

HARP: a submillimetre heterodyne array receiver operating on the James Clerk Maxwell Telescope

Authors: H. Smith¹, J. Buckle¹, R. Hills¹, G. Bell¹, J. Richer¹, E. Curtis¹, S. Withington¹, J. Leech^{1&4}, R. Williamson¹, W. Dent², P. Hastings², R. Redman³, B. Wooff³, K. Yeung³, P. Friberg⁴, C. Walther⁴, R. Kackley⁴, T. Jenness⁴, R. Tilanus⁴, J. Dempsey⁴, M. Kroug⁵, T. Zijlstra⁵, T. M. Klapwijk⁵

¹Cavendish Astrophysics Group, Cavendish Laboratory, University of Cambridge, J J Thomson Ave., Cambridge CB3 0HE, UK

²UK Astronomy Technology Centre, Blackford Hill, Edinburgh, EH9 3HJ, UK.

³Herzberg Institute of Astrophysics, 5071 West Saanich Road, Victoria, BC, V9E2E7, Canada

⁴Joint Astronomy Centre, 660 N. A'ohoku Place, Hilo, HI, 96720, USA

⁵Kavli Institute of Nanoscience, Faculty of Applied Sciences, Delft University of Technology, Lorentzweg 1, 2628 CJ Delft, The Netherlands

ABSTRACT

This paper describes the key design features and performance of HARP, an innovative heterodyne focal-plane array receiver designed and built to operate in the submillimetre on the James Clerk Maxwell Telescope (JCMT) in Hawaii. The 4x4 element array uses SIS detectors, and is the first sub-millimetre spectral imaging system on the JCMT. HARP provides 3-dimensional imaging capability with high sensitivity at 325-375 GHz and affords significantly improved productivity in terms of speed of mapping. HARP was designed and built as a collaborative project between the Cavendish Astrophysics Group in Cambridge UK, the UK-Astronomy Technology Centre in Edinburgh UK, the Herzberg Institute of Astrophysics in Canada and the Joint Astronomy Centre in Hawaii. SIS devices for the mixers were fabricated to a Cavendish Astrophysics Group design at the Delft University of Technology in the Netherlands. Working in conjunction with the new Auto Correlation Spectral Imaging System (ACSIS), first light with HARP was achieved in December 2005. HARP synthesizes a number of interesting features across all elements of the design; we present key performance characteristics and images of astronomical observations obtained during commissioning.

Keywords: submillimetre, array, heterodyne, focal-plane, JCMT, SIS, ACSIS

1- INTRODUCTION

HARP is a 16-pixel heterodyne focal-plane array receiver, operating in the submillimetre, which has recently been commissioned on the James Clerk Maxwell Telescope (JCMT). HARP provides simultaneous mapping and high resolution spectroscopy when used in conjunction with the ACSIS (Auto Correlation Spectral Imaging System) back-end [1, 2]. With submillimetre heterodyne receivers approaching background-limited performance, the scientific need for focal plane heterodyne array receivers is imperative. Through building multiple detectors, mapping speed can be increased by at least an order of magnitude, enabling observations which require samples of useful statistical size. HARP operates in the submillimetre frequency band, spanning 325-375 GHz, a frequency range which contains transitions from nearly all the most abundant molecules in interstellar gas.

In late 2005 HARP was installed on the right-hand Nasmyth platform of the JCMT. Its 16 SIS detectors are arranged in a 4x4 grid layout of 30 arcseconds spacing. At 345 GHz the beam size is 14 arcseconds and the field of view is 180x180 arcseconds. HARP and ACSIS, along with the upgraded JCMT Observatory Control System, allow for rapid observing and data taking. Image de-rotation optimises observing efficiency for objects similar to or smaller than the array field of

view by ensuring no edge detectors miss the source emission. For larger objects, and particularly for raster mapping, the array can be orientated with respect to the scan direction to provide good sampling. Fast rastering enables data to be taken continuously at up to 10 Hz, and it is possible to make fully-sampled maps of 1 square degree in less than one hour. HARP uses a cold SSB (single sideband) filter to minimise the system temperatures and improve calibration accuracy. HARP sensitivity has exceeded expectations, with the combination of receiver noise temperature and beam efficiency significantly better than its original design specification of 330K (SSB). HARP is easily and rapidly tuned across the operating range of 325-375 GHz via automated control software in 10 to 30 seconds. ACSIS offers wide bandwidths; up to 1.9 GHz for each of 16 IF channels. Alternatively, high spectral resolution modes can be used, with channel spacing as small as 31 kHz, or 0.03 km s^{-1} , in the HARP frequency range. In addition, the system can provide up to two sub-bands per channel, allowing simultaneous observations of lines within the HARP 1.9 GHz IF bandwidth.

2 - THE HARP INSTRUMENT

To deliver the performance reported above, the HARP instrument itself comprises three major sub systems: Optics, Imaging Array (mixers, LO and IF units) and Support Systems (cryostat, control electronics and software). Figure 1 shows a 3-D representation of the system on the JCMT. HARP utilises a number of innovative features across all elements of the design, the most important of which are presented in the following sections; key features of the designs are discussed together with their physical implementation and performance.




