Keck Adaptive Optics Note #366

Development Team Response to the NGWFC Preliminary Design Review Committee Report

Stomski, Johansson, van Dam, Wizinowich, Sumner, Biasi

October 27, 2005

1 Introduction

On September 7, 2005 a preliminary design review (PDR) was held for the Next Generation Wavefront Controller (NGWFC) project here at W. M. Keck Observatory. Subsequent to that meeting, the review committee issued a report detailing their evaluation of the NGWFC preliminary design. This document is the development team's response to the PDR report. The PDR report itself is available as KAON 365. Since the report is quite short, the complete text of each section is included below for reference as we discuss the development team responses to the points made in that section.

The committee's report is organized into comments under 8 headings. This response is formatted with sections that have a 1-to-1 correspondence to those of the committee report. In the appendices of this document, we have attached the KAONs describing analysis done to address issues brought up by the review committee in the original PDR report and referenced in this response.

1.1 Table of Contents

2 Responses to points made in the *"Summary"* section

> We believe the Keck Observatory NGWFC project has reached a technical and managerial maturity appropriate for a Preliminary Design Review. The key MicroGate subcontract appears well established with good communications with the team. The technical design of the MicroGate real-time compute engine appears capable of meeting all functional requirements. The Project Scientist appears to have a good understanding of the current AO system limitations and the expected benefits of the NGWFC upgrade, although we continue to be concerned that some functional requirements for this project are not adequately justified in terms of expected science gain. Moreover, we expect the scope of the upgrade to be too ambitious in terms of both available resources and implementation schedule, particularly on Keck 2. Accordingly, we recommend that project and program management proactively undertake, with significant Keck Adaptive Optics Working Group (KAOWG) science input, a detailed review of functional requirements and provide a science-driven prioritization of capability, should project resources prove insufficient.

2.1 Scope too ambitious in terms of schedule and resources

We agree that this is a valid concern. We have attempted to strike an appropriate balance between being aggressive and mitigating risk. With the input of the review committee in mind we will re-visit this issue and present a plan that we believe is reasonable at the DDR.

2.2 Prioritization of functional requirements

In addition to this mention of the prioritization of the functional requirements in the summary section of the of the PDR Committee report, it is also mentioned in the section dedicated to "Functional Requirements". Our consolidated response appears in section 3.1 of this document.

3 Responses to points made in the *"Functional Requirements"* section

> The committee is concerned that functional requirements for this project have not been derived from formal analysis of current laser guide star (LGS) system performance. Although we appreciate that key members of the NGWFC team (notably the NGWFC Project Scientist) have other critical Observatory responsibilities, we remain concerned that some requirements are derived from anecdotal sources. We believe a prioritization of functional requirements, with strong science oversight (such as may be provided by the KAOWG), will help solidify both the rationale behind certain requirements (e.g. number of subaperture pixels, maximum wavefront sensor frame rate), and help CARA to plan for the LGS performance improvement tasks not in the direct scope of this upgrade.
>
> As the primary example of the above, we are concerned that more advanced centroiding algorithms may be necessary to fully realize the top-level project goal of improving LGS. This is likely to be the first 're-programming' of the NGWFC by CARA staff, so CARA should endeavor to understand/implement this functionality soon after hardware implementation. Plans for this should be developed by DDR, to ensure no hardware limitations are imposed.
>
> Another example of loose connection between science and functional requirements is the

dark current requirement for the WFS CCD. The derivation of the requirement is at best vague in the PDR documentation.

Similarly, we strongly feel that denominator-free centroid operations are valuable for extending faint NGS operations. On the other hand, we question whether 2,000 frames per second operation of the wavefront sensor camera and real-time reconstructor are justified as benefiting LGS operation, the stated goal of the NGWFC upgrade.

Specifically, we recommend that SNR calculations for CCID56 8x8 mode and CCD39 4x4 mode be compared prior to implementation.

3.1 Prioritization of functional requirements, with strong science oversight

We would welcome additional input from the AOWG on the relative priorities of system capabilities. We will discuss how to best involve the AOWG on this issue with the AOWG co-chairs. Obviously there are time constraints involved, so initiating and conducting this exchange will need to happen quickly. Our current thinking is that the most efficient way of soliciting AOWG input in this area would be for us to draft a document describing our recommendation for the relative priorities of the capabilities and get one or more members of the AOWG to comment on it.

In the Detailed Design phase, we plan to go through the functional requirements and establish relative priorities for each. The science implications of each requirement and the related capability will be of significant importance in determining the relative priorities. Low priority items may be de-scoped or be deferred during the project planning or implementation as needed

We will have to consider the relative priorities from several other perspectives as well. The metrics used to judge relative priority will include:

- Science implications
- Operability and maintainability implications
- Schedule and human resource implications
- Cost implications
  - Short term costs (during the development)
  - Long term costs (operational and future upgrade costs)

3.2 Feasibility and risk assessment for 8x8 correlation based centroiding

We have documented the algorithm in terms of storage and computational requirements. We have shared this with Microgate and will work with them to verify the feasibility and risks involved in implementing this capability as a future upgrade. This will be well understood by DDR.

3.3 Is the 2 kHz framerate justified?

We understand the committee's concern regarding the 2 kHz framerate. However, relaxing that requirement will not change the project in any fundamental way. Analysis shows that the Microgate HW and SW are capable of performing at a 2 kHz framerate. Therefore, the Microgate deliverables will not change, and neither will the CARA scope of work or schedule. Hence, while it may be desirable from a science point of view to prioritize (or possibly de-scope) the 2 kHz requirement in relation to other functional requirements, doing so will not have any real effect on the project.

3.4 SNR calculations and dark current requirements for the WFS CCD

Subsequent to receiving feedback from the committee, we performed SNR calculations for both the CCD-39 and the CCID-56 for various binning modes and frame rates. We also addressed the issue of dark current in the CCD. The results of these analyses are presented in **KAON 359: Wavefront sensor observing modes**, and will be included in the DDR documentation.

3.5 Denominator free centroiding

The current design provides the necessary stubs within the Microgate-supplied RTC to implement denominator-free centroiding. Specifically, requirements 13 and 137 call for a vector of user specified denominators (one per subaperture) to be supported. This gives us the framework to load a denominator vector into the RTC. The current requirements do not call for its usage, leaving the implementation to CARA as an upgrade outside the scope of this project. The implementation is a simple mod to the centroid algorithm to use the user-specified denominator vector value in place of the calculated value if the user-specified value is non-zero. Operationally, the values in this vector can be managed by a process that could be running on the WCP or in an external IDL tool.

Since receiving this feedback from the committee, we evaluated the use of denominator-free centroiding using end-to-end simulations. The results are documented in **KAON 357: Simulations of denominator-free centroiding**. The conclusion is that there is some benefit at very low light levels (<=20 counts per subaperture).

