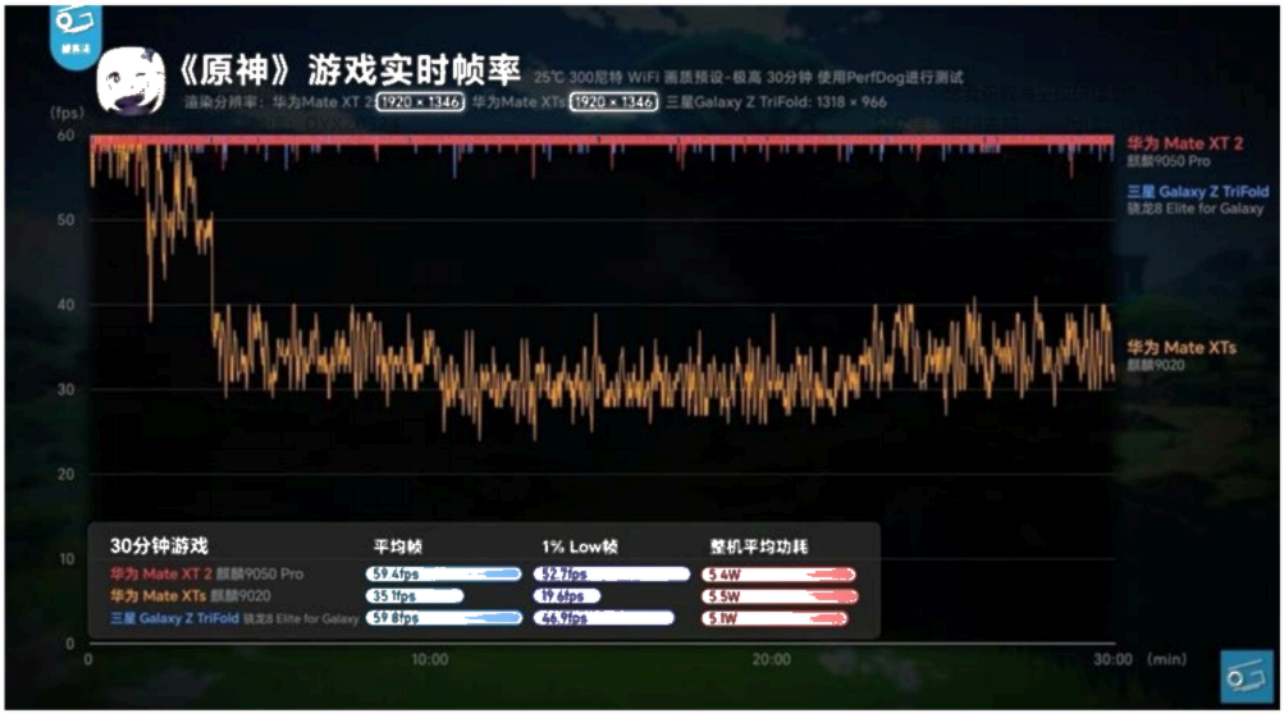
On these three games, Genshin Impact, Ananta and Wuthering Waves, Mate XT 2 matches the Samsung Galaxy Z Trifold on average frames per second. For example, when running Genshin Impact is capped at 60 FPS (Exhibit 12): Kirin 9050 Pro holds 59.4 FPS versus 59.8 FPS on Snapdragon 8 Elite, even though Huawei is rendering at a clearly higher resolution (1920 x 1346 versus 1518 x 966). Mate XT 2 stutters less: the slowest 1% of frames is 52.7 versus 46.9 on Samsung Galaxy Z Trifold. Ananta is the fairer like-for-like: similar resolution, Kirin 9050 Pro holds 39.4 FPS versus 38.9 FPS on Snapdragon 8 Elite. Wuthering Waves is the other way round: while Samsung has higher resolution, Kirin 9050 Pro still matches (38.2 FPS versus 37.7 FPS). Mate XT 2 uses more power than the Samsung Trifold on Ananta and Wuthering Waves (6.4 W versus 6.0 W, and 6.9 W versus 6.2 W). On Genshin, the gap is smaller at the higher resolution (5.4 W versus 5.1 W).

Compared with the results of Mate XTs (with Kirin 9020 chip), Mate XT 2's frames per second rise by about 70-100%. Three upgrades have been achieved together on Mate XT 2: HarmonyOS 7, the new Maleoon 955 GPU, and LogicFolding on Kirin 9050 Pro. We cannot say how much of the game lift comes from each of these upgrades. Furthermore, the tests are conducted only on trifold phones, and a possible future non-foldable flagship with Kirin 9050 Pro is not in the set. A non-

foldable flagship would be thicker and would shed heat more easily, which means Kirin 9050 Pro could hold peak for longer. In other words, Mate XT 2's FPS may therefore look weaker than the same chip would in a non-foldable phone.

