MONSOON Software PDR Report

July 25, 2003

Presented to David Sprayberry (NOAO, Head of Engineering)

Executive Summary

Monsoon is an extremely ambitious and visible hardware/software project that is critical to many of NOAO's future aspirations. If successful, Monsoon will advance array control in ways that have only been a dream before, and could have tremendous impact far beyond NOAO. To fulfill that hope, though, it deserves attention at the highest levels within NOAO.

The panel was very impressed that the hardware and software are already working together in a lab system with so few personnel involved. We like the fact that the software is being developed in an open fashion with a good approach to documentation on an open source operating system.

However, the Monsoon team is resource-starved, and consequently, appears to be designing and implementing to meet the demands of immediate projects while sacrificing long-term design goals. The Monsoon team needs to develop a design that encompasses the broader, larger goals of Monsoon. In particular, the supervisory and DHS layers need to be integrated more thoroughly into the software designs to address long-term requirements.

FIRST RECOMMENDATION: Project Management -- A project manager is essential to guide the development of this ambitious project. The project needs a credible short- and long-term schedule and plan to succeed. The project manager, having developed a management plan, will likely find that ~2 additional software personnel are necessary to meet those schedules. If NOAO is truly serious about the long-term goals for Monsoon, both within NOAO and for outside observatories, a greater level of commitment will be required for both the software and hardware components, including more internal and external scientist and instrument builder involvement.

SECOND RECOMMENDATION: Integration of the Design -- The central core of the Monsoon software is being handled well, but the surrounding layers do not appear to be integrated. The supervisory, data handling, and microcode layers need to be developed more fully at the conceptual level. The panel did not see that these areas were being addressed by the design. Until the entire software chain is integrated, the Monsoon controller system will require a huge effort to bring to every new application.

THIRD RECOMMENDATION: Requirements and testing -- A number of the requirements need more detailed definition, and in many cases, a set of metrics for a test plan to compare against. A larger Monsoon team than is currently available will likely be needed to expand the requirements, define metrics, and fulfill the testing plan on the current schedule for short-term needs.

FOURTH RECOMMENDATION: Software practices -- Long-term development costs can be reduced by adopting existing high level software interfaces and modern concepts (e.g., object oriented styles).

Caveat: There is confusion over what constitutes the Monsoon project. It is referred to as a "pixel server", which is a more limited concept than a full-up astronomical array controller. The ultimate goal is to "control" arrays of detectors, and so, the panel views the Monsoon project as an "array controller", complete with the hooks, interfaces, and software to provide full functionality. Consequently, some of our concerns lie with the restricted scope of the project rather than the design for that scope. We feel that NOAO is making a serious strategic error by de-scoping Monsoon to the level of available resources rather than allocating the resources needed to produce a fully functional array controller system.

Overview

A. Charge to the Committee

1. Are the requirements as spelled out:
   - complete enough to address foreseeable needs and risks?
   - clear enough to guide the software development?
2. Does the proposed design adequately address the requirements?
3. Is the testing procedure adequate?
4. Are the resources and schedule for producing the software sufficient and credible?
5. Are there any risk areas in the design, testing plan, resources, or schedule that have not been adequately addressed?

B. Committee Membership

Bret Goodrich (NSO/SOLIS)
George Jacoby (Panel Chair, WIYN Observatory)
Bob Kibrick (Lick Observatory)
German Schumacher (NOAO/SOAR)
Chris Smith (NOAO/CTIO)

C. Meeting logistics

Meeting Date: July 1, 2003

Presentations were held at NOAO/Tucson. Two of the panelists were on-site while video conference connections were made to Chris Smith and German Schumacher at NOAO/CTIO and to Bob Kibrick at Lick Observatory.

Introductory Comments

The Monsoon project being carried out by NOAO represents an ambitious controller package that is intimately tied to many of NOAO's future large focal plane ventures (e.g., NEWFIRM, LSST, GSMT ODI). In addition, Monsoon is expected to replace some of NOAO's aging array electronics (Arcons, 2901s, Wildfire). The project is vital to NOAO's future and will be highly visible in the broader community. There clearly is a need within astronomy for the fast, multi-channel, extensible hardware/software capabilities that Monsoon hopes to realize, and the panel strongly endorses this effort and recognizes it as an appropriate project for the national center. The bar for success has been set far higher than for any prior controller system, and yet, if the goals can be met, Monsoon will have far-reaching impact beyond NOAO.