Figure 1: HARP on the JCMT

2.1 - HARP Optics

The optical system for HARP has three sections: the K-mirror; the ‘warm’ optics (optical relay and calibrator external to the cryostat) and the ‘cold’ optics, which are inside the cryostat. The basic design & analysis of the HARP optical system has been presented in earlier papers (Design: [4], Analysis: [5]) and so this paper will present only an overview of the system to aid comprehension, with more detail on the physical implementation and results of various laboratory and ‘on sky’ measurements. [Note for reference that in this paper each of detectors making up the 4x4 array grid - designated ‘receptors’ in the working instrument - is labelled sequentially ‘A’ to ‘P’].

2.1.1 - K-Mirror

The K-mirror acts as an image/polarization rotator and brings the beam through the encoder onto the Nasmyth platform; it forms an image of the secondary mirror at a point near to the JCMT elevation encoder. It consists of three large mirrors, all powered but with only modest curvatures. Rotation of the three-mirror unit about the axis of beam propagation through an angle θ causes that image (and signal polarization) to rotate through an angle 2θ . On the JCMT the usable range of motion was found to be about ± 57.5 degrees, so the total amount of image rotation available is about ± 115 degrees. The K-mirror has very low losses and aberrations, especially for the wavelength range and field of view of HARP. In order to meet the needs of possible future JCMT instrumentation, its surfaces have been finished so that they are compatible with the highest operation frequencies of the JCMT, which are around 870 GHz. The details of the

K-mirror design can therefore be decoupled from those of the rest of the system. Figure 2 shows an image of the actual K-Mirror as installed and a table of its main performance characteristics.

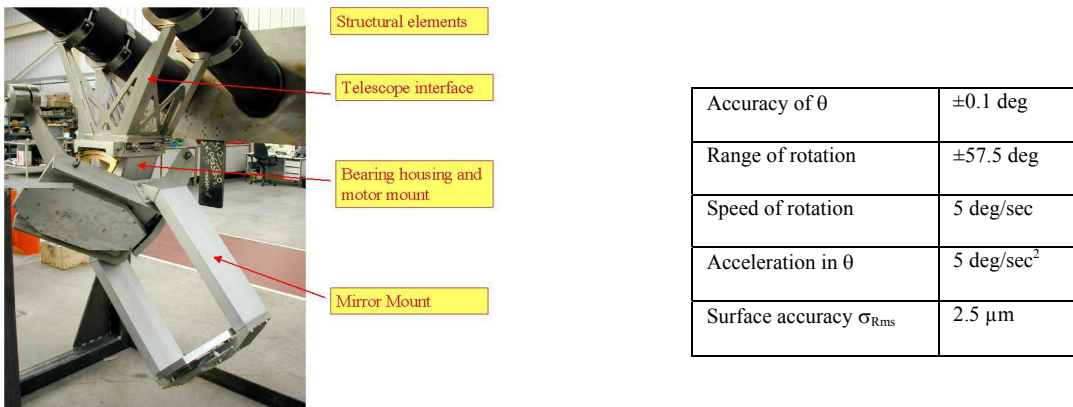


Figure 2: The K-Mirror and its performance

2.1.2 – Optical Relay, Cold Optics & Calibrator

Considered essentially as a single system to maximise performance, these are shown diagrammatically in figure 3.




Figure 3: The HARP main optics & calibrator system (note that the line source is shown as above for clarity. In the real system it is injected from a point perpendicular to the plane of the paper above the centre of CAL2)

Optical Relay

The primary purpose of the relay optics (M4 & M5) is to produce a beam which is small enough to pass through the window into the cryostat and then to form, inside the cryostat, a ‘pupil’ (i.e. a further image of the secondary mirror) which is at a good location for a cold stop. M5 also allows rapid switching to a calibration position.

Cold Optics

The main component of the cold optics is a polarizing Mach-Zehnder interferometer, which separates the signal and image sidebands. The HARP IF centre frequency of 5 GHz sets the requirement that the difference between the paths of the two beams through the interferometer should be 15 mm and this, together with the field of view of the array, sets the overall scale of the system. The interferometer is arranged so that the image sideband is terminated on an internal cold load, the 'sideband load' or 'SSB dump'. The other purpose of the cold optics is to truncate the sidelobes of the feed pattern so that they see cold absorber inside the cryostat rather than thermal emission from the region around the elevation encoder. This is done by using a reflective cold stop; the mirror C1 surrounded by absorbing tiles. The critical elements of the interferometer were the two grids and the fixed and moving mirrors (the moving mirror is required to tune the interferometer). These four items are mounted on a single stiff frame, while the input grid, the cold stop, and the focusing mirror for the SSB dump, are mounted as a separate unit, known as the cold fore-optics. The SSB dump also has a separate mounting. The final grid set at 45 degrees just below the array serves to terminate the cross-polarized response of the mixers and is built into the 20 K heat shield. The SSB dump is a roughly cubical isothermal cavity lined with mm-wave absorbing tiles. It is cooled by the second stage of a CTI cooler and the requirement here is simply to get it as cold as possible. In use temperatures of the order of 12 K are observed. It was assumed that there would be a non-negligible level of IR leakage into the Cryostat via the window and so a suitable input baffle was designed and fabricated. The purpose of this is to limit the IR radiation passing into the 60 K volume occupied by the interferometer. It uses a series of ellipsoidal surfaces to reflect as much as possible of the incoming radiation back to the window. The Dewar window itself is ~6 mm quartz with quarter-wave Teflon coating for RF matching.