The result demonstrate that even with about 3 years of gap in chip performance, the gap from real user experience is already minimal compared to other flagship chips, even when running on performance sensitive high-end game. Beyond the chip performance improvement, the software-hardware co-design is also a very important contribution to the actual user experience.

EXHIBIT 12: **Genshin Impact - Huawei at higher render resolution, still locked 60 FPS, better 1% low**



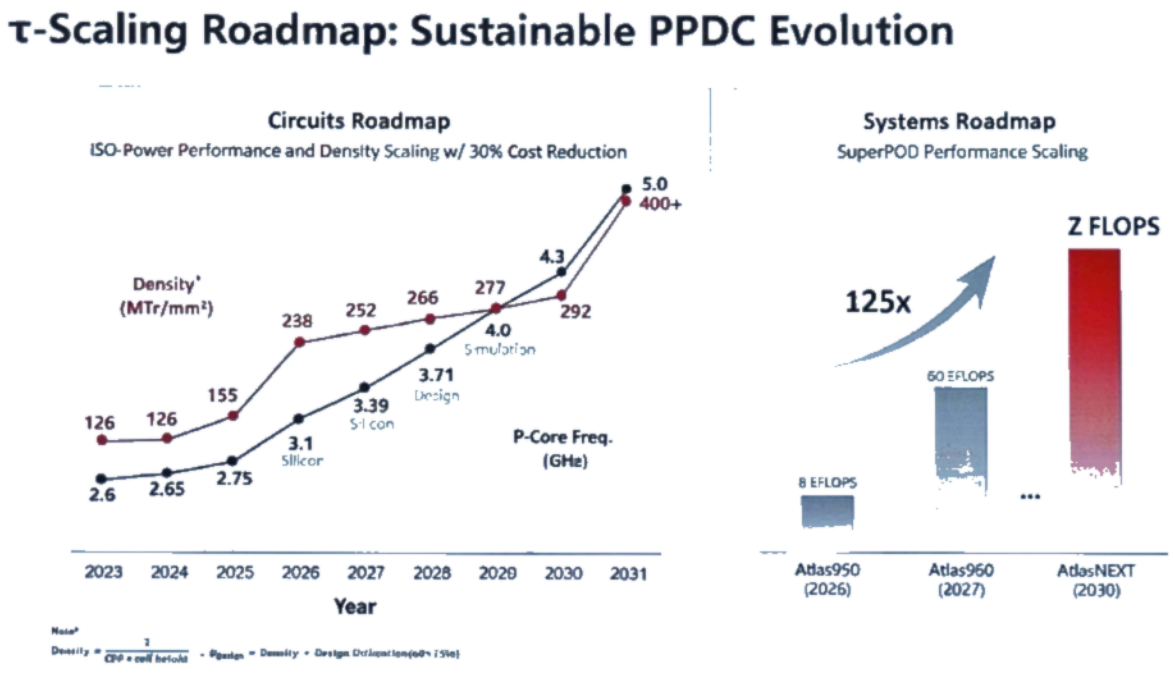
Source: Geekerwan, Bernstein analysis

TECHNOLOGY BEHIND THE SCENES: LOGICFOLDING

Huawei has measured the performance of folded 9050-generation silicon against planar Kirin 9030 Pro and recently presented in the paper: Huawei's τ Chip Was Supposed to Melt? The paper shows the laboratory tests results for the first LogicFolding Kirin (labbled Kirin 2026 in the exhibits, which is Kirin 9050 pro in commercial deployment) on NPU, GPU, CUP P-Core, Geekbench 6 and DSP. Two protocols have been applied: match the 9030 Pro's score and record power; then let the folded chip run faster. More details on the logic folding technology was also illustrated in detail on the previous full technical paper: A Time Scaling Theory for Multi-Layer Electronic Systems.

Holding each block to the same score with Kirin 9030 Pro, the power is down 66% in NPU, 58% in GPU, and 41% in P-core at the folded chip. Geekbench 6 at the same 2.75 GHz clock is a milder 23% power cut on the CPU as a whole. LogicFloding shortens the wires, so the circuit finishes each clock cycle with time to spare. And Huawei takes that extra time to lower voltage, which is why each of these blocks uses less power at the same work. NPU and GPU are parallel: extra transistors become extra copies of the same engine, so the Kirin 9030 Pro's TOPS or FPS can be hit at a much lower clock and voltage on each copy. A single-thread P-Core job cannot be split across copies and there is still one pipeline. So the same-work saving only comes from the shorter wires and a 9% clock cut, the smallest of the hot blocks. More details in Exhibit 13 to Exhibit 19.