4 Responses to points made in the *"Documentation"* section

> Although the large number of technical notes (i.e., KAONs) provided most of the information needed for our review, we found it to be somewhat cumbersome and at times redundant. We recommend streamlining future documentation into reasonably complete and self-contained documents for the purposes of formal project documentation and reviews. We believe a simpler document repository, containing, for example,
>
> a. System Requirements Document
> b. Software Requirements Document
> c. Conceptual Design Document (if a conceptual design is performed)
> d. Preliminary Design Document
> e. Detailed (or Critical) Design Document
> f. As-built documentation (mechanical drawings, schematics, PC board fab drawings, source code, etc.)
> g. Final Test Document (including test plan, procedures, and results)
>
> would be a collection more easily referred to by the NGWFC project team.

This is a new approach to project design documentation being used for this project. We appreciate the feedback and we do intend to streamline it for the DDR.

5 Responses to points made in the *"Detailed Design Review"* section

> The committee is concerned that DDR is scheduled to follow PDR so closely. We understand management constraints, but feel that the following must minimally be provided for a successful DDR:
>
> a. A more formal description of the software design, in particular the MicroGate portion, should be provided. This should include structured English (also known as pseudocode) for key algorithms, and a top level flow diagram. Equally important, it seems to us that the specific MicroGate deliverables for DDR require

better definition. (We would expect certain documentation to be formally delivered prior to DDR.)

b. A more formal description of the guaranteed telemetry response time that will be provided by the MicroGate Processor. The requirement states 10 clients and a begin response with 0.1 seconds. How long until the data is served, as a function of the request? The discussion at the review did not provide any formal analysis or testing describing how this would be achieved. In particular, the choice of whether requests would be handled via monitors or polling was not yet decided.

c. A more detailed set of functional requirements and design specifications for the user tools. For most of the items in the WBS dictionary section 12.1 (page 21 of KAON #347), a one or two page specification, possibly including screen layout, should be provided.

d. Complete list of new and deleted channels and keywords.

e. A more complete top level functional block diagram similar to that given on page 10 of KAON #289. In particular, provide more detail on external systems, label arcs, and provide sub-diagrams for the larger components (e.g., WFP).

5.1 DDR is scheduled to follow PDR too closely

Our original plan was to hold the DDR on or about November 7, 2005. That date was derived based on the work we expected to do in the detailed design phase. The committee's specific suggestions in this section for material to be provided at DDR, along with the response and follow-up activities associated with other issues raised in the committee's report will require additional work by the project team. With all of this in mind, as well as the committee's concerns about the insufficient length of the detailed design phase, and our own commitment to making the DDR an unqualified success, we have revised our proposed schedule. Our current recommendation is to have the DDR early enough in December that we are assured to receive the committee's report before the holidays. This should give us enough time to complete all the work needed and at the same time should not impact the final Microgate delivery date. In the DDR report we will present our recommended implementation and I&T schedule, revised as needed in light of all the work done in the detailed design phase. The actual date and logistical details for the detailed design review itself will be worked out in collaboration with the review committee and the project team by Sean Adkins.

5.2 A more formal description of the software design

We plan to add flow diagrams and pseudo-code for the DDR documents and will ask Microgate to do the same for portions of code where that level of detail is considered worthwhile.

5.3 Clarifying the Microgate DDR deliverables

We believe the Microgate deliverables for DDR are well defined and understood by the development team. Perhaps we did not communicate them clearly enough to the review committee. They include:

- Detailed preliminary acceptance test procedure
- WIF-WCP interface detailed definition
- WIF-WCP interface command list
- Command Interpreter and Executer detailed definition (flow charts)
- Schematics of all proprietary boards, including component placement
- Assessment of HVC analog simulation

- Results of HVC analog circuit breadboarding
- TRS query detailed definition
- Results of preliminary tests on TRS query performance (not on final implementation)
- Validation of assumed computational times through testing of centroid computation algorithm written in C
- Detailed integration and test planning
- Final choice of COTS components (jointly with CARA)
- Test of IP over Fibre Channel interface test between RTC and TRS

5.4 A more formal description of the guaranteed telemetry response time

We will address this topic in the DDR documents.

5.5 Functional requirements and design specifications for the user tools

We will provide these at the DDR.

5.6 Complete list of new and deleted channels and keywords

We will provide a complete list of existing keywords for the DDR showing which are to be reused as is, which are to be modified and reused, and which are to be deleted. We will also provide a list of all new keywords to be implemented.

Regarding the EPICS channels, we plan to list the channels that are tied to keywords or that are needed to support interfaces to other external subsystems such as the SC and IF. Given that there are over 560 channels used internally in the current WFC our list will focus only on those that need to be supported by the NGWFC.

5.7 A more complete top level functional block diagram

We will expand the diagram from KAON 289 and add sufficient detail to the interfaces and data paths to make it more comprehensive. We will also identify the specific “external subsystems” that will interface with the NGWFC.

6 Responses to points made in the *“System Architecture”* section

> We commend the team for the attention given to monitoring system health, but encourage the design team review the audible voice alarms that have been added to the existing AO system, after the fact by the operations team, with an eye toward designing this upfront into NGWFC.
>
> The committee feels that use of the UTT strain gauge feedback is not currently supported by scientific gain and recommend management review its necessity/desirability.
>
> The issue of driving the DM and tip tilt mirror at rates higher than in the current system has not been addressed. This is separate from the design decision to include strain gauge feedback in the tip tilt controller. This may be outside the scope of NGWFC, but appears to us a 'hole' in the plan for improved science with NGWFC.
>
> One committee member has learned from PI that you can get another 10-15% of range by exploiting the negative voltage capability of the actuators. You might want to consider

building this into the new FSM electronics.

We recommend retaining the center subapertures in the controller calculation for purposes of alignment.

6.1 Audible alarms for system health

We are familiar with the audible alarms implemented by the operations team. In fact, one of our team members helped in the design of the WFC crash detector. We plan to implement as much fault monitoring as possible at a low level in the WCP. Fault occurrences which are detected can easily be tied to many different audible alarms at a higher level (i.e., user tools). Furthermore, we will work with the operations team to make sure that appropriate high-level audible alarms are designed for the NGWFC.

6.2 Justification for TT strain gauge feedback

The justification for DTT and UTT strain gauge feedback are addressed separately below. In the current NGWFC design, the DTT and UTT controllers are two instantiations of the same basic controller. The strain gauge feedback is built into the high voltage amplifiers of these controllers and its use is selectable through a software switch.

We did investigate the possibility of de-scoping this feature for both the DTT and UTT, and learned that doing so would result in a savings of at most 1 or 2 man months of effort for Microgate according to Roberto Biasi. The revised system would still be based on Microgate's existing hardware that is capable of reading the strain gauge feedback, but the software would not use the information.

DTT: Subsequent to the review, we revisited the issue of strain gauge feedback in the DTT controller and determined that it is indeed required for optimum tip-tilt system performance and for preventing uncontrolled oscillations in the tip-tilt mirror. This is documented in **KAON 355: Characterization of the tip-tilt mirror**. After considering both the cost issues and the performance benefits, we have decided to retain our current design approach, which is to implement and use the strain gauge feedback in the DTT controller.

UTT: We agree that the performance benefit of using strain gauge feedback in the UTT controller is unclear at this point and requires further analysis. Microgate has estimated that de-scoping the UTT portion but retaining it for use on DTT will result in virtually no savings of time or money. Hence, we have decided to keep the strain gauge feedback in the system for possible use on UTT as well.

