Hardware in the Loop Observing Simulation Platform Development

Stuartt Corder (PI), Tzu Shen (co-PI), Alejandro Saez, Rodrigo Amestica

Executive Summary

The Hardware in the Loop development project is intended to return considerable (~5%) successful observing time annually to the astronomy community through reductions in allocation of production hardware test resources. This is accomplished by providing a simulation platform that attempts to incorporate, in a flexible way, as much of the actual hardware as feasible. In the most realistic case, four actual antennas would be correlated with 12 simulated antennas to provide some real, on-sky data, running on a completely different software/firmware setup, without any impact whatsoever on PI science. It is also a potential source of time saving in the commissioning of major development projects. Finally, it is likely that the facility will result in improved uptimes due to more flexible time available for testing stability or efficiency improvements and troubleshooting problems with specific antennas. We expect the ~800 kUSD (~1% cost of operations) investment in the facility will return a minimum of 5% increase in the time available to science annually and will also provide flexible resource allocation to maximize return on other developments (e.g., upgraded correlator), improvements, and investigations providing further benefit to the community.

Motivation

The Hardware in the Loop project (HiL) concept was born out of the need to satisfy many boundary conditions within ALMA Operations. Chief among these are the conflicting requirements that the ALMA Department of Computing (ADC) and Integrated Computing Team (ICT) utilize as little time on the production hardware environment as possible all the while providing increasing capabilities and improving operational robustness and efficiency.

The HiL project provides flexibility not offered by subarrays, as regardless of subarray distribution, a common software version must be run on all of them. Therefore, to test new software deployments, an entire system must be migrated to the upgraded version. Subarrays are useful for testing new capabilities in their earliest phases where calibration schemes or performance investigations can be done on subsets of antennas. Subarrays have been very useful to share array resource between Engineering activities (troubleshooting, maintenance, verification of single antennas, etc.) and some Extension and Optimization of Capabilities (EOC) activities, but the annual upgrade process still requires considerable activity to be carried on the full production hardware, taking time away from science observations.

Steps have already been taken to minimize the impact of full system software testing. Afternoons are allocated to computing testing, a time during which phase stability is the poorest and wind conditions often restrict the activities at the high site by Engineering teams. It should be noted, however, that compact conditions, science is still usually viable and that in extended conditions, observations at bands 1-4 would often be possible. Other test times, particularly full system tests of the online system candidate for the next cycle, require a broad range of test scenarios, weather conditions and specialized skills, meaning that it is impossible to miss only "poor weather" time.

Currently simulation capabilities at ALMA are restricted in a few critical ways. Firstly, the correlator hardware and firmware are not simulated. Improvements to key software simulation layers have been made, but overall the simulation capabilities are restricted. Antenna simulations are limited and the critical photonics are simulated in a limited way. Because of this, ADC is scheduled for two five-hour windows per week on the production environment. In addition, the release cycle of online software, and specifically the candidate for a given observing cycle, requires considerable dedicated time on the production hardware. To mitigate the impact of this, two antennas are usually removed from operations from January through March to install on the 2-station correlator at the base camp. While providing valuable insight into more fundamental issues in the new release without impacting the full array[1], the limitation of this facility, mostly in correlator and photonic hardware, only catches a fraction of the issues that would have been found in production environment testing.

We anticipate that an upgrade of the simulation platform will dramatically reduce the number of hours required for software testing on the production hardware on an annual basis. Our concept allows for a 16-station single correlator quadrant located at the base camp facilities to interface with an actual photonic system at the high site, providing up to 4 real inputs from actual antennas at the AOS, giving real correlation data to test calibration and correlator software changes. The system running the simulation environment will use as much of the real hardware as is available, without requiring system restarts or halting PI driven science during the testing period.

The platform will also serve as a vital test bed for hardware upgrades that can be partially introduced into the system without production installation. An excellent example of the utility of this facility would be in the commissioning of the upgraded baseline correlator, a project currently under review. Other opportunities to support development also exist. For example, during Band 5 commissioning, an upgrade to the online system was required to change the control software to optimize the bands appropriately during integration at the high site. In the current set up, a patch, which is risky and requires test time, is

[1] This type of testing is critical because it allows real data to be collected on sky, identifying critical issues in control and data formatting. It does not, however, test the correlator in a real way because the 2-station setup is fundamentally different from the baseline correlator.

installed to allow integration to continue in the production environment. In the HiL, the test on actual antennas could have been done without patching the production environment, allowing integration to continue and providing better timing for patching and verification in production. The benefits will come in big pieces and small additions for development projects.