Most of the software presentations at the meeting were given by Nick Buchholz. In addition, Phil Daly discussed the library structure and Peter Moore briefly described the Monsoon hardware. The committee wishes to commend all the presenters for providing very clear descriptions of their work and for responding to our questions with equal clarity while being concise. The software component of Monsoon is a very daunting and challenging task; we recognize that there are very few software engineers assigned to the project, and the results to date are remarkable and impressive given that limitation.

As a result of the excellent progress in some areas, several members of the panel felt that the presentations demonstrated a level of detail and development far beyond that appropriate for a preliminary design review. In fact, a prototype version of Monsoon is operating. Re-coding on the basis

of the panel's comments may no longer be an option considering the short-term demands for Monsoon controllers within NOAO, and some panelists felt constrained in their comments by the advanced state of certain system components. On the other hand, aspects of the software design have not yet reached the conceptual level (e.g., DHS, MLS), let alone to the point of a preliminary design. Thus, the panel report may be of less value to the project than if the work had been organized more uniformly. We believe that this dichotomy is rooted in NOAO's urgency for a short-term solution to its controller needs for Orion testing and NEWFIRM while dedicating an understaffed work force to the long-term ambitions for Monsoon.

Detailed Response to the Charges

Charge 1. Are the requirements as spelled out:

- complete enough to address foreseeable needs and risks?
- clear enough to guide the software development ?

The requirements are very extensive, reflecting the attempt for Monsoon to address the needs of many instrumentation projects. In some cases, the statements are overly general (e.g., "Provide efficient science operation support", "Support efficient boot-up and initialization") when a quantitative metric should be assigned. Those requirements that are general or nebulous should be recast by a small group of instrument scientists and observers to be much more specific and to identify missing requirements. The inability to change exposure times during the middle of an exposure is an example of where science operation efficiency was lost because the requirement was unclear.

We recommend that the Monsoon team draft a set of metrics to attach to the requirements. Because some requirements may compromise or compete with others (e.g., security vs remote observing access), we suggest that the team also assign priorities to the requirements.

An integral part of requirement definitions is a way to validate, or test, that the metric was achieved. That is, a test plan for each requirement must be clear. In some cases, these are self-evident. In others, the test plan should be part of the re-stated requirement.

For NEWFIRM, which is on NOAO's immediate horizon, the requirements are likely to be more clearly enumerated than any other application. Nevertheless, some gaps were identified. Within the FPRD, we note the following specific items:

2.4.2.1 Detector Configuration Database Access and Management
2.4.2.2 Exposure Configuration Database Access and Management

No requirements were specified for these elements

2.4.3.1 Device selection (1 to N, for mosaics)
2.4.3.2 Mosaics handled as a single focal plane

No requirements specified.

2.4.4.1 Integration time selection

What is the requirement for the resolution to which the integration time must be specified (e.g., milliseconds)?

Is there a requirement that one must be able to change the integration time of an in-progress integration without first pausing and resuming the integration?

Is there a requirement that the mechanism by which the integration time is measured takes account of (and corrects for) any latencies in the actual opening and closing of the shutter?

Is there a requirement to support different integration times for different detectors that compose a single mosaic? For example, a spectrograph illuminates a mosaic such that light from the blue end of the spectrum falls on one detector while that for the red end falls on another.

For a given 1-hour exposure on a science object, it may be desirable to read out the red detector of the mosaic several times (e.g., two 30-minute exposures or three 20-minute exposures) while reading out the blue detector only once at the end of the 1-hour exposure. This involves pausing the exposure on the blue detector of the mosaic during the periods when the shutter is closed and the red detector is red out, and it involves keeping track of multiple exposure times within the context of a single mosaic potentially operated by a single DHE and PAN. If Monsoon is truly intended to address the foreseeable needs of instruments over the next decade, it should consider unusual requirements such as these, since some facilities are already faced with such requirements.

2.4.4.3 Device output selection

No requirements are specified for this element.

What are the requirements for selecting a readout of a subset of channels?

What are the requirements for coping with situations where one or more of these readout channels is inoperative?