EXHIBIT 13: **Huawei's tau scaling roadmap**



Source: Company reports, Bernstein analysis

EXHIBIT 14: **The comparison of power at iso-performance between Kirin 9050 pro and Kirin9030 pro**

	Kirin9030 Pro	Kirin2026
Temperature	25 °C	25 °C
Voltage	1.1 V	0.9 V
Frequency	2.75 GHz	2.5 GHz
Normalized power at iso-performance	1	0.59
Normalized area	1	0.625
Normalized power density	1	0.944

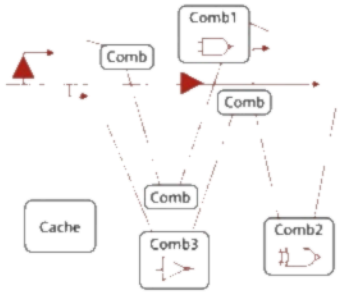
Source: Huawei, Bernstein analysis

EXHIBIT 15: **Trend of the operating frequency of Kirin CPU performance core**

	SoC	Architecture	Frequency	State
2023	Kirin9000s	Planar	2.6 GHz	Mass product
2024	Kirin9020	Planar	2.65 GHz	Mass product
2025	Kirin9030 pro	Planar	2.75 GHz	Mass product
2026	Kirin 2026	LogicFolding	3.1 GHz	Silicon
2027	Kirin 2027	LogicFolding	3.39 GHz	Silicon
2028	Kirin 2028	LogicFolding	3.71 GHz	Pre-silicon
2029	Kirin 2029	LogicFolding	4 GHz	Pre-silicon