6.3 Concerns regarding driving the DM and tip tilt mirror at higher rates

Deformable mirror: Tom Price at Xinetics has indicated that the resonance of the DM is in the 20-25 kHz range, and most likely cannot be excited by the controller. He further states that customers routinely run Xinetics DMs at 5 kHz. Finally, Xinetics have even measured the small signal response of the mirror up to 20 kHz with no problems. Our conclusion is that there are no issues with driving the DM at rates up to 2 kHz with the NGWFC.

Tip-tilt mirror: The tip-tilt mirror has mechanical resonances at approximately 700 Hz and 1 kHz which are easily excited, even with our current system (see section 6.2 above and KAON 355). At loop rates higher than our current system maximums (672 Hz WFC, 1 kHz STRAP), exciting these modes will become even more problematic. We believe driving the tip-tilt mirror at these higher rates cannot be separated from the strain gauge feedback issue. Strain gauge feedback is required to properly tune the controller to dampen the response at the resonant frequencies. Furthermore, driving the mirror at high rates without the strain gauge feedback will most likely lead to mirror instabilities and oscillations which can only be mitigated by turning down the controller gain, resulting in poor tip-tilt performance and lower Strehl. Our conclusion is that strain gauge feedback is required in the DTT controller in order to drive the tip-tilt mirror at rates above our current maximums.

6.4 Exploiting the negative voltage capability of the DTT actuators

We have received e-mail from PI confirming the ability to drive the P-841.60 actuators beyond their normal limits to achieve as much as 40% additional stroke. Microgate is performing experiments to verify that the strain gauge readback functions properly in conjunction with overdriving the actuator. There may be some additional cost with this option, as a new power supply will have to be selected, and possibly two power supplies may be required, because the lower limit is no longer 0V. We will present the results of our work on this at the DDR.

There are two main performance benefits to be derived from additional stroke:

- Better tip-tilt performance during bad seeing. We routinely “rail” the tip-tilt mirror during poor seeing conditions. Increasing the stroke will result in better tip tilt control and better Strehl. We currently do not have any data to quantify the improvement.
- Increased dynamic range for the IF chopper. The IF team would like to chop at ±600 mas, which is 75% of the ±800 mas total stroke of the mirror. We are currently unable to do that with the current system due to instabilities and poor performance; hence we chop at ±300 mas. Increased dynamic range from the additional stroke combined with improved stability from closed-loop mirror control (strain gauge feedback) will result in improved chopping.

6.5 Retaining the ability to use the center subapertures for alignment

Since the PDR, we have investigated this option (see **KAON 358: Reading the subapertures behind the central obscuration**), and have concluded that the benefits are not enough to warrant a change order for to the RTC. We can however still implement this feature using the WFS raw frames in the telemetry stream and processing in IDL if needed.

7 Responses to points made in the *”Telemetry Data Product”* section.

> To make the improved NGWFC telemetry capability useful to observers, we recommend the project discuss with potential users the format, frequency, and content of the telemetry stream(s). While delivery of telemetry products is understood to be outside the scope of this upgrade, we believe that more of this kind of scientific contextualization is necessary to justify stated performance requirements.

7.1 Getting user input on the telemetry data products

As we were considering this issue, we realized that there is actually a more general issue to be addressed with regard to managing NGWFC data. While some of this effort will be operational in nature, we do want to help develop a plan to deal with it. The more general issue is that the system on the summit will be capable of holding ~5 Terabytes of data, which is enough for a 5 night AO run, and we need a plan to manage that data. Specifically we need to look at:

- When is the data moved from the summit to HQ?
- After moving the data to HQ, how is it archived? What format? On what media?
- Does the data need to be easily accessible? For how long?
- What is the long term data archive strategy?
- What data might the astronomers want to take with them? In what format?

To address all of these issues, we recently held a meeting with the support astronomers. From this interaction, we expect to have a draft data management plan for DDR. The plan will be flexible enough to allow the support astronomers to evolve aspects of the plan as we learn more through use of the systems.

8 Responses to points made in the *"Microgate Technical Design"* section

> We have found the MicroGate PDR documentation quite thorough and of high quality. Most committee technical issues were discussed at the review, to our satisfaction. However, a few outstanding questions/comments remain
>
> a. Although there is a discussion of what centroiding entails in the requirements, as listed in the MicroGate Design Document (see, for example, requirements 10 through 15 on page 104), there is no corresponding algorithmic discussion in the Software Design chapter (chapter 6).
>
> b. Why is the number of centroids updated after every 8 pixels (section 6.1.1)? If the number of pixels needed for a full set of centroids was not evenly divisible by 8 (an odd number of 2x2 sub-aps, for example), wouldn't that leave a sub-ap hanging at the end? Why not simply bump up the number of computable centroids every time n usable pixels are received, where n is the current number of pixels per centroid?
>
> c. We didn't recall seeing the (true) sustained storage rate for the Triton array in a RAID 3 configuration given. It wouldn't hurt to check that against requirements again.

8.1 Software design for centroiding

For the PDR, Microgate did not detail the algorithm because it did not differ in any significant way from the algorithm shown in the Keck requirements document, which is already exhaustive in its description. This was done to avoid repeating the same information in the various documents.

8.2 Question: Why update number of centroids after every 8 pixels?

The '8 pixel update' comes from the optimization of the algorithm on the DSP, that makes use of the SIMD (Single Instruction, Multiple Data) feature embedded in the device. In fact the number of pixel must be a multiple of 8, and dummy pixels are added at the end of the stream and simply discarded. This implementation is the one that actually maximizes the DSP throughput.

8.3 Sustained storage rate for the Triton array?

For the PDR, Microgate tested the Triton array only to verify that the throughput requirement of 13MB/s could easily be met, which it could. This ensures that the TRS has sufficient headroom for our requirements. The sustained data storage rate will be tested more extensively for DDR.

9 Responses to points made in the *"Management and Integration Plans"* section

> The committee considers both the K1 and K2 I&T schedule to be very risky. The time between first actuator motion on K1 and resumption of science using K1 appears

> insufficient. Premature resumption of K1 AO operations carries the risk of not understanding a key issue, then suffering the same problem on K2. The total allocation of "no-AO time" for K2 is likely also insufficient, based on our experience upgrading AO systems, with which we have also been very familiar. The panel strongly recommends that the Observatory consider allocating longer AO downtime for both K1 and K2. CARA should explore a strategy of greater upfront allocation, combined with a return of engineering time to science operation, upon successful I&T on both K1, then on K2.
>
> The committee recommends the I&T plan formally include a decision point, based upon Keck 1 implementation success, prior to dismantling of the Keck 2 AO system. CARA management should protect the project team from rushing into the Keck 2 upgrade, based on schedule demands, should significant issues from Keck 1 remain outstanding.
>
> Full regression testing of NGWFC functionality should be given sufficient resources to optimize use of telescope engineering time.