Calibrator

When calibration is required, the mirror M5 is tilted. The two positions of M5 can be seen in the diagram. Only a small rotation (10 degrees) is needed. The curved mirror CAL1 refocuses the beams to make them reasonably compact at the input to the calibration loads. The rotating mirror 'CAL2' is a flat which rotates about the axis shown to provide four 'ports' for calibration. The main calibration loads are two isothermal cavities ('warm' & 'cool') with conical feeds. Their temperatures are set to give a difference of ~60 K, with the cool load set to a few degrees below ambient. A third position on the calibrator allows the receiver to look at a sheet of absorber tiles at ambient. The final calibrator port contains the optics for insertion of a low level signal in the HARP band for test and alignment purposes – the 'line source'. Position switching between calibration loads takes less than 1 s.

Optics Testing & Beam patterns

To test the overall quality of the optics we measured the shape of the far-field patterns. It was most convenient to perform a focal plane measurement using the image of the array that is produced 10.5 meters beyond the position of the JCMT elevation encoder. For this test, signal was injected via the relay optics; an analogy to how HARP will work at the telescope. In practice, the measurement was made accessible by using two large flat mirrors to turn the beam around on itself within the laboratory and thus provide an easily accessible plane for the measurement. A second measurement plane was available at the pupil where all the beams are co-incident at the notional position of the JCMT elevation encoder. A physical optics simulation of HARP was performed early in the design phase of HARP using GRASP8 [5]. By comparing the beam patterns in both focal and pupil planes with the results of that simulation it was possible to check whether HARP was performing as expected. In theory, mapping the amplitude and phase in one plane could have completely characterised the beam. However, as the beams at the pupil and focal planes have a Fourier transform relationship, it was sufficient, and much more convenient to map the beam power in these two planes, and not to measure the phase. The measured beam patterns were obtained in the laboratory by mounting a test source on a motorised stage with freedom to move in two independent axes X & Y at the measurement planes. HARP was then tuned to an appropriate frequency and the output from one of the pixels at a time connected to a spectrum analyser, acting as a narrow band filter and detector. The HARP calibration line source unit was used to generate the test signal. Its feedhorn having a half power beam width of 15 degrees, making it much wider than the angle subtended by the array, it was not necessary to take its gain into account when analysing the beam patterns.

Results

The techniques discussed above were used to map the beams as the instrument was being built, allowing us to make adjustments to the detector array mirrors; the cause of much of the initial misalignment. A more minor misalignment of the cold stop was also detected. A large volume of experimental data was collected and reduced, hence only one or two examples of the final measurement results are presented here owing to space constraints. Figure 4 shows focal plane

beam pattern measurements of pixel ‘K’ in the array, showing a full pattern centred on its design X-Y position (left), and pixel ‘E’, showing power levels at larger radii from its beam centre (right). The axes give the nominal position in millimetres and the colours indicate power level relative to the peak value (see scale bar).




Figure 4: Beam pattern measurement for Pixel ‘K’ (left) and at larger radii from pixel ‘E’ (right)

Figure 5 shows the final focal plane beam map across the entire array (15 pixels only). All are of the correct size in the correct positions and have good sidelobe performance.




Figure 5: HARP Beam maps measured in the laboratory

Figure 6 shows ^{12}CO integrated intensity images of the compact carbon star IRC+10216 from several detectors, extracted from a single commissioning observation taken in raster mode. The axes are now in RA and DEC, and the area shown covers 210x210 square arcseconds. The plot is orientated so that detector position in the plot matches that of the laboratory results shown in figure 5. Three detectors are not shown.




Figure 6: On sky ^{12}CO integrated intensity images of a compact source.

Cold Interferometer Performance

Laboratory testing of the interferometer was done for all pixels using the calibration line source. Example results for rejection at 340 GHz and 362.5 GHz versus moving mirror position are shown in figure 7; excellent sideband rejection levels can be seen across all pixels in the array, easily meeting the design requirement of simultaneous rejection of >15 dB for all pixels. On-sky tests of the sideband rejection were carried out by tuning HARP so that the emission from a strong line source appeared in the image sideband. The tuning was adjusted to place the emission line at different distances from the centre of the IF. Figure 8 (a) shows the resulting spectrum in the centre of the band, where the emission line has an integrated flux of $513.7 \text{ K km s}^{-1}$ and a peak intensity 24.04K . Figure 8(b) shows the spectrum when HARP is tuned so that the line is in the centre of the image sideband. In the image sideband, there is a possible detection of the line, with a peak intensity of 0.2 K , and the integrated flux over the frequency range covered by the emission line is 1.94 K km s^{-1} . These measurements were carried out for all HARP detectors at several frequencies across the band, and demonstrated a better than specification performance. Overall, across all of the detectors and the full tuning range, the sideband rejection is better than 19 dB.




Figure 7: Sideband rejection at 340 GHz (left) and 362.5 GHz (right) as measured in the laboratory




Figure 8: (a) Spectrum of ^{12}CO taken towards IRC+10216. (b) The residuals from the image sideband. The image sideband frequency is shown on the top scale. Note the difference in the intensity scales.

2.2 - SIS Imaging Array Unit





Figure 9: The HARP imaging array

Illustrated in figure 9, this comprises:

- 16 mixer blocks (horn-reflector antennas) in a 4x4 array configuration with F7 beams at 345GHz; with fixed-tuned, DSB, SIS mixer devices, including a facility for rapid heating of the mixers to allow speedy removal of trapped flux.
- 4-6 GHz IF, with approximately 30 dB of gain internal to the cryostat and noise temp of ~ 7 K;
- 'meander-line' LO injection unit
- Phase locked auto-tuning LO mounted externally and injected via a window on the cryostat.