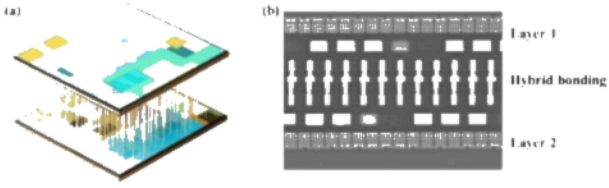
Source: Huawei, Bernstein analysis

EXHIBIT 16: **The schematic illustration of LogicFolding**



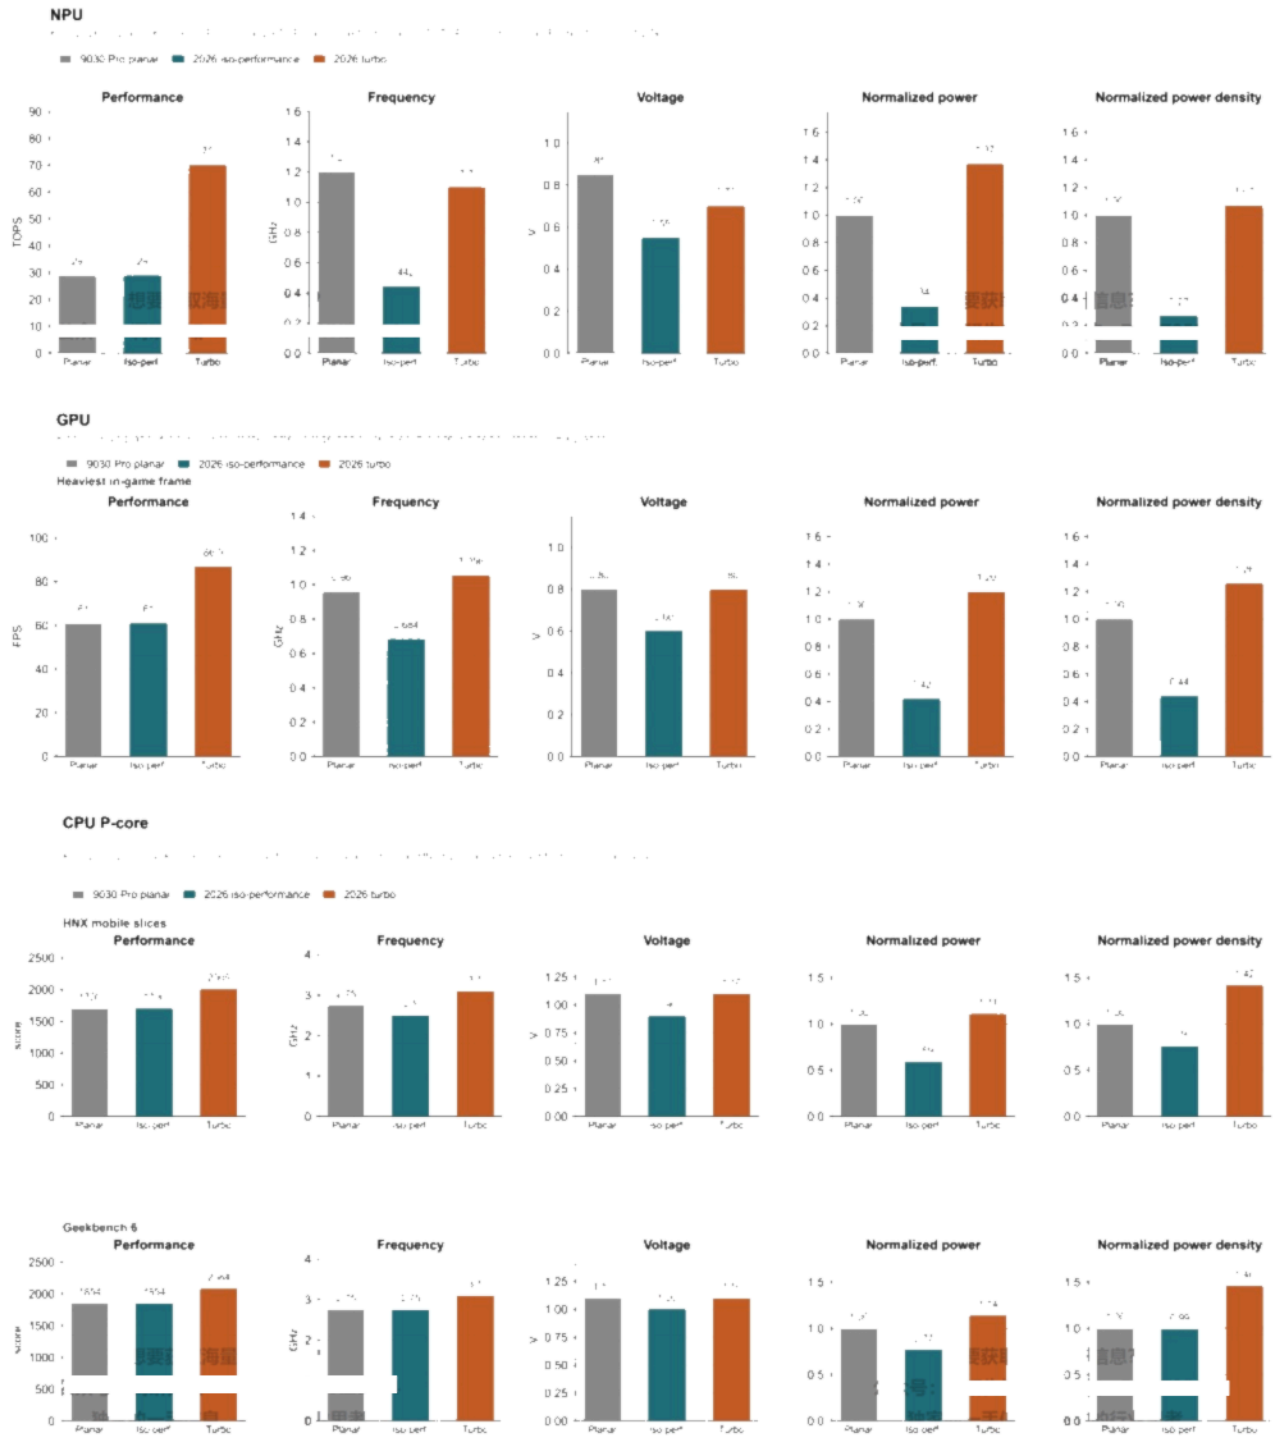
Source: Company reports, Bernstein analysis

EXHIBIT 17: **(a) The schematic illustration of the next-generation Kirin SoC platform; (b) cross-section image of its bonding interface**