9.1 Risk mitigation for K1 and K2 I&T

During the detailed design phase we will look carefully at ways that we can balance the desire to do this upgrade quickly with reasonable risk mitigation strategies. We intend to explore options spanning from very slow and conservative to quick and aggressive. A very conservative plan would be guaranteed to take longer, but would have minimal risk of additional unplanned schedule slips. A more aggressive approach has the potential to complete the upgrade much sooner, but carries a higher risk that technical problems arise and significant unplanned schedule impact would be possible.

Our final plan will likely be somewhere in the middle of these extremes. The goals for the I&T plan itself are to:

- Get the NGWFC upgrade done reasonably quickly
- Have as small an impact on the observing community as possible during I&T
- Be as flexible as possible

The way we plan to accomplish this is to…

- Have a list of de-scope / phased development / human resource contingencies ready
- Have the goals and test plans for each step as detailed as possible
- Specifically consider likely problem scenarios
- Use additional day time and lab testing where appropriate.

9.2 How much "no-AO time" to allocate

We have already engaged Barbara Schaefer, who does the telescope scheduling here at Keck, for her help in coming up with a roadmap for how we will get to the ultimate recommendation for when the AO systems should be declared unavailable to astronomers. This decision will require the participation of several stakeholder groups, including the project staff, operations staff, Keck management and the user community.

There are several approaches being considered. One point raised is that the current NGWFC schedule has us doing K2 I&T at the very end of semester '06B. (December '06 and January '07) which would require us to make the decision to commit to a schedule for the "no-AO time" in February or March of '06. If, however, we scheduled K2 I&T for the beginning of semester '07A (February and March of '07) then we would not have to commit to the "no-AO time" schedule until August or September '06. So, basically delaying K2 I&T by two months would allow us to get an additional 6 months of system development behind us before we had to make the final decision on scheduling "no-AO time" for K2. This should reduce the risks since we would be making decisions based on better information, and also the extra 2 months before K2 I&T could be used to further reduce risk by allowing us more time for lab and K1 sky testing. The timelines associated with these two options are shown graphically in the figures below.

These options, as well as others, are being considered. We will have a recommendation by DDR.

Timeline as Presented at PDR

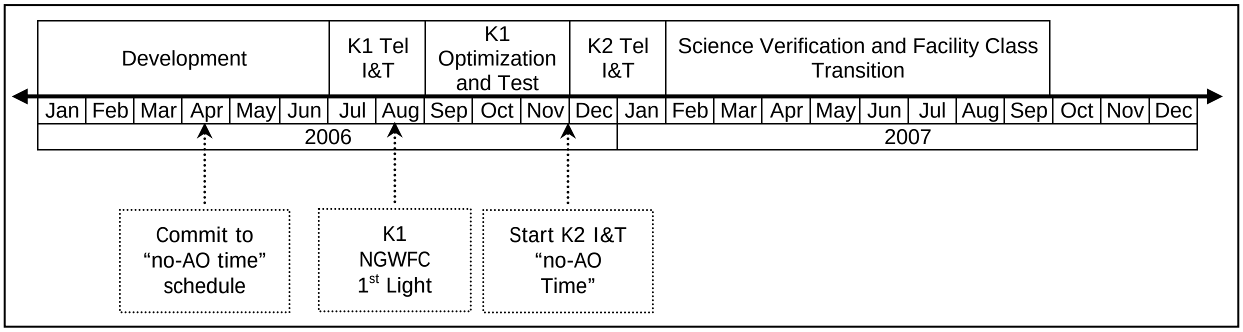

Timeline Under Consideration

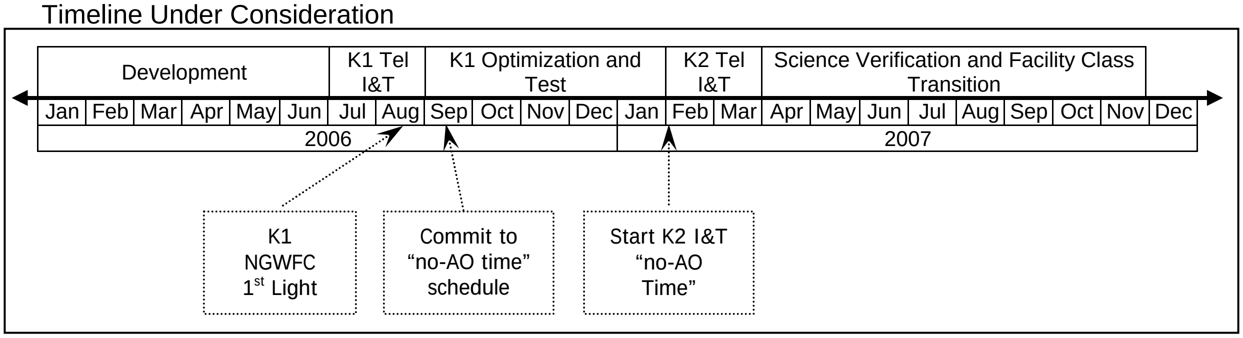

9.3 Provide sufficient resources for "regression testing"

In the detailed design phase we will develop a test plan that is comprehensive and repeatable. Our general approach will be to have a plan to verify each and every requirement. Some of these requirements will need to be verified by tests. Test procedures will be written and archived in the Sybase data base already in use here at Keck for engineering tests. Test results will also be archived with that tool. Sufficient resources will be provided to support this testing.

10 Responses to points made in the *"Other Comments/Suggestions"* section

If the wavefront sensor is aligned in the lab, it should be formally reevaluated at the summit, or the alignment itself should occur at ambient temperature (e.g. at the summit)

Remember that the current AOA chips/boards are floating—use care when moving.

Remember also that the chip on one of the AOA cameras is tilted, and that the lenslet array was tilted to compensate.

If attempting to measure a wavefront sensor flat-field, note that you can't use the pupil imaging mode of the WFS; aberrations (mostly from the Na notch filter) show up as intensity variations in the pupil.

> Another known problem in the existing WFS alignment is that the field stop is out of focus with respect to the star image when the beam is collimated.

The development team would like to thank the committee for these specific and very helpful comments and suggestions. These suggestions will definitely be taken into consideration.

11 Conclusions

We would like to thank the review committee members, Rich Dekaney, Al Conrad, Chris Neyman, Scott Acton, Dave Palmer and Bruce Macintosh for their time and effort in making the NGWFC PDR a very productive review. We would also like to thank the other attendees of the review both at Keck and our colleagues who participated remotely via polycom.

12 Appendices

This section includes the new KAONs written in response to issues brought up in the NGWFC PDR Committee Report and referenced in the body of this document.

Keck Adaptive Optics Note #355

Characterization of the tip-tilt mirror

Marcos van Dam

19 September 2005

1 Introduction

Some concern was expressed at the Preliminary Design Review about the lack of justification for the need for closed loop tip-tilt mirror positioning using strain gauge read back. This report shows why this feature is needed, at least for the downlink tip-tilt mirror.