2.2.1 – Mixer blocks

Each mixer comprises an SIS device mounted in rectangular waveguide fed by a corrugated waveguide horn with a reflector mirror at its aperture (the basic imaging 'pixel'). The reflectors are designed to produce a focal plane slightly in front of the array, commensurate with the requirements of the LO injection system. For HARP, the feed-horn, waveguide, device and IF output slots forming a 'mixer block' are manufactured as split aluminium blocks, giving blocks that are light weight, low thermal mass, free from waveguide flanges and straight-forward to machine. The corrugated horn design was extensively refined to make the manufacturing process as fast and reliable as possible. This was achieved by using a design that has fewer corrugations per wavelength and constant depth slots along the whole of the length compared to those of a canonical design. The final 350 GHz mixer block is shown in the image in figure 10 with specific mechanical details for the corrugated horn in the table.



semi-flare angle	9.36°
aperture radius	5.4 mm
no. of grooves	82
slot depth	260 μm
slot width*	210 μm
tooth width*	150 μm
start radius of throat	370 μm
length of start taper*	971 μm

Figure 10: Photograph of a HARP mixer block. Table: Details of final HARP horn design. Dimensions with an asterisk denote lengths parallel to the symmetry axis of the horn.

An example of measured vs. calculated beam patterns for an individual block (feedhorn) at each end of its operating band are presented in figure 11. Beam pattern measurements of a HARP mixer block both with and without reflectors indicated good performance across the HARP band. A useful additional feature is a small resistive heater element included in the centre of the air-cored coils (to suppress Josephson currents) at the back of the blocks. These can be fed with current to allow the detectors' temperature to be raised by ~ 10 K so that trapped flux can be easily removed. In use this process and subsequent re-cooling takes only ~ 10 s.




Figure 11: Beam patterns for HARP mixer blocks at each end of the band. Measured:- H-plane (dark grey), E-plane (light grey). Simulated (modal matching):- E-plane (solid line), H-plane (dashed line).

SIS devices

Work at Cavendish Astrophysics on designing and developing SIS devices for HARP concentrated entirely on the radial-probe design [3]. These give extremely broadband operation without the need to reduce the height of the waveguide. In the radial-probe mixer, a single-sided radial probe extracts power from a rectangular waveguide. Radial probes are able to match tunnel junctions to full-height waveguide. One of the other advantages of the radial-probe mixer is that many devices are fabricated during a single device-processing run, and these devices have similar characteristics. Since all SIS devices are desired to have identical noise levels and are fed via a common local oscillator, tight fabrication control was needed and sufficient immunity for process variations. Several runs of wafers were realized at Delft, gradually allowing iteration to devices matching the mixer-design and a yield sufficient for our laboratory test & development programme, 16 final mixers and a number of spares. Typical SSB noise temperature performance for active receptors in the HARP instrument is shown in figure 12.

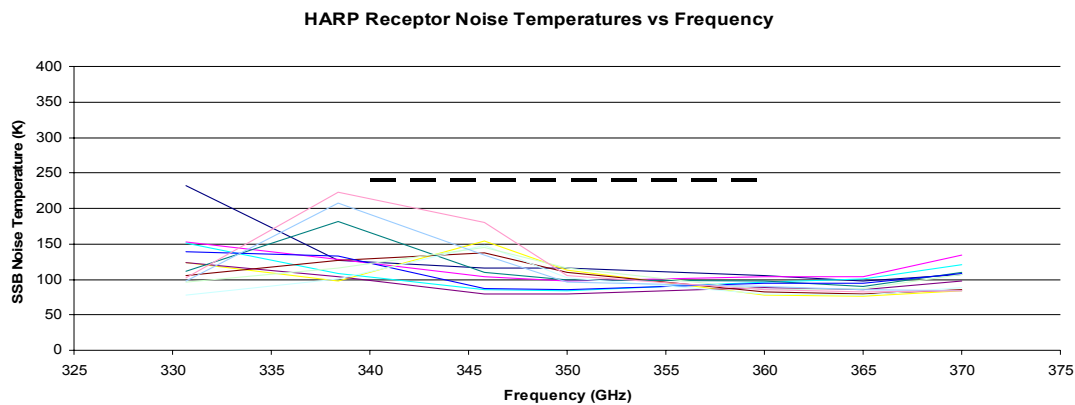


Figure 12: SSB Noise temperatures of HARP active pixels taken March 2008 (Dashed line indicates original specification requirement referred to the measurement conditions)

The curves on the graph are the results of a noise temperature measurement run with the HARP instrument on the JCMT in March 2008. The dashed line indicates the original design specification adjusted to take into account the measurement conditions; as can be seen, the performance of the HARP mixers exceeds this requirement.

