



26 January 2026

MIMOSIS-3 Design Review Report

Stefan Biereigel, Vladimir Gromov

Keywords:

Summary

This document summarizes the discussions at the MIMOSIS-3 design review held ahead of the production readiness review (PRR) to be performed in January 2026.

Contents

1	Review scope and contents	2
2	Overview of known issues in MIMOSIS-2.1	2
3	Modifications addressing data transmission reliability issues	3
3.1	Reduction of random jitter generated in the PLL	3
3.1.1	Modifications proposed by the designers	4
3.1.2	Recommendations by the reviewers	4
3.2	Reduction of activity-dependent jitter in the output data path	5
3.2.1	Modifications proposed by the designers	5
3.2.2	Recommendations by the reviewers	6
3.3	Bypassing of the PLL	7
3.3.1	Modifications proposed by the designers	7
3.3.2	Recommendations by the reviewers	8
4	Summary	8

1 Review scope and contents

The MIMOSIS-3 design targets being the final submission (and hence the mass production design) of the MIMOSIS sensor ASIC. Both schedule and budget do not allow for a re-spin of the design in case of further issues. For this reason, the submission is to be treated as 'high-risk' and the review strategy has been changed accordingly. The design review held on 11 December 2026 focused on a set of well-defined remaining issues. It is not understood as a review of the complete ASIC design, so reviewer exposure to potential issues is limited to those discussed. The comments provided in this report are therefore also concerned only with these issues.

2 Overview of known issues in MIMOSIS-2.1

A set of issues in the existing MIMOSIS-2.1 design were presented to the reviewers, which are briefly summarized below in decreasing order of applicability to the review scope:

- Output data jitter: The data transmitted to the GBTx is impaired by both excessive random jitter as well as jitter components that are strongly correlated with on-chip sequencer activity.
- 320 MHz clock backup: The ASIC cannot be operated using only a single 320 MHz clock (which is considered as a backup option), but requires an additional 40 MHz input clock for the slow control blocks. This additional clock signal cannot be made available in the final environment.
- I²C pad issues: The used CMOS pads are suspected to create bus conflicts, leading to large transient current draw affecting other circuitry.
- SLVS pad issues: The used differential receivers have no built-in fail-safe network (weak pull-up/down resistors on the differential lines), making them susceptible to noise pickup. External termination can not be incorporated on the FPC cables.
- Power monitors: Power connections for some on-chip power monitors were incorrectly implemented when modifying the supply scheme, leading to them being non-operational. Additionally, the output of the power monitors can not be disabled on demand, prohibiting the sharing of a common output line between multiple ASICs.
- Matrix reorganization: Some changes to the sensor matrix must still be implemented, unifying their size and making the corresponding changes to the threshold adjustment DACs.

Most time during the review was spent on the first two issues listed, in order to ensure that clocking and data transmission of the ASIC can function reliably. The remaining listed issues were discussed to ensure the reviewers correctly understood them and the proposed mitigation strategies for these were discussed qualitatively.

3 Modifications addressing data transmission reliability issues

The presentation highlighted that two separate issues are degrading the quality of the transmitted data signal. When operated in the intentional mode with the built-in LDO, higher-than-expected random jitter (about 130 ps rms, i.e. $1.8 \text{ ns}_{\text{pk-pk}} @ \text{BER}=10^{-12}$) is observed. Additionally, activity-dependent deterministic jitter components with a magnitude of up to $300 \text{ ps}_{\text{pk-pk}}$ have been observed. Both issues combined lead to a substantial horizontal eye closure for the 320 Mb/s data link, leading the GBTx-based back-end to lose tracking of the frame.

3.1 Reduction of random jitter generated in the PLL

The designers described a set of extensive simulations understanding the nature of the excessive random jitter observed at the PLL output, in particular when the LDO is used. To allow a more realistic assessment of the random jitter at the PLL output, a model of the reference clock phase noise was incorporated as recommended in a previous review.

The supply strategy was reviewed, and the designers confirmed that all blocks of the PLL are supplied from a single supply voltage, which is not shared with the digital functional blocks of the ASIC. At the system level, these domains are supplied by independent FEAST DC/DC regulators. It was concluded that this should minimize any inter-domain coupling, in particular since also the on-chip ground network for these supplies is separated and shared paths are only present on the PCB.

Spectre 'pss' and 'pnoise' simulations have been able to reproduce the increased jitter when the LDO is active. Simulations and measurements agree well for these simulations. Simulations further showed that increasing the bias current provided to the LDO does not result in a decrease of the PLL jitter. A stability analysis on the LDO circuit was performed and shown that the LDO itself is stable when providing the nominal VCO supply current. A deficiency was found when very lightly loaded (VCO disabled).