2 Response of the Keck II tip-tilt mirror

The response of the Keck II tip-tilt mirror was measured using a spectrum analyzer before it was shipped to Keck. The response is displayed in Figure 1. The accompanying documentation states that the first resonance occurs at 700 Hz with a large resonance at 1 kHz, thereby supporting closed loop bandwidths as high as 100 Hz. [1] In addition, there appears to be a phase lag at around 120 Hz, further limiting how hard the mirror can be driven. We would like to obtain closed loop bandwidths higher than 100 Hz, and this does not appear to be possible with the current mirrors driven in open loop (i.e., without closed-loop mirror positioning)

The closed-loop response transfer function was measured using the WFC and STRAP. The tip-tilt loop was closed on the wavefront sensor noise, with a negligible amount of turbulence on the AO bench. It was found that the measured power spectrum from the diagnostics of the residual centroids behaved exactly as one would expect when the frame rate was less than about 200 Hz, as can be seen in Figure 2. However, at higher frame rates, there is a large overshoot, which is almost certainly due to the mirror. The response for a frame rate of 660 Hz and gain of 0.5 is shown in Figures 3 and 4. A gain of 0.6 led to instability in the mirror. This instability does not occur on Keck I, which leads one to believe that it is not a latency issue, since the software and computers are identical. The response on Keck I with a gain of 0.6 is displayed in Figure 5. The results from STRAP, which are not presented here, agree with the results using the WFC. Since the latency for STRAP is negligible, it leads one to believe that the instability is due to the tip-tilt mirror.

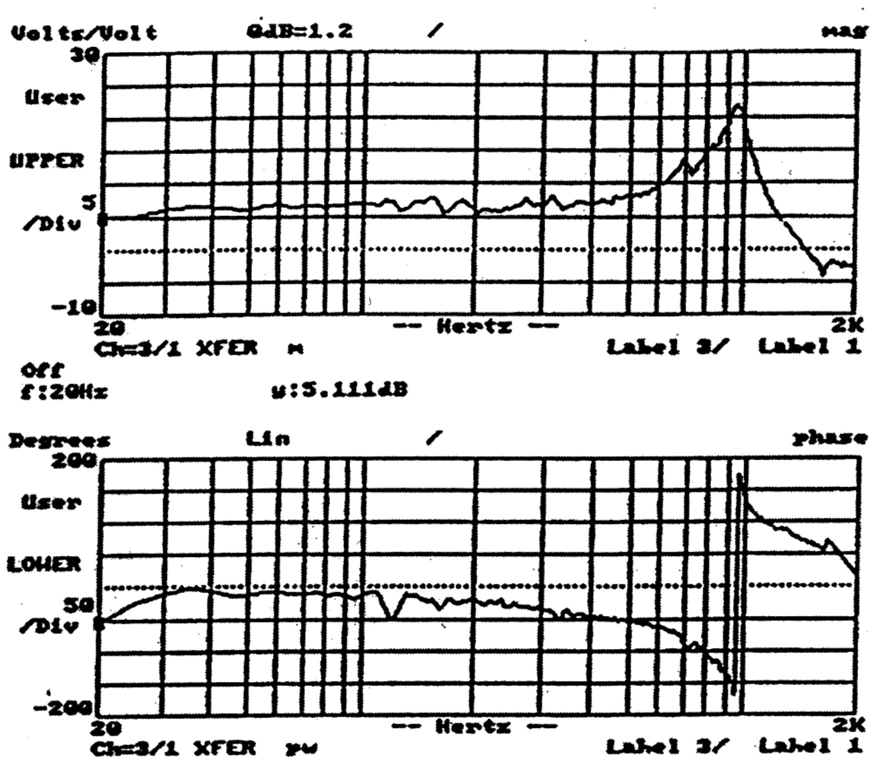

Figure 1 – Measured dynamic response of the Keck II tip-tilt mirror. [1]

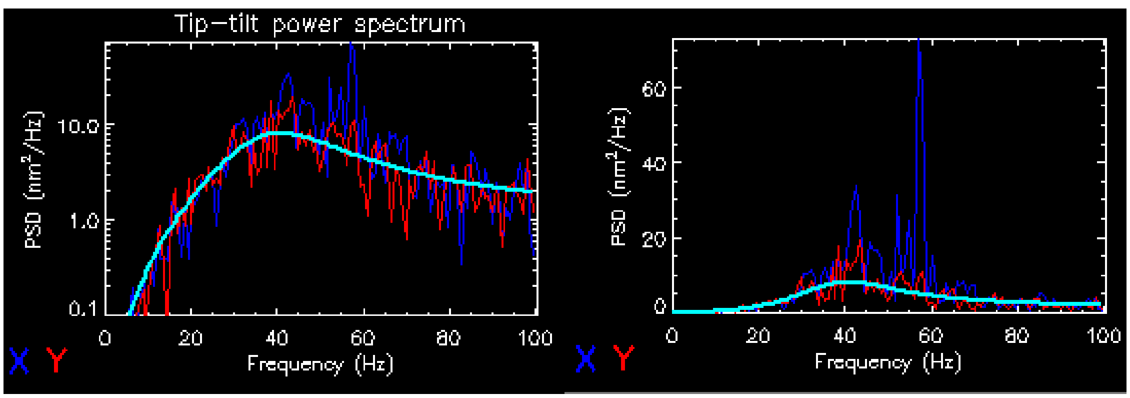

Figure 2 – Tip-tilt power spectrum on Keck II. The frame rate was 200 Hz and the tip-tilt gain was 0.8. The light blue line is the modeled power spectrum. Note the good agreement between the curves.

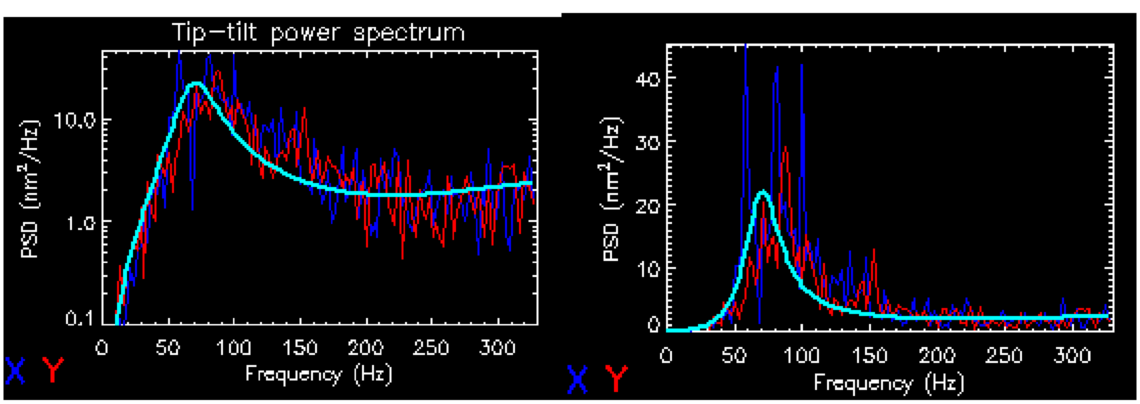

Figure 3 – Tip-tilt power spectrum on Keck II. The frame rate was 660 Hz and the tip-tilt gain was 0.5. The modeled power spectrum assumes tip-tilt mirror dynamics.