Benefits

Throughout the year, a variety of types of testing utilize the production hardware in an exclusive or disruptive way. We anticipate that production hardware software development testing utilizes, in terms of antenna hours:

1. Two antennas for 1100 hours of nominally PI science allocated time in January and March to do basic, initial (e.g., typically highly disruptive) testing of the next cycle's online software candidate. Given nominal assumptions, this would be 2200 antenna hours.
2. Five hours per week of the entire array (~45 antennas) for 48 weeks a year, given nominal assumptions, this is 9720 antenna hours.
3. It is expected that approximately 50% of the problems found in the production environment testing at the final phases of testing the candidate online release would be discoverable on the HiL facility. Approximately 1.5 nights (12 hours total) every other week is used for this phase (full system) over a period from April to June, at which point the frequency drops to 1.5 nights per four weeks for July and August. This totals 13 nights or 105 hours of full array use or 4700 antenna hours.
4. Additional savings is expected on an antenna-by-antenna basis or in investigation of problems, but this timesaving is difficult to predict. It is likely that this savings would be offset by the number of hours that ADC/ICT would be using the four antennas in terms of testing and we call this a net zero gain when weighed against that value. Although it should be noted that much of the computing test time with the four antennas could be done on "engineering days" which would then not limit the PI science in any way.

All told, the HiL platform will return, directly to the PI, 2.5 weeks of successful, full array observing time per Cycle. Making typical assumptions, this would imply ~30 additional normal project completed. Or another way, 2.5 weeks out of 48 weeks currently scheduled for PI time represents a savings of ~5%. Either way, the modest 1% of annual budget investment for a 5% return on observing time is more than worth the cost for a single year, let alone for a system that will return this year after year. It could be argued that we are returning the bulk of this time during poor weather periods, with the saving from items 2 and 1 above are in afternoons and during the poorer weather period of the year, respectively. While this is true, the 4700 antenna hours saved during April through August represents the saving of high quality observing time, and it should be noted that a considerable fraction of the requested time at ALMA comes in bands, 3 & 4, and would provide more time for bands 1 & 2 to operate during long baseline configurations. Even discounting the ~70% of antenna hours by a factor of two because it is returning poor weather time (its contribution as being "less significant" to the annual budget), it still represents a quality weighted 1.6 weeks a year, or more than 3% of time.

In addition to the savings directly in terms of antennas and full array time used, the improved stability allowed by expanding the ADC/ICT test time during engineering periods, should provide more rapid efficiency improvements and a more robust system overall. The current steady state time allocation for ALMA assumes that during PI scheduled time, up to 5% could be lost for technical downtime. Our current performance is slightly out of line with this goal but it is expected that the HiL could return 0.5-2% of this downtime back to the PI. As this comes directly into the estimates against the total allocated time, the return is directly in proportion to the overall allocated time. Taking the estimates above from ~5% of total time to 5.5-7% of total time returned to the PI. As technical downtime is typically front loaded early in the cycle (October-December), and is uniformly distributed during the day, there is no need to discount this time as "poor weather" return, meaning even in the discounted for weather quality case, a savings of 3.5-5% overall, per year.

Certain development projects, from new receiver bands to new photonics or new correlators, would benefit from the HiL capability. While use of the HiL capability for this sort of testing may remove the benefits above for a brief period of time (e.g., for a new band it the benefits would compound, for central components the benefit would replace those above rather than complement), the resulting improvements could be dramatic. As an example, ALMA is currently considering an upgrade of the baseline correlator. Without a facility like the HiL, the components would be tested purely in simulation in Charlottesville and then shipped to the ALMA site. After basic simulation tests, the hardware would need to be installed into the current infrastructure at the high site, effectively stopping all PI science. While it is difficult to estimate the total downtime, experience from construction (e.g., upgrade from two to four quadrants), it could be significant. Subtle timing, sequencing, and load issues will not be shown until real data is examined. The HiL facility would allow timing and sequencing issues to be relieved before installation, with PI science continuing in the process. Conservative estimates of the required downtime with HiL run at 2-3 months. It is possible that the time required to commission without the HiL could run 4-5 months, representing a one time saving of 1-2 months, or 10-20% of the 11 month observing season at the AOS.

Summary

Overall, the HiL will provide guaranteed time saving on the array, paying back its investment in its first year of operation. It also serves as a platform for testing upgraded hardware, which, in cases where the upgrade is a central resource within the array, could result in considerable reductions in commissioning time.

Estimates of the reduction in full system lost time for this upgrade range from factors of 50% to factors of several. At a minimum, the HiL simulator will likely save ~2 months of downtime in the commissioning process. Given that ALMA's annual operations budget is ~80 MUSD, an investment of 1% of this budget in the HiL, will return ~15% of the year we spend commissioning the new correlator and will return ~5% more time, annually, to PIs worldwide, in terms of successful executions.