Jitter was found to decrease (with the LDO in use) when increasing the charge pump current and keeping the LDO bias current constant. This is a somewhat surprising finding, since the jitter performance with the same charge pump current is much improved when the LDO is bypassed. It does not appear to be fully understood why this is the case, and the simulations presented appear to have inconsistencies (w.r.t the reference clock phase noise model) that make interpreting the results difficult. A better understanding could be obtained by inspecting the per-device noise contribution from the pnoise simulation.

The existing ASIC prototype does not allow for independent adjustment of the two currents since they are generated as a fixed multiple of a single, externally provided reference current. This complicates substantiating these simulation results with measurements. Measurements have however confirmed that increasing both bias currents together is generally a viable strategy. Both simulations and measurements predict a substantial reduction of the jitter (by a factor of 3-6) by increasing the charge pump current, making operation with the LDO a viable option.

3.1.1 Modifications proposed by the designers

Three key modifications have been proposed by the designers:

- The new design foresees separate trim registers for both the charge pump and the LDO bias current. This will allow the bias currents to be increased to a value that provides sufficient jitter performance. One consequence of this modification is the required addition of additional slow control registers for these trim bits, increasing the implementation effort required. **To be clarified:** What is the range and resolution of this trimming? What is the power-on default value (ensuring correct PLL startup)? What is the expected impact on ASIC power consumption, if any?
- The new design adds a 200 microampere load to the LDO output when the VCO is disabled, such that the LDO is unconditionally stable when the ASIC is powered. Because an instability of the LDO has not so far been proven to cause the excessive jitter seen in simulations, it might not be necessary. It is however considered a beneficial addition for the proposed operational backup solution of bypassing the PLL circuit completely and as such the reviewers suggest implementing it.
- On the system level, the designers consider bypassing the LDO to be a possible operational mitigation for the jitter issue. This is already possible in the existing ASIC design, but it remains to be proven that this solution will provide sufficient performance in the detector environment. Since discussions between designers and reviewers did not reveal a particular justification for the presence LDO in the design and encouraging test results, the reviewers suggest continuing to study this option as a backup should the prior mitigation not be successful at reducing the jitter to the required level.

3.1.2 Recommendations by the reviewers

The reviewers generally believe that a credible path for reducing the output jitter was identified and that the proposed modifications will result in significant performance improvements. Both the increase of the charge pump current as well as the bypassing of the LDO for the final system are likely to reduce jitter to a level acceptable for reliable data transmission (when also addressing the power-supply induced jitter in the output data path, discussed next). The reviewers recommend the following additional actions to be taken to increase confidence in the success of the revised design:

- To increase confidence in the proposed bias current adjustment, the reviewers recommend an analysis of the per-instance noise contribution and pinpointing the key contributing devices. Better understanding the origin of the random jitter contributions in the critical frequency range (around 1 MHz) could prove crucial for assessing the reliability of this modification and viability for the final system.
- The reviewers recommend performing a steady-state loop transfer function analysis with the proposed charge pump current value in addition to the noise analysis, to ensure jitter peaking around the loop bandwidth is not excessive and that sufficient stability margin exists. As an alternative, a transfer function measurement may be

carried out on the existing measurement setup (using a phase-modulated reference clock signal and measuring input-to-output jitter transfer).

- To guarantee success of the backup strategy of bypassing the LDO, the reviews suggest carrying out further systems tests, during which the ASIC is powered from a representative power supply network (using FEAST DC/DC converters, realistic supply impedance) to demonstrate that sufficient jitter performance can be obtained in the final environment with the PLL LDO bypassed. Especially line noise contributors (such as DC/DC switching noise) should be critically reviewed to ensure that PSRR performance of the PLL alone will provide acceptable performance.
- The reviewers agree with the proposal of not modifying the clock divider, since gains here appear to be negligible compared to the major contributors.

3.2 Reduction of activity-dependent jitter in the output data path

A major finding with respect to the last review meeting is a credible explanation of the activity-dependent jitter components observed on the data outputs. The designers conclude that the origin for this jitter component is to be found in supply variations on the digital supply voltages affecting the propagation delay of buffers ('clock tree' and 'data path'). A power integrity simulation has identified internal voltage ripple to approach 100 mV. The sensitivity of the clock and data path amounts to about 5 ps / mV, giving rise to the amount of jitter that is observed.

Simulations were presented that demonstrated very good agreement between the modeled clock/data path delays and its interaction with the supply voltage ripple generated by the sequencer activity. It was concluded that addressing the large supply voltage ripple itself is rather difficult at this stage, potentially requiring deep architectural changes to the ASIC. Instead, mitigation efforts were focused on the sensitivity to this supply noise.