2.2.2 – Local oscillator injection unit - the ‘meander line’

An array of Mylar beam splitters was adopted for reasons of bandwidth and ease of fabrication. An arrangement is used which folds the LO beam back and forth across the array, passing through the mylar beam splitters as it goes. This unit is termed the ‘meander line’. The basic concept is to have the beam splitters close to the projected apertures of the horn-reflector antennas. The horn-reflector antennas produce well-collimated beams, and therefore an array of small beam splitters can be used. This arrangement has a large RF bandwidth, and injects LO power in a highly efficient manner. Figure 13 shows a schematic plan view of the LO injection system (left) together with a photograph of the final unit (right). Single mode Gaussian modelling of the meander line scheme showed that it was desirable to alter the focal length of the feed horn reflectors to even out the coupling between the mixers. This led to the design of a novel type of ellipsoidal reflector for the individual mixer feeds. The final design has a differential coupling of around 15 % with 1 % beam splitters; final receiver testing showed that this coupling was within the required levels for acceptably uniform mixer noise performance.

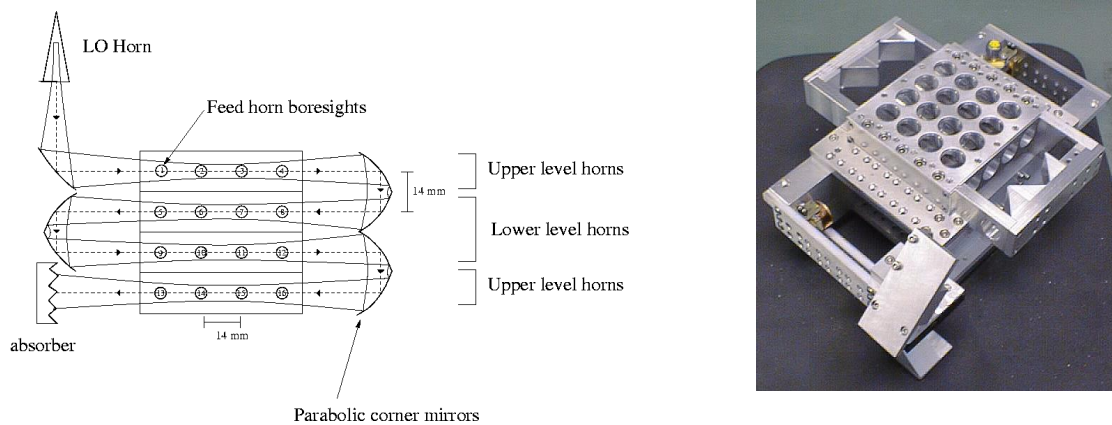


Figure 13: The LO injection system (Left: schematic. Right: the final unit)

2.2.3 - Local Oscillator

The LO system is an integral part of HARP. Its primary function is to generate a local oscillator signal in the range 330-370 GHz with a minimum power level of 1 mW across the frequency band, have good phase noise performance and be capable of frequency switching appropriate to the available ACSIS observing modes. It is comprised of two major subsystems; the LO RF Chain and a motorized micrometer operation & control system. The LO Chain is the RF system that generates the local oscillator signal. A Gunn Oscillator is used to generate a signal between 110-123.5 GHz. A portion of this RF signal is mixed down to a lower frequency and fed into a digital phase-lock system. The PLL compares it with a 350 MHz reference from a frequency synthesizer and produces an error voltage to phase-lock the Gunn output. A frequency Tripler is used to convert the microwave RF signal to a B-band LO signal to be fed to the LO injection system of the mixer array through a feedhorn. Power to the mixers is adjusted via a ferrite modulator under the control of the LO micro (see also section 2.3.3). Two micrometers are available to tune the frequency and the coupling efficiency of the Gunn output and a further pair allow adjustment of the coupling efficiency of the multiplier’s input and output. A standard ‘off the shelf’ motor controller system is used to operate and control the four tuning micrometers, which are actuated by d.c. servo motors on anti-backlash drives.

2.2.4 - I.F. System

The IF system per pixel connected to each mixer consists of five distinct units (leaving out interconnecting cables and level setting attenuators). These are, a bias-tee and a cold isolator/IF amplifier unit (1st stage) connected to the 4 K stage of the cryostat, and two ambient (un-cooled) IF amplifiers (2nd and 3rd stage) mounted externally on the Dewar wall separated by a band pass filter. A centre frequency of 5 GHz was chosen with a usable bandwidth of ~2 GHz. The typical overall mid-band IF gain is about 73 dB with noise temperature ~7 K. The output power level at the final IF

amplifier is set to approximately -17 dBm (-20 dBm/GHz, receiver terminated at ambient temperature load), as required by the Auto Correlation Spectral Imaging System (ACSIS).

2.3 – Support Systems: Cryostat, control electronics and software



Figure 14: The HARP Cryostat & Control Electronics

2.3.1 – Cryostat

The cryogenic systems design for HARP was intimately connected with the overall instrument layout and was carried out as part of that design process. The optical & array designs led to a straight-forward cryostat design. The HARP optics are in one plane in order to allow good aberration performance. This meant that the cryostat could be constructed as a lower 60K cold optical bench with optics on one side and the cryocooler on the other, with a 4 K unit on top to house the array. The cryostat uses both a two-stage CTI cooler and a three-stage Daikin cooler to maintain the operating temperature of the array at 4 ± 0.25 K and the optics package at 60 ± 5 K. There are five individual controller units containing support electronics that are mounted at suitable points on the cryostat exterior in close proximity to their individual terminating connectors on the vacuum vessel. Figure 14 shows the cryostat and control electronics on the Nasmyth mount shortly before shipping to the JCMT.