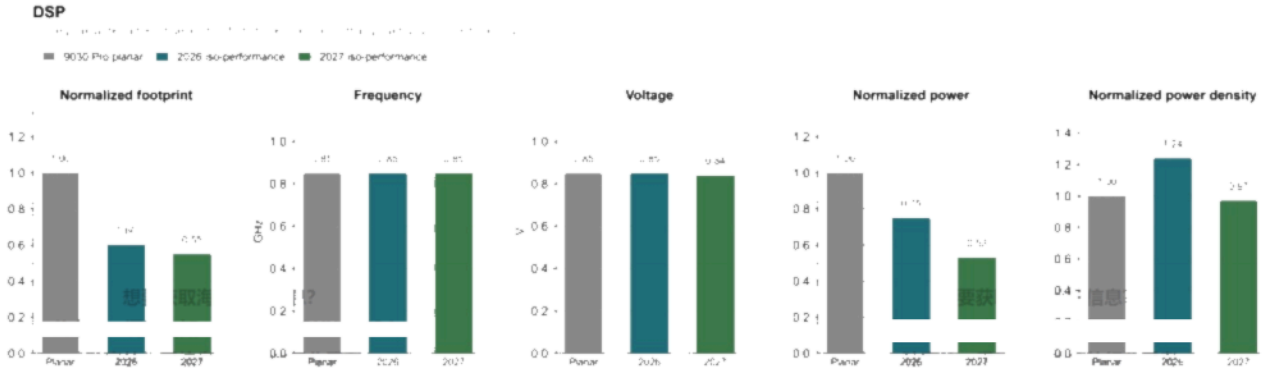
Source: Company reports, Bernstein analysis

EXHIBIT 18: **Kirin 2026 vs. planar Kirin 9030 Pro - iso-performance and turbo**



Source: Company reports, Bernstein analysis

EXHIBIT 19: **DSP is the warning case - first-generation fold shrank area faster than power**



Source: Company reports, Bernstein analysis

Tests of Kirin 9050 Pro versus Kirin 9030 Pro performance on the phone follow the same pattern as the laboratory test. The lab test first held each block of Kirin 9050 Pro to the Kirin 9030 Pro's score and recorded power reduction, then let the chip run faster (turbo). On the phone, the CPU same-clock test is the power side of that: the Kirin 9050 Pro prime, locked at 2.75 GHz, does the same integer work as the Kirin 9030 Pro prime and uses more than 30% less power. The lab's CPU power cuts were 23% on Geekbench 6 at that same 2.75 GHz clock, and 41% on the P-Core when the score was held to the Kirin 9030 Pro's. The GPU and NPU tests on the phone are the turbo side: Steel Nomad Light increases 39% versus that of Kirin 9030 Pro, against the lab's +42% frames per second; INT8 NPU raises about 150% versus that of Kirin 9030 Pro, against the lab's +141%. CPU peak on the phone is a smaller step than GPU or NPU, and it is below Huawei's launch claim. Geekbench 7 increases 14% on single-core and 20% on multi-core versus that of Kirin 9030 Pro, in the same range as the lab's P-Core turbo of +18%.

At turbo, NPU throughput is up 141%, GPU FPS increases 42%, and P-Core score is 18% higher, with power density above the planar Kirin 9030 Pro. That is 70 TOPS, 86.9 FPS, and a 3.1 GHz P-Core. In other words, the same silicon that is cooler at matched work will become hotter once the slack is spent on clock.

NPU power falls 66% even as transistor density rises 55%. The predecessor was one big NPU core plus two efficiency cores. LogicFolding puts more transistors into the same top-down chip area; those extra transistors are used to place four big NPU cores rather than to run a single core faster. At the Kirin 9030 Pro's 29 TOPS, each of the four can clock much slower and at lower voltage (from 1.2 GHz to 0.44 GHz, and from 0.85 V to 0.55 V), so block power falls even though density is up. Turbo is the other use case of the same four cores: 70 TOPS at 0.70 V, still below the predecessor's 0.85 V.

P-Core Power is down 41% at the Kirin 9030's score, the smallest saving among NPU, GPU and the P-Core. Frequency only moves from 2.75 GHz to 2.50 GHz. One integer thread still walks one pipeline, so LogicFolding cannot split that work across extra cores. The paper's path to a 5 GHz-class P-Core is a denser vertical connection - bond pitch 1.5 µm this year, 1.0 µm on 2027 silicon, then 720 nm - plus more stacked layers, which would shorten the remaining long paths. Kirin 2026 generation only folds selected slow paths inside the core, not the entire CPU into 3D.

DSP is the block where the first-generation fold made heat per area worse. Power fell 25%, but the area shrank 40%, so watts per square millimeter rose 24%. NPU, CPU and the P-Core all cut power faster than area at the Kirin 9030 Pro's score; but DSP did not. 2027 silicon is claimed to bring that heat density back below the unstacked Kirin 9030 Pro.