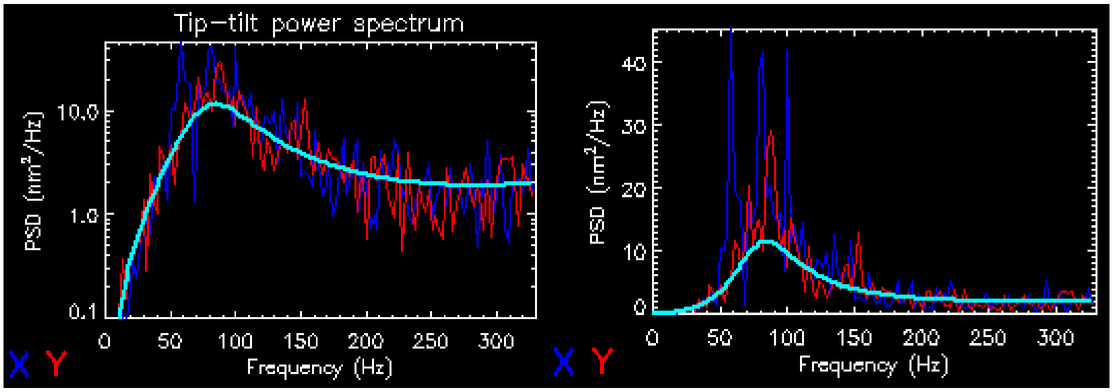

Figure 4 – Tip-tilt power spectrum on Keck II. The frame rate was 660 Hz and the tip-tilt gain was 0.5. The modeled power spectrum does not assume tip-tilt mirror dynamics.

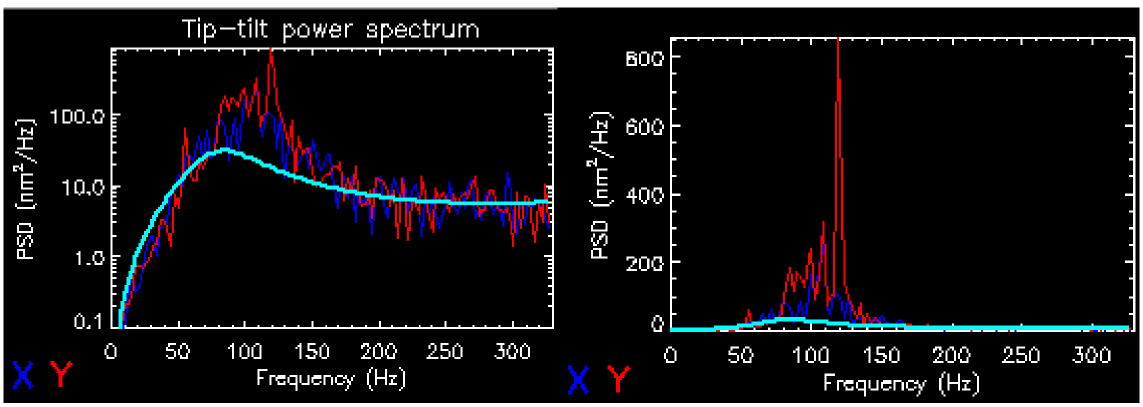

Figure 5 – Tip-tilt power spectrum on Keck I. The frame rate was 660 Hz and the tip-tilt gain was 0.6. The light blue line is the modeled power spectrum if the mirror has no dynamics.

3 Conclusion

The down-link tip-tilt mirrors on both the Keck I and Keck II have dynamics that prevent the mirror from being driven as hard as the control loop allows. Hence, they are currently the limiting factor in the maximum attainable tip-tilt loop bandwidth, especially for Keck II. Having a closed-loop positioning system would largely eliminate the effect of the tip-tilt dynamics, as well as the effect of hysteresis, by dampening the overshoot and ensuring that the transfer function of the tip-tilt mirror is closer to unity. The effect of the closed-loop mirror positioning on the up-link tip-tilt mirrors is not yet well-understood.

4 References

[1] Keck AO Wavefront Control – Hardware Manual.

Keck Adaptive Optics Note #357

Simulations of Denominator-Free Centroiding

Marcos van Dam

04 October 2005

1 Introduction

Denominator-free centroiding has been implemented on a handful of adaptive optics system to reduce the measurement noise when the noise on the wavefront sensor is high [1]. This report investigates how the performance of the NGWFC in NGS AO mode will benefit from this feature via end-to-end simulations.

The standard quad cell centroid formula takes the form:

$$s_x = \frac{I_1 + I_2 - I_3 - I_4}{I_1 + I_2 + I_3 + I_4},$$

where s_x is the x-centroid I_n are the pixel intensities. The denominator-free centroid, on the other hand, replaces the denominator by the expected value of the denominator

$$s_x = \frac{I_1 + I_2 - I_3 - I_4}{E\left[I_1 + I_2 + I_3 + I_4\right]}.$$

This has the effect of reducing the variance of the estimate by eliminating the fluctuations in the denominator. It also has the positive effect of biasing the estimates towards zero when the measured light intensity on a given subaperture is low.

2 Simulations

The simulations were idealized by not including effects such as calibration error, telescope vibrations and segment errors, etc. The simulations were run with an r_0 of 0.15 m and a wind speed of 10 m/s. The CCD parameters used assume the CCD 39 with binned quad cells. The frame rate was kept constant at 100 Hz for all the simulations to see the effect of the denominator-free centroiding. In each case, 200 iterations were run after convergence. This is a very small simulation, so the numbers are only valid as a comparison. In each case, the loop gains were roughly optimized. The true gain, which is the product of the nominal gain and the centroid gain, is quoted. The results of the simulations are displayed in Table 1.

Magnitude	Photons/subap	True Gain	Wavefront Error (nm)	
			Without DFC	With DFC
14.0	38.56	0.6	306.5	305.9
14.5	24.33	0.5	375.8	371.9
15.0	15.33	0.4	477.8	438.0

Table 1: Simulation results comparing the wavefront error with and without denominator-free centroiding (DFC).

The denominator-free centroiding results in a significant improvement when the number of photons per subaperture is less than about 20. The difference in the effectiveness between the two methods is expected to be reduced in practice due to the following reasons:

The denominator used will not be perfect in practice because the pupil is rotating.
When there are non-zero offset centroids, which was not the case for this simulation, the denominator-free centroid produces biased estimates.

In addition, the difference in the error can be reduced by lowering the frame rate when the guide stars are faint.

3 Conclusion

An implementation of the denominator-free centroiding was simulated and was found to reduce the wavefront error significantly for the case where there are fewer than 20 photons per subaperture. The benefit is such that this feature should be an upgrade to the system after the NGWFC is commissioned, and the hooks should be left in place.

Denominator-free centroiding would be even more beneficial on STRAP, where the problems with centroid offsets and changing subaperture flux levels are not an issue and having fewer than 20 counts per readout is common. This possibility should be investigated.

4 References

[1] J. C. Shelton, Denominator-free control algorithms for adaptive optics, SPIE 3126 (1997).