The designers acknowledged that the existing delay on the data path (and hence the sensitivity of the data path) is much larger than required and could be reduced substantially. Particularly the use of delay cells on the data path was stated to be a major contributor to the existing situation. Discussions on system-level impacts of changing these delays revealed that no particular need for matching delays among different data outputs exists, and also that introducing further delays (e.g. from pipeline stages or FIFOs) is not considered problematic.

The designers acknowledged that while the shown simulation models the data jitter profile well, it does not account for the 'undershoot' observed in measurements. This was explained as being a response of the PLL to the supply drop, which is not included in the simulation model.

3.2.1 Modifications proposed by the designers

The proposed modifications for this issue are at a less advanced stage. Two conceptual solutions to the problem were discussed:

- Introduction of an asynchronous FIFO located very close to the output pad. The output path of this FIFO would be clocked with the same clock frequency, but distributed there using a much shorter clock tree. The asynchronous nature of the FIFO

reduces concerns relating to timing closure between clock 'domains' with very different insertion delays. While the intention would be to use a 'design ware' (pre-verified) IP block for this FIFO, the potential impact of single-event effects and the consequence of non-deterministic output delay (repeatability across reset cycles) must be kept in mind.

- Alternatively, a more simple resampling/pipeline solution was proposed. A D-flip flop placed close to the output could re-sample the data, again using a clock signal obtained from a much shorter clock tree. The challenge with this approach is that timing towards this resampling flop needs to be closed in all corners.

3.2.2 Recommendations by the reviewers

- While both proposed modifications can potentially suppress the issue to the same extent, the reviewers recommend prioritizing the pipelining/resampling approach over the addition of a FIFO on the output data path. While the FIFO solution might give the impression of a simpler implementation, the reviewers perceive it as a higher risk solution due to potential implications on SEE tolerance and verifying of design constraint correctness. Timing closure seems possible to obtain with the resampling solution, even if potentially two or three pipeline stages may be required to fully absorb the clock and data delay variations. Even in this case, the power and area implications of this solution should be smaller than the alternative FIFO implementations.
- Since the proposed mitigation only reduces the timing jitter but does not address the underlying supply voltage variations, the reviewers suggest remaining mindful of vertical eye closure induced at the SLVS pad. Compatibility with the GBTx inputs should be verified as far as possible.
- The reviewers suggest incorporating a model of the off-chip PDN (bondwire inductance, important off-chip decoupling capacitors, PDN resistance) into the simulations to better assess the margins and catch potential issues from PDN resonance effects. The updated estimates of dynamic IR drop can be used both for an updated prediction of the jitter but also of the expected vertical eye opening.
- Given the proposed data path modifications are done outside of the PLL control loop, it should not be expected for the 'undershoot' feature observed in prior measurements to disappear. While this is not the dominant issue currently, the reviewers recommend verifying that the remaining supply-induced jitter component is acceptable for reliable data transmission.
- Finally, the review did not cover whether the shown 'sequencer activity' profile already presents the worst-case supply ripple condition expected during operation. Possible activity scenarios should be tested/simulated as far as possible) to make sure that the margins (3x IR drop) assumed in some simulations actually envelop the expected worst-case dynamic and static IR drop.
- On a wider note, the reviewers recommend critically reviewing the digital design timing reports for sufficient margins to absorb the relatively large dynamic IR drop. Timing

models must be derated from their nominal values to properly account for instance voltage drop in addition to statistical variations. Measurements at elevated temperatures and reduced supply voltages (below the foundry corners or outside of the expected operational voltage ranges) should be considered to ensure that margins are appropriate.

A complementary approach to substantially improve the power supply sensitivity could be the provision of an additional supply voltage connection to be used only for the delay-critical clock and data paths. Using this approach, the value of the power supply voltage would not be affected by the voltage-drop-related fluctuations due to data generation on the chip. When the power supply voltage for the important cells (clock and data trees towards the output) is stable the delay instability in the clock-tree will be feasible to control properly. The reviewers acknowledge that this is a modification difficult to implement at this stage of the design and a careful assessment of the associated risk must be performed.

3.3 Bypassing of the PLL

A third avenue explored to improve reliability of the data transmission is the operation of the ASIC completely without the PLL. Even though an external 320 MHz clock can be provided, the existing design cannot operate with only this clock signal, since a 40 MHz clock signal is required additionally to operate the slow control logic.