2.3.2 - Control Electronics

The HARP control electronics needs to address widely varying requirements. In order to control and monitor the 16 mixer systems we needed to overcome the problem of many large and unreliable wiring harnesses with associated connectors. However, many of the other ‘one off’ control requirements are accomplished via simple, cheap, off the shelf modules, leading to a multiplicity of different standard interfaces (e.g. GPIB, RS232, Ethernet, ‘TTL’, analogue, etc.) for the control

computer to deal with. For the multiple ‘16-off’ control requirements the solution chosen to solve the reliability problems and eliminate large external wiring harnesses was to use a fast serial bus with robust protocols to distribute commands to many small individual embedded microcontrollers located near the centres of operation in the form of the Controller Area Network (CAN) bus. CAN is a fast, well-established protocol which is very widely used in small micro-controllers and other devices. The CAN protocol deals with all the necessary bus negotiations, to ensure that message packets are delivered across the bus without supervision from high level software. In HARP a common CAN bus connects several electronic modules together in a daisy-chain configuration, each particular module performing one aspect of the control of 16 devices. The HARP control electronics is built around this central CANbus backbone, which is driven by the HARP microcomputer. Bought-in items such as motor controllers for the Mach-Zehnder interferometer, K-mirror and auto-tuning of the LO are controlled via Ethernet to a terminal server or simple parallel TTL in order to minimize the number of types of interface.

2.3.3 - HARP Control Micro and Software

The HARP receiver has a distributed control system. At the top level, the computers in the observatory control system (OCS) issue commands to the numerous subsystems that comprise the observatory using four standard commands:

- INITIALISE – initialize each device into a known state and read all configuration files
- CONFIGURE – configure each device for the next set of observations
- SETUP_SEQUENCE – set up the device for the next observing sequence
- SEQUENCE – run an observing sequence

On some telescopes the last two steps are amalgamated into a single command that might be called OBSERVE, but at the JCMT it was found useful to synchronize the completion of the set-up actions to ensure that they were all complete before the observation itself started. The interface to the HARP physical hardware is provided through a set of simple programmable controllers on the CANbus and a set of commercial devices that communicate through RS232 serial ports. The programmable controller CPUs are halted during the operation of the SEQUENCE command to prevent electrical interference with the receiver while an observation is in progress. The interface between the OCS and the hardware controllers is supplied by the HARP microcomputer, which interprets the OCS commands into detailed instructions for the hardware. This program was largely written in EPICS, an open-source set of software tools for the development of soft real-time control systems. This piece of code is commonly referred to as “the HARP control system”, but with the understanding that some parts of the complete HARP control system are actually running in the embedded controllers.

Tuning tables

Experience in tuning the receiver showed that many of the receiver components influenced each other. For example, within the local oscillator (LO) most of the components were tuned to maximize the LO output power, and the LO power varied with the positions of the Gunn tuner and backshort, the multiplier input and output backshorts, the multiplier bias voltage and the modulator drive voltage. Potentially, this could have resulted in a six-dimensional optimization when tuning the receiver to each new frequency. However, the settings of some of these components were much more sensitive than others and it was found that only the Gunn tuner actually needed dynamic optimization. The initial characterization of the receiver in the lab had to deal with the full dimensionality of the problem, but the result was a set of tuning curves that gave the optimal settings for the components as functions of LO frequency. During each CONFIGURE action, the components are all sent to their nominal setting, after which the Gunn tuner is dynamically adjusted to lock the LO to the observatory frequency standard.

A versatile tuning curve system was therefore developed for HARP. A good tuning curve system should have several properties:

- The functions are defined on the whole real line, i.e. there are no gaps within the fitted range and it is possible to extrapolate past the end points.
- The function representation should be generic, capable of fitting structures on any scale.
- It should be easy to estimate the effects of errors in the coefficients of the fitted curves.
- The effect of a noisy data point on the fitted curve should be restricted to a small interval around the bad data point and should not propagate across the entire fitted curve.
- The tuning curves should be stored in a machine-readable text file with a humanly readable syntax.
- It should be easy to fit the tuning curve to measured data.

Cubic splines are used to implement the HARP tuning curves. The spline and its first derivative are continuous at each node point, but it is not required that the second derivative be continuous, as is done for the more conventional “natural cubic splines”. We supplement the usual definition of a spline function by extrapolating the curve linearly outside the range defined by the first and last nodes. With these definitions, the curves are completely defined by the positions of the nodes and the values of the function and its first derivative at each node.

After considering several custom systems to fit these useful but unusual functions to measured data, the simplest solution turned out to be an Excel spreadsheet. A set of macros were developed to evaluate the spline functions and the Solver add-in provides a simple system to do non-linear minimization with constraints of the form required to introduce kinks and discontinuities. An example of the kind of complex data that can be fitted quite easily with these functions is shown in the figure 15, which plots the modulator drive voltage versus local oscillator frequency.