Keck Adaptive Optics Note #358

Reading the subapertures behind the central obscuration

Marcos van Dam

04 October 2005

1 Introduction

Currently, the region on the CCD chip where the telescope's central obscuration falls is read but not used. The subapertures that are read out are marked with a "1" in Figure 1.

```
0 0 0 0 0 0 0 1 1 1 1 1 1 0 0 0 0 0 0 0
0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0
0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0
0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0
0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0
0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0
0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0
1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 0 0 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 0 0 0 0 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 0 0 0 0 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 0 0 1 1 1 1 1 1 1 1 1
1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0
0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0
0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0
0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0
0 0 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0
0 0 0 0 0 1 1 1 1 1 1 1 1 1 1 0 0 0 0 0
0 0 0 0 0 0 0 1 1 1 1 1 1 0 0 0 0 0 0 0
```

Figure 1: Location of subapertures in the CCD

It was suggested by the Conceptual Design Review committee that the 12 obscured subapertures could improve calibrations. The benefits, implementation and additional work required are outlined in this document.

2 Benefits

There are several benefits to calibration of having the 12 additional subapertures present. The main benefit would be the image sharpening algorithm, since the image sharpening could take place in closed loop.

Currently, the loops can only be closed on 240 subapertures, which means that it is impossible to drive all the actuators to their correct location since many of them are slaved. Having 304 subapertures would greatly alleviate this situation, but several actuators in the central obscuration would continue to be slaved. Another potential use of the obscured subapertures lies in the DM-to-lenslet registration. Here, a waffle pattern is placed on the DM and the lenslets are shifted until the amount of waffle seen by the WFS is at a minimum. By using more subapertures, a more accurate measurement could be made. However, the DM-to-lenslet registration is repeatable at the 1% level, which is more than sufficiently accurate. There might also be a small benefit in the alignment routines.

3 Implementation

There are two possible ways of implementing the obscured subapertures: in the real time code or using telemetry of raw pixel data.

The real time code would require an additional 12 subapertures, or an increase of 4% in the number of subapertures. From the perspective of the real time software and hardware, this does not appear to be a problem. The telemetry stream would then contain 316 subapertures, 632 centroids, etc. There would be a modest amount of work involved in changing the size of the keywords, the scripts that display the intensities, centroids, etc. All the calibration routines would have to be modified to make use of all 316 subapertures.

Alternatively, the pixel data could be obtained from the telemetry server and the centroids obtained in IDL. The calibration routines would then have to be modified to use these centroids rather than the centroid values from the keyword. The DM loop could be closed using an IDL routine that uses all 316 subapertures instead of using the real time controller.

4 Conclusion

If the AO system was designed from scratch, the centrally obscured subapertures should probably be retained. However, the amount of work needed to retrofit this into the existing code is not warranted. A partial solution, which will improve the wavefront sensor calibration, is to close the DM loop in IDL by deriving the centroids in IDL from pixel data in the telemetry. This could be implemented later if necessary.

Keck Adaptive Optics Note #359

Wavefront sensor observing modes

Marcos van Dam

10 October 2005

1 Introduction

There are many possible operating modes for each of the two NGWFC wavefront sensing CCDs. At the preliminary design review, concern was expressed that there are too many modes and that the number of lenslet/CCD readout combinations should be reduced by only considering a subset of these. Reducing the number of modes will eliminate a substantial amount of work in the user tools and will simplify operations. This KAON compares the different configurations and suggests which modes should be used

2 CCD-39

There are three possible operating modes for the CCD-39, summarized in Table 1.

Mode	Binning	Binned CCD Size	Subap. Size (after binning)	Spacing (pixels)	Max frame rate (Hz)	Max data rate MB/s
1	1	80 x 80	2 x 2	4	2000	24.4
2	2	40 x 40	2 x 2	2	2000	6.1
3	1	80 x 80	4 x 4	4	2000	24.4

Table 1: Operating modes for the CCD-39 [1]

In deciding which configuration to use, there are four issues to be considered: read noise, dark current, charge diffusion and linearity. Table 2 displays the different read modes developed by SciMeasure in conjunction with JPL for the CCD-39. The full table provided by JPL is Table 8 in the appendix.

It can be seen that binning the pixels allows one to read the camera more slowly, thereby reducing the read noise. However, there is no way to read just the central 2x2 pixels of a 4x4 pixel subaperture, so the readout rate for Mode 1 does not take advantage of the reduction in pixels.

Program	Pixel Rate	Image Size	Binning	Filter	Filter TC	Max Frame Rate	Read noise
Units	kHz/port				Ns	Hz	e-
0	203	40	2	4	1100	455	3.7
1	382	40	2	3	690	832	4.4
2	735	40	2	2	48	1520	5.8
3	1136	40	2	2	48	2219	6.4
4	204	80	1	4	1100	129	3.5
5	455	80	1	3	690	283	4.5
6	1136	80	1	2	48	699	5.7
7	2038	80	1	1	30	1226	6.6

Table 2: The readout modes for the CCD-39 [2].

A comparison of the noise performance of the two quad cell modes is summarized in Table 3. The spot size quoted assumes the seeing is 0.6" and that the pixel size is 1.0" for the unbinned case and 0.6" for the binned case. The charge diffusion is 0.6 pixels FWHM and the dark current measured on the Palomar system is 200e-/pixel/sec at -40 degrees Celsius [3], which differs slightly from Eq. (1). It can be seen that binned pixels results in better performance for all stars. KAON 284 shows that quad cells are better than 4x4 centroids for all stars except the very brightest. However, if the plate scales are designed for binned pixels they will also be usable in 4x4 mode. It is recommended that only binned quad cells be used, as is the case with the JPL WFC at Palomar.

Bin	Mag	Frame rate	Photons	Read noise	Dark current	Pixel noise	Centroid noise	Spot size	Absolute noise
1	11	455	134.3	5.7	0.44	12.93	0.024	0.85	0.0109
2	11	455	134.3	3.5	1.76	12.18	0.023	0.70	0.0084
1	13	200	48.4	4.5	1.00	8.35	0.043	0.85	0.0195
2	13	200	48.4	3.5	4.00	8.04	0.042	0.70	0.0155
1	14	100	38.6	3.5	2.00	7.27	0.047	0.85	0.0213
2	14	100	38.6	3.5	8.00	7.67	0.050	0.70	0.0185
1	15	50	30.7	3.5	4.00	6.85	0.056	0.85	0.0252
2	15	50	30.7	3.5	16.00	7.68	0.063	0.70	0.0233

Table 3: Comparison of the noise performance for the CCD-39 quad cell modes

Another issue of interest is the temperature to which the CCD chip should be cooled. The dark current, *D*, in electrons/pixel/second, was calculated using the formula in the EEV documentation [4]:

$$D = 75000 \times 122 T^3 \exp[-6400/T] \quad (1)$$

For the case of binned quad cells, the centroid error was calculated as a function of temperature for two examples: a magnitude 14 star with the WFS operating at 100 Hz and a magnitude 15 star at 50 Hz. The results are displayed in Table 4 and show that the chip needs to be cooled to a temperature of -30 ºC or lower, with diminishing returns below -40 ºC.