This is a potentially complex area given the range of detectors and mosaics that Monsoon is intended to support. A clear elaboration of these requirements is vital.

2.4.4.4 Binning

This is another potentially complex area given the interactions between binning, non-imaging pixels (e.g., prescan pixels on CCDs), and regions of interest. Is there a requirement to allow different detectors of a given mosaic to be read out using different levels of binning, or is a single binning factor imposed for the entire mosaic?

2.4.4.7 External Exposure Mode support

While 10 such modes are listed, no requirements are specified for any of them.

2.4.4.9 Region of Interest Definition

This specifies "Multiple Region of Interest Support". But what is the requirement? What is the minimum number of multiple regions that must be supported? Is there a requirement that there be no restriction on the overlap between these regions, or is such a restriction permitted?

This can get quite involved for mosaics. The requirements here need to be elaborated.

2.4.4.10 Charge Shifting - Orthogonal transfer imaging

No requirements were provided. These need to be fleshed out.

2.4.4.11 Exposure Control

Is there any requirement for providing a parameter to control the number of erase cycles that occur at the start of an exposure, and if so, what are the required minimum and maximum values for this parameter?

Is there a requirement to support the "Abort" operation during either the erase cycle or the readout cycle, or is there only a requirement to support "Abort" during integration?

What is the requirement with respect to supporting multiple exposures as a single image (i.e., erase detector, repeat M times<open shutter, integrate N seconds, close shutter but don't read out detector, change instrument and/or telescope settings>, open shutter, integrate N seconds, close shutter and readout detector)?

What is the requirement for providing to upper-levels of software (e.g., instrument and/or observatory control layer) asynchronous notification of state changes in the exposure cycle so that those upper layers can synchronize their activities to the various changes in exposure state (e.g, to take a snapshot of UT, ST, RA, DEC, HA at the moment that the shutter opens at the start of an exposure)? Specifically, what is the requirement as to which state changes will provide such asynchronous notification capability? Are there any requirements on the timeliness of such asynchronous notification, or a requirement that such notifications be time stamped?

2.4.6 Shutter Mechanism Control

Is there a requirement to monitor and report the actual state of the shutter (i.e., did it actually open fully when commanded to do so, did it fully close when commanded, did it spontaneously open or close without being commanded to change state, etc.)? Is there a requirement to monitor and report the state of the individual blades of a dual-bladed shutter?

2.4.8 Detector Monitoring, Maintenance and Protection

Is there a requirement to provide control of the detector temperature? (Section 2.4.12.1 hints that there is, but does not elaborate). If not, what entity is responsible for maintaining the detector temperature, and how is this integrated/cross-checked with the temperature monitoring functionality provided by Monsoon?

Mention is made of provision of power controls for the Detector bias voltages. Is there a similar requirement for power controls for the Detector clock voltages? (For example, for some detectors it necessary to power up various bias voltages (e.g., Vopg and Vref) in a specific order relative to the powering up of certain clock voltages (e.g., reset gate).

2.4.11.1 Descrambling

No requirements were provided. Again, this can become quite involved, and has interactions with the requirements regarding detector and detector output selection, binning, etc.

2.4.11.2 Centroiding

No requirements were provided. This is a potentially large area to cover, given the variety of algorithms possible for various types of guiding.

2.4.11.3 Data Scaling

No requirements were provided.

2.4.11.4 Data Processing Modes

Although modes were identified, the requirements for each modes were not provided.

2.4.11.5 FITS Packaging

No requirements were provided. This is a potentially huge area (especially where mosaics are involved) and the requirements need to be specified at two levels:

1. The minimal FITS packaging required to support detector/ instrument development in the lab (e.g., integration time, binning, readout gain)
2. The maximal FITS packaging required to support science observations on the telescope (i.e., images written by DHS).

While the latter imposes requirements perhaps beyond the scope of the Monsoon software itself (assuming the DHS is an entity that is external to Monsoon), those requirements are relevant to Monsoon in that Monsoon must demonstrate that it provides sufficient hooks to the DHS so that it can carry out its job. It is vital that science images written to disk as FITS files contain sufficient header data to enable science, to enable automated data pipelines, and to provide an archival record.

Will multi-DHE images be a single FITS image? If so, what happens when the file sizes grow beyond the 32-bit limit of 4 Gbytes? If not, how will the different DHE outputs be organized?