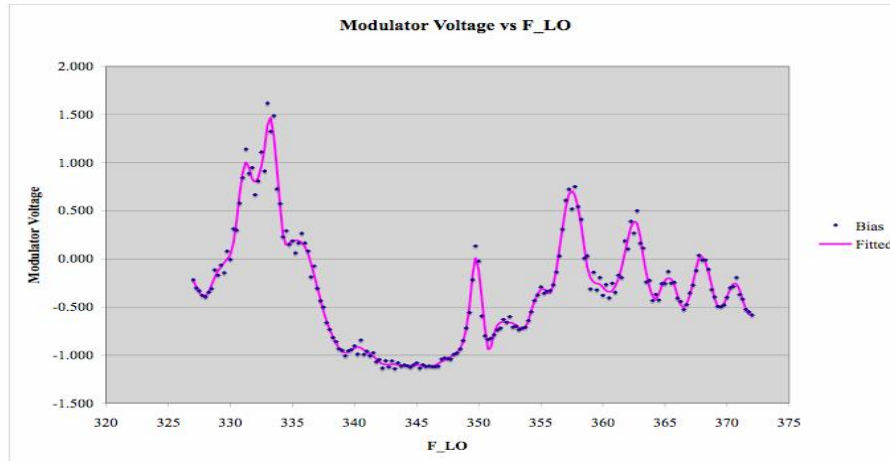


Figure 15: Curve fitting of modulator drive voltage characteristic to produce a tuning curve

3 - HARP SCIENTIFIC PROGRAMME

First light with the HARP system was achieved on December 11th 2005, with a pointed observation on W75N. The resulting spectrum is shown in figure 16.




Figure 16: First light spectrum, a detection of ^{12}CO towards W75N.

The major impact of HARP has been the ability to scan large regions rapidly using a raster scanning technique, delivering high quality spectral images in a fraction of previous observing times. In the raster observing mode [1], the array is rotated at an angle of 14.04 degrees to the scan direction. In this geometry, each detector scans a separate, non-overlapping area of the sky, so the observed region is fully sampled, with a spatial sampling of 7.3 arcseconds. The scan spacing, scan direction, map size, map orientation and integration time per pixel are all controlled by the user. This map can then be repeated multiple times to achieve the desired noise. To minimise the effects of striping in the data due to atmospheric changes, a basket-weave technique can be used, where the observed region is scanned alternately in perpendicular directions. Blank sky reference positions are observed every N rows, where N is typically 1-3 depending on sky conditions and the length of the row. Astronomical standard observations are made regularly, along with observations to check pointing offsets, which are necessary every hour or so. Therefore, the maximum map size for a single observation is governed by calibration necessities and is effectively 1 square degree, observable in less than an hour.

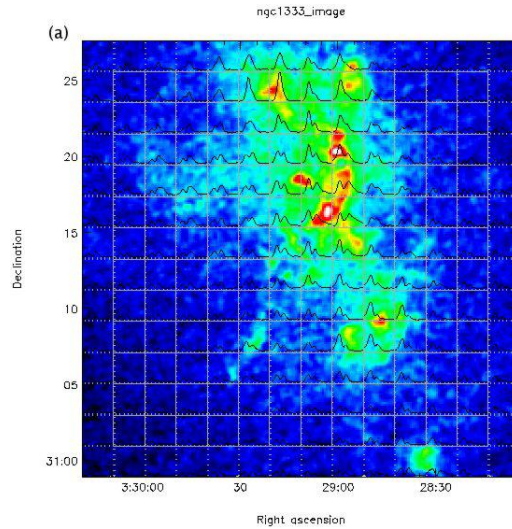


Figure 17(a): Integrated ^{12}CO intensity image of the star forming NGC1333 region, overlaid with binned spectra.




Figure 17(b): ^{12}CO volume rendered map of the NGC1333 region, highlighting the velocity space in the 3D image.

Shown in two views in figure 17 is a 0.5×0.5 square degree map, with a pixel size of 7.5×7.5 square arcseconds and a velocity resolution of 0.42 km s^{-1} , observed in ^{12}CO towards the low mass star forming region NGC1333. There are $257 \times 257 \times 2048$ pixels in the map, or 1.3×10^8 in total. In Figure 17(a), the morphology of emission is shown, in an integrated intensity image, overlaid with average spectra from several tens of pixels. In Figure 17(b) the map has been volume rendered in order to highlight the velocity space of this 3 dimensional data (the image was produced using the GAIA software package). The directions of each of the 3 axes are shown on the plot. Distinct cores and blue and red outflow lobes from young protostars can be seen in the morphology of the peaks and troughs in the image velocity space.

Molecular transitions in the 325-375 GHz band available with HARP trace the moderate density and temperature regions of molecular clouds, typically 10^3 - 10^7 cm^{-3} , and 10-100 K, and probe the warmer, more excited gas associated with energetic events such as star formation. This science is highly complementary to the wide-field continuum studies of cold dust in the earliest stage of star formation that have been carried out with SCUBA, and will be continued with SCUBA-2. However, HARP offers the added benefit of being able to derive dynamical parameters from the line profiles of appropriate molecular lines.

For smaller sources, less than 2 square arcminutes, there are three observing modes available, the grid position switch, the jiggle-chop, and the jiggle position switch [6]. All of the observing modes can be used with a variety of velocity resolutions and bandwidths, utilising the capabilities of the ACSIS back-end.