Temperature (ºC)	Centroid error	
	V=14, 100 Hz	V=15, 50 Hz
0	0.1508	0.2606
-10	0.0985	0.1634
-20	0.0686	0.1047
-30	0.0542	0.0734
-40	0.0486	0.0598
-50	0.0468	0.0551

Table 4: Centroid error as a function of CCD temperature

Given that the 2x2 binning mode will be used, end-to-end simulations were run to determine the optimum plate scale. In each case, the loop gains were optimized using the algorithm that is used on the sky by the gain optimization tool. 1000 iterations were run after convergence. The results are summarized in Table 5. Note that the plate scale quoted is the total extent of the two binned pixels. Note that the results using plate scales higher than 2.06” require an artificially long wavelength (1.4 microns) in order to Nyquist sample the spots on the WFS. Since D/r0 > 1 in each subaperture, using longer wavelengths is probably acceptable.

Mag	Frame rate	R0	Plate scale (”)	Error (nm)
7	2000	0.10	1.03	285.2
7	2000	0.10	1.55	258.5
7	2000	0.10	2.06	246.9
7	2000	0.15	1.03	195.9
7	2000	0.15	1.55	180.8
7	2000	0.15	2.06	175.0
11	500	0.15	1.03	223.0
11	500	0.15	1.55	214.7
11	500	0.15	2.06	212.6
13	200	0.10	1.03	352.3
13	200	0.10	1.55	339.6
13	200	0.10	2.06	344.0
13	200	0.15	1.03	247.2
13	200	0.15	1.55	248.2
13	200	0.15	2.06	257.5
14	100	0.15	1.03	294.2
14	100	0.15	1.55	300.4
14	100	0.15	2.06	312.5
14	100	0.15	3.09	352.6

Table 5: Simulated performance as a function of plate scale for the CCD-39. The plate scale is the total extent of the two binned pixels.

The simulations were simplified by neglecting effects such as telescope vibrations, calibrations errors, primary mirror aberrations, etc. It is interesting to note that the bigger the plate scale, the better the performance on bright guide stars. This is probably due to the fact that the charge diffusion acts as an anti-aliasing filter and increases the linearity of the WFS. Based on these results, it is likely that we only need two plate scales for the CCD-39: a 1.5” plate scale to be used with faint targets and a 3.0” plate scale for bright point sources, extended objects and the laser guide star. Hence, the plate scales should be 0.75” and 1.5” per unbinned pixel. The 1.5” plate scale could also be used in the future with 1x1 binning with very faint guide stars to reduce the dark current if this turns out to be the limiting factor.

3 CCID-56

Table 4 displays the six possible modes for the CCID-56. The 8x8 mode will be supported, since it is needed for LGS AO and may be beneficial for NGS AO. The 4x4 configurations will not be supported initially, since quad cells perform better than 4x4 centroiding. The supported quad cell configurations will depend on the noise properties of the CCD and readout electronics.

Mode	Bin	Binned CCD Size	Subap. Size (after binning)	Spacing (pixels)	Max frame rate (Hz)	Max data Rate MB/s
1	1	160 x 160	2 x 2	8	1000	48.8
2	2	80 x 80	2 x 2	4	2000	24.4
3	4	40 x 40	2 x 2	2	2000	6.1
4	1	160 x 160	4 x 4	8	1000	48.8
5	2	80 x 80	4 x 4	4	2000	24.4
6	1	160 x 160	8 x 8	8	1000	48.8

Table 6: Operating modes for the CCID-56

At this stage, it is reasonable to assume that the supported configurations will be 2x2 and 4x4 binned quad cells and the 8x8 mode. For the LGS AO case, we would like the total extent of the WFS to be 6”x6”. This fortuitously results in a 2x2 binned quad cell of 1.5”, which was shown to be the most versatile quad cell size for the CCD-39.

A comparison was made of the 8x8 centroider with the quad cells for a bright guide star. It is found that the 8x8 centroid results in much better performance than the centroider, as can be seen from Table 5. The improvement is thought to stem from the fact that the quad cell is sensitive to changes in the spot shape (as well as size), while the 8x8 centroider is not.

Mag	Frame rate	R0	Plate scale (“)	Error (nm)
7	1000	0.10	0.257	231.3
7	1000	0.10	0.387	226.9
7	1000	0.10	0.516	227.9
7	1000	0.10	1.547	244.6
7	1000	0.10	3.090*	236.8
7	1000	0.15	0.257	163.8
7	1000	0.15	0.387	162.8
7	1000	0.15	0.516	164.5
7	1000	0.15	1.547	172.9
7	1000	0.15	3.090*	168.2
11	500	0.15	0.257	190.3
11	500	0.15	0.387	196.1
11	500	0.15	0.516	204.3
11	500	0.15	1.547	181.4
11	500	0.15	3.090*	184.0

Table 7: Simulated performance as a function of plate scale for the CCID-56. The plate scale is the total extent of the two binned pixels. *Indicates that the plate scale was set to be effectively 3” by unbinning the pixels, but the total extent was still 1.5”.

Note that, since the frame rate between the magnitude 7 and 11 stars differed, the turbulence used in the simulations differed. Hence, the results between the different magnitudes are not directly comparable.

The optimal pixel size for 8x8 centroiding would appear to be roughly independent of the plate scale from Table 5. However, the linearity of the WFS is much better when the spot is well sampled. For this reason, and in order to be able to calibrate using a diffraction-limited source, I recommend that the plate scale be 0.25".

4 Conclusion

Only one mode for the CCD-39 will be used: the binned quad cell. Two plate scales are needed to support all observing modes: 0.75" and 1.50" per unbinned pixel. Since the WFS spot is not well sampled, there is no benefit to having a smaller plate scale for the Keck I LGS AO system.

The CCID-56 will support two modes: 8x8 centroiding and 2x2 binned pixels. 8x8 centroiding will be used for LGS AO and for bright natural guide stars. The two plate scales required are 0.25" and 0.75". On Keck I, an additional plate scale of 0.55" would be useful for LGS AO.

The 0.75" plate scale for the CCD-39 can probably be reused for the CCID-56.

5 References

[1] KAON 289: System Design Manual
[2] Mitchell Troy, JPL
[3] CCD39-01 Back Illuminated High Perforamnce CCD Sensor product sheet.
[4] Mitchell Troy, "Marconi CCD39 Dark Current" (2003)

Program	Pixel Rate		Image Size	Binning	Clamp	Sample	Filter	Filter TC	Max Frame Rate		Read Noise	
Units	kHz/port							ns	Hz		e-	
	Requested	**Actual**							**Requested**	**Actual**	**Requested**	**Actual**
0	178	203	32	2	50	0	4	1100	261	455	3.5	3.7
1	400	382	32	2	70	130	3	690	560	832	4.7	4.4
2	1140	735	32	2	0	115	2	48	1388	1520	5.9	5.8
3	2000	1136	32	2	20	150	2	48	2114	2219	7.5	6.4
4	178	204	64	1	100	100	4	1100	132	129	3.5	3.5
5	400	455	64	1	150	0	3	690	288	283	4.7	4.5
6	1140	1136	64	1	30	10	2	48	743	699	5.9	5.7
7	2000	2038	64	1	75	5	1	30	1176	1226	7.5	6.6

Table 8: Full table for the reading modes for the CCD-39 provided by Mitch Troy, JPL