2.4.13.1 Clock Driver Requirements

In addition to the items listed, is there an explicit requirement to support clock shaping (e.g., sloped clocks)?

Is there a requirement to continue running serial clocks during integration?

Is there a requirement to support anti-blooming parallel clocks during integration?

Is there a requirement to support sequencing of bias voltages relative to clock voltages?

2.5 Major System Conditions
2.7 User characteristics
2.8 Assumptions and Dependencies
2.9 Operational Scenarios

No requirements were specified for any of these elements. In particular, it would have been extremely helpful to have some examples of various operational scenarios envisioned for the Monsoon component of the NEWFIRM instrument.

Charge 2. Does the proposed design adequately address the requirements?

Given that many of the requirements are not stated with specific metrics, the panel found aspects of this charge difficult to evaluate. No where was this seen more unanimously than with the DHS issue. To quote one panelist, "It is the responsibility of the MONSOON project to design and deliver a DHS interface."

The panel was strongly in favor of having a design that was responsive to requirements for the Supervisory layer, multi-PAN systems, and multi-DHE systems. Without these requirements and personnel to develop a design, the current status is at the pre-PDR level. Similarly, the interactions, the division of functionality, and the interfaces between the MSL, DHS, and multiple PANs are not well defined. The kinds of questions that we could not answer are, for example:

- How does the DHS know how many PAN data streams to expect?

- How does the DHS communicate that it needs a re-transmission of data from one PAN?
- Will there be any tools for developing micro-code or will detector engineers be required to use VHDL?
- Could the MSL inadvertently start another exposure if there has been a problem in the DHS?

On the other hand, for those requirements that were specified with enough detail, the design does appear to be adequate. While this may be sufficient for the short-term goals of Monsoon (e.g., NEWFIRM), the implementation (which takes us beyond the PDR stage) utilizes software engineering practices that may be difficult to maintain over time and may introduce unnecessary costs downstream. Alternatives include, for example:

- communication via higher level protocols than sockets (e.g., RPC, BEEP, ICE, NDDS)
- an object-oriented approach offers advantages for development and maintenance
- text files may prove to be inadequate for logging and a more sophisticated data base is likely to be needed for the highly multiplexed many-DHE systems of the future

---

Charge 3. Is the testing procedure adequate?

For the more immediate projects, the test plan appears to be thought out well. Testing of the DHE and PAN layers in the lab looks fine. There was no well-developed plan for the multi-PAN system with a real MSL and a real DHS at the telescope. There may be issues related to timing and communication errors/contention that aren't evident in the simpler implementation of a single PAN-DHE.

Charge 4. Are the resources and schedule for producing the software sufficient and credible?

The panel was unanimous about this item. The project is seriously compromised with only 1.5 programmers, no project manager, and limited scientific input. In order to meet the needs of the immediate project deadlines (e.g. NEWFIRM), the limited personnel have focused on the short-term and could be sacrificing long-term design goals.

The schedule that was presented was limited in detail and there was no long-term plan discussed.

Both of the above issues are red flags. Our strongest recommendation for improving the likelihood of success of this highly ambitious, highly visible, and NOAO-critical project is to enhance the management team for Monsoon that integrates the hardware and software efforts. A key component of the management team will be adequate scientific input from both IR and optical instrumentalists and future users. The first appointment should be a project manager, and this should be done very soon. As it stands, the project schedule and long-term goals are at risk.

Charge 5. Are there any risk areas in the design, testing plan, resources, or schedule that have not been adequately addressed?

We identified several risk areas that were not addressed at all:

- A great risk, that could be mitigated by the recommendation under 4, is the push to meet short-term demands at the expense of long-term design concepts. For example, a considerable amount of software was implemented prior to the PDR.
- Given that this project is largely a one-person effort, the risk to the schedule is extreme should Nick become unavailable (e.g., illness, depart the organization, be co-opted into another NOAO project).

- A related risk with a single-person operation is that the code is not reviewed by a team and could contain idiosyncrasies that lead to long-term maintenance costs.

It is difficult to comment on risks to the schedule beyond the short-term planning for NEWFIRM. For example, is it planned that a fully-developed Monsoon system will be available to replace the Arcons on the two NOAO Mosaic cameras in N (2-3?) years?