One of the features of ACSIS is the ability to split the output from HARP into two frequency sub-bands, using an LO internal to ACSIS [1]. This allows simultaneous observing of lines that are separated by up to 1.9 GHz, the HARP IF bandwidth, and is known as the sub-band mode. Wide-band imaging utilises this sub-band mode. Up to 1.9 GHz of bandwidth is available with a resolution of 0.977 MHz, or 0.78 km s^{-1} to 0.90 km s^{-1} across the HARP frequency range. The frequency overlap between the sub-bands can be set by the user. Spectral imaging surveys are clearly an important research area that will exploit this new capability (e.g. the JCMT Spectral Legacy Survey [7]), offering a much better understanding of the molecular content and physical conditions of a variety of astronomical objects. Figure 18 shows a 1.9 GHz spectrum towards OMC1 at the rest frequency of 345.79 GHz, where the rich molecular content of this frequency band, surrounding the strong $^{12}\text{CO } J=3-2$ emission line, is clearly demonstrated. Spectral surveys can be efficiently observed, with the ability to tune HARP quickly, in 10-30 s, through automated control software.

Wide-band observations are also important in extra-galactic studies, and the new capabilities of HARP means that distant galaxies can be fully observed in a single spectral imaging observation at good resolution. This is vital for studying the distribution and properties of molecular clouds in a variety of galaxies. Figure 19 plots a volume-rendered image of M82 observed in ^{12}CO , where the colour scale indicates peaks in intensity in space across the galaxy and in channels across the spectrum. The directions of the 3 axes are shown on the plot – with the axes rotated to highlight the two peaks in velocity, and the differences in shape as the galaxy is seen in red- and blue-shifted emission. The embedded square shows the galaxy image in a single spectral channel.




Figure 18: Wide-band spectrum of OMC1




Figure 19: Wide-band ^{12}CO volume rendered image of the M82 galaxy. The embedded square shows the image plane in a single spectral channel.

One of the aspects of the sub-band mode with most impact in star formation research is the ability to observe astronomical important molecules simultaneously. An example of these are the C^{18}O and $^{13}\text{CO } J=3-2$ transitions at 329.331 GHz and 330.587 GHz respectively. C^{18}O and ^{13}CO are important tracers of the kinematics of star forming regions, and this observing mode is being exploited by the JCMT Gould Belt Legacy Survey [8]. High spatial and velocity resolution are vital in studying the detailed morphology and dynamics of small, clumpy, fragmented core regions, while wide-field imaging is crucial in studying the molecular outflows and the extent of the turbulent interaction across star forming regions. The sub-band mode can also be used to observe the same molecular species at different resolutions and bandwidths. Figure 20 shows 4 spectra from two different observations taken towards a star forming core in Serpens, both of which were observed in this mode during science verification for the JCMT Legacy Surveys. Figure 20 (a,b) shows ^{12}CO , (a) taken with low velocity resolution and wide bandwidth, of 0.85 km s^{-1} and 1 GHz, and (b) taken at high resolution, 0.05 km s^{-1} . Figure 20 (c) and (d) show spectra of ^{13}CO and C^{18}O , taken towards the same source. Again these two spectra are from the same spatial pixel, but different sub-bands, observed simultaneously. The velocity resolution in both of these spectra is 0.05 km s^{-1} .




Figure 20: Data taken simultaneously at two velocity resolutions (a,b), and at two different frequencies (c,d).

In addition to the increased mapping speeds, one of the other major benefits of HARP is the decrease in calibration errors and increase in image quality, since the multiple detectors observe the source simultaneously. Relative calibration of individual detectors is more accurate, as data are taken simultaneously through the same atmospheric path. Observations on the Moon show the relative calibration errors for the individual detectors to be less than 5%. Calibration errors from point to point, and pointing errors, are greatly reduced, and by making overlapping maps, pointing drifts between maps can be removed, as is currently carried out with optical and infrared imaging.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of the astronomy research funding agencies of the UK, Canada and The Netherlands

REFERENCES

- [1] Dent, W., et al., *HARP and ACSIS on the JCMT*, ASP Conf. Ser. 217: Imaging at Radio through Sub-millimeter Wavelengths, 2000, p33
- [2] Lightfoot, J. F., Dent, W. R. F., Willis, A. G. and Hovey, G. J., *The ACSIS Data Reduction System*, ASP Conf. Ser. 216: Astronomical Data Analysis Software and Systems IX, 2000, p502
- [3] J. W. Kooi, G. Chattopadhyay, S. Withington, F. Rice, J. Zmuidzinas, C. Walker and G. Yassin . *A Full-Height Waveguide to Thin-Film Microstrip Transition with Exceptional RF Bandwidth and Coupling Efficiency*. Int Jnl IR & Millimeter Waves, Volume 24, Number 3, March, 2003
- [4] Smith, H., Hills R., Withington, S., Richer, J., et al. *HARP-B: a 350 GHz 16-element focal plane array for the James Clerk Maxwell telescope*. Proc SPIE conference on Astronomical Telescopes and Instrumentation, Hawaii, 25-28 August, 2002
- [5] Williamson, R., Hills, R., Smith, H., *Physical Optics Analysis of HARP-B*. Proc SPIE conference on Astronomical Telescopes and Instrumentation, Hawaii, 25-28 August 2002
- [6] Tilanus R., JCMT, http://www.jach.hawaii.edu/software/jcmtot/het_obsmodes.html
- [7] Plume, R., Fuller, G. A., Helmich, F., et al., *The James Clerk Maxwell Telescope Spectral Legacy Survey*, 2007, PASP, 119, 102
- [8] Ward-Thompson, D., Di Francesco, J. Hatchell, J. et al., *The James Clerk Maxwell Telescope Legacy Survey of Nearby Star-forming Regions in the Gould Belt*, 2007, PASP, 119, 855