In case the PLL issues can not be successfully addressed, introducing a backup operation mode that can run the ASIC fully on an externally provided 320 MHz clock would be beneficial, especially if it comes at a low complexity increase for the ASIC design. The external 320 MHz clock (CLKRESCUE) will be the source of the clock tree. Assuming the CLKRESCUE clock is always aligned to the data sampling clock in the GBTx and the data generated on the chip is aligned to the CLKRESCUE clock, proper alignment of the DATA to the sampling clock will be guaranteed. In this mode, there would be no need for the PLL or LDO blocks except the use of the feedback divider to derive the 40 MHz clock used in the slow control subsystem.

On the system level, it is acknowledged that this introduces the challenge of having to supply a 320 MHz clock across FPC cables without good impedance control, raising issues of signal integrity and EMC.

It was discussed that while this mode of operation has been tested on the existing ASIC to some extent, it has not been fully qualified to work in the final detector system. Issues related to other issues (coupling into floating SLVS input pads) have been observed in this mode of operation, and these have not been fully understood or followed up yet.

3.3.1 Modifications proposed by the designers

The designers proposed to introduce the following modifications into the design to enable operation from a single, externally provided 320 MHz clock:

- Introduce multiplexing to allow for direct provision of 320 MHz instead of the PLL output, plus use of the built-in divider circuit to drive the 40 MHz clock signal used in slow control applications.

- Re-purposing of an unused input pad to switch between the two modes of clocking the ASIC.

3.3.2 Recommendations by the reviewers

The reviewers strongly support the implementation and validation of this additional mode of operation. In case the modifications presented in the previous sections fail to deliver the expected performance improvement, this option offers an additional avenue towards a reliable detector readout. Even though it is not to be considered the baseline for operating the ASIC readout and while it comes with potential complications at the system level, the reviewers suggest clearly identifying any potential showstoppers for this mode of operation to make it a viable fallback possibility.

In particular, a study of signal integrity and crosstalk issues observed when using FPC cables should be performed. It should be assessed whether these can be successfully mitigated by addressing the SLVS pad issues or if further issues should be expected.

While this operational mode offers the promise of eliminating the random jitter contributed by the PLL (which, according to the simulations presented, is currently the dominant fraction), it imposes a separate set of constraints of its own. Coupling from the ASIC output data pins to the 320 MHz input clock signal must be kept under control, such that no excessive jitter is introduced at the ASIC clock input as soon as the ASIC is powered up. As far as it's feasible, separation of the clock and data lines on PCBs and cables towards the ASIC should be maximized. Simulations or measurements should be carried out to understand potential coupling problems ahead of the final implementation.

4 Summary

The reviewers believe that the design team has gotten a good handle on the issues affecting the data transmission of the MIMOSIS-2.1 prototype since the last review. The team is following a sound strategy for reproducing the issues in simulations, implementing changes and independently reviewing these changes both internally and externally. A risk/impact assessment of each proposed modification has been performed, with which the reviewers broadly agree.

All issues appear to be understood at the level of being reproducible in simulations. Mitigations for all issues have been identified that were prototyped both in simulations and in measurements using the existing ASIC prototype. As of today, there seems to be no definitive root-cause-level understanding of the random jitter affecting the PLL when operated with the LDO enabled. Suggestions have been made by the reviewers to potentially identify the responsible circuit component should this be considered a critical shortcoming of the existing design. More importantly, the reviewers believe that a sufficient number of fallback solutions are being proposed for this issue (operating without the LDO or bypassing the PLL completely). This mitigates some of the risk of new changes being introduced late in the design cycle. Some additional work might still be needed to validate that these fallback solutions are actually viable in the final detector environment.

Some analysis and implementation work is still required for addressing the activity-dependent data jitter introduced in the clock/data tree of the digital design partition. While

the presented strategy here appears credible and possible to implement in the short amount of time available with the presented resource profile, the amount of improvement attainable must be quantified more accurately.

Ideally, a jitter budget accounting for all expected sources of jitter (e.g. random jitter, power-supply induced jitter per voltage domain, FPC cable channel including losses and crosstalk) should be established for various operational scenarios (typical and worst case module environments, different ASIC operating modes) to assess whether the combination of all proposed modifications will be sufficient reduce the total jitter sufficiently for the output data links to operate at or below the desired bit error rate with sufficient margins. Particular focus should be put on better understanding the data channel and supply impedance characteristics (in particular for the longest cables expected), such that any additional degradation can be quantified and incorporated in this assessment. Verification through analysis (SI simulations using EM solvers) or measurements (e.g. comparative eye diagram measurements using the existing ASIC or network analyzer measurements for channel modeling) can both be considered appropriate for this task.

References