

The NOVA Fringe Tracker: a second generation cophasing facility for up to six telescopes at the VLTI

Jeffrey A. Meisner^a, Walter J. Jaffe^a, and Rudolf S. Le Poole^{ab}

^aLeiden Observatory, Niels Bohrweg 2, 2333 CA Leiden, The Netherlands

^bTNO Science and Industry, Stieltjesweg 1, 2628 CK Delft, The Netherlands

ABSTRACT

The NOVA Fringe Tracker (NFT) is a proposed solution to the call by ESO for a second generation fringe tracking facility. This instrument at the VLTI will enable the cophasing of up to 6 telescopes simultaneously. Using broad band optics with detection from 1.2 to 2.4 microns, a unique configuration is employed that eliminates so-called “photometric crosstalk.” This refers to imbalance in the beam combiner which results in fluctuations of the incoming wavefronts and the proportion of power accepted by a spatial filter masquerading as a visibility, a common problem afflicting previous interferometric instruments and fringe trackers. Also proposed for use in “science instruments” (for the measurement of visibility), the “Polarization-Based Collimated Beam Combiner,” with its achievement of photometric symmetry in hardware, is particularly suited for combined use of the smaller AT (1.8 meter) telescopes with the UT (8 meter) telescopes involving a 20:1 intensity ratio of the interfering beams, and also for fringe tracking using highly resolved sources having a very small visibility. Recent enhancements to the proposed fringe tracker include selectable modes which detect only a single quadrature phase, both quadrature phases, or an uneven combination of the two. Optimization of partial spatial filtering using pinholes has been performed using a wavefront simulator and simulated tracking loop. Aiming for an instrument achieving the best limiting sensitivity, analysis and simulations predict that reliable cophasing will be obtained using the 1.8 meter AT telescopes tracking on an unresolved reference star with a K magnitude of 10

Keywords: Astronomical interferometry, Beam combiner, Fringe tracking, VLTI, Phase referencing, Visibility estimation, Near infrared

1. INTRODUCTION

The NOVA Fringe Tracker (NFT) is a proposal and hardware design¹ for a fringe-tracking facility for the VLTI. It was submitted in response to a 2009 ESO request for phase-A studies towards a so-called second generation fringe tracker (thus going beyond the existing FINITO² and PRIMA³ facilities) to be installed at the VLTI. Unlike those earlier instruments, the new fringe tracker would be required to cophase 4 or in the future up to 6 telescopes simultaneously in order to meet the demands of, for instance, MATISSE⁴ (which accepts 4 input beams) or future science instruments which may take advantage of the VLTI infrastructure’s eventual capability of delivering beams from 6 telescopes to the interferometric laboratory. By enabling coherent integration^{5,6} through hardware tracking of the randomly varying OPD or through use of offline phase correction,^{7,8} the ability of a science instrument to obtain accurate visibilities on dimmer targets is greatly enhanced.⁹ Therefore the sensitivity limit of an interferometric instrument is determined by the sensitivity of a fringe tracker which it relies on to permit coherent integration. The NFT was designed, above all, in order to maximize this limiting sensitivity and permit interferometric observation on the largest number of science targets.

Using the dual-beam facility of the VLTI¹⁰ fringe tracking can also be performed on an arbitrary star⁵ very near to the science target for so-called *off-axis fringe tracking* in which case the sensitivity limit for science observations is yet extended, albeit with a reduction of precision dependent on the separation between the tracking star and the science target. Thus the sensitivity of the fringe tracker will have a large effect (according to its sensitivity raised to the 3/2 power) on the proportion of the sky that will be within a given angular distance to a suitable tracking star, again allowing for a greater number of particularly dim science targets to be observed interferometrically. The NFT will operate in this mode as well when the beams from the off-axis star are directed into it, thereby allowing it to control the common delay-line carriages which simultaneously removes the random OPD fluctuations from the paths of the light from the science target (and/or supplying a phase reference for offline coherent integration of the science target’s visibility).

Send email correspondence to meisner@strw.leidenuniv.nl

Phase-tracking using stellar light from a long-baseline interferometer was first demonstrated in 1979.¹¹ Although the potential of coherent integration to enable visibility estimation on faint targets was also early appreciated, its routine use has been limited by various challenges, especially due to the increase in telescope apertures beyond r_0 which (lacking sufficient adaptive optics correction) confounds coherent techniques. Much effort involving off-axis fringe tracking has been for the purpose of narrow-angle astrometry which faces its own set of issues which cannot be addressed by the fringe tracker itself. Phase referencing between different wavelengths requires measures to address random dispersion fluctuations. Realizing the full potential of these techniques thus requires work in a number of areas, however it starts with creating a fringe tracker which is not only sensitive, but robust, able to identify and track a particular fringe peak over an extended period, able to simultaneously estimate dispersion fluctuations, and able to ascertain and report the confidence of its tracking periods. The proposed NFT and the algorithms referred to in this paper satisfy these requirements and can be the basis of moving forward to realize the potential gains obtainable through these techniques.

1.1 Timeline of the NFT Proposal

In response to the call by ESO to conduct phase-A studies towards a second generation fringe tracker at the VLTI, three different groups were awarded contracts in 2009 to pursue such work. A consortium was formed under the leadership of NOVA (an association of Dutch astronomy departments) which also included Leiden Observatory, the Technical University of Delft, TNO Science and Industry, and Landessternwarte Heidelberg. The NOVA consortium began work in mid-2009 which resulted in the submission of the NFT proposal to ESO in mid-2010, around the time of the previous SPIE conference on astronomical interferometry at which time a paper was presented¹ which concentrated on the unique beam-combining topology through which “photometric crosstalk” (section 2.1) is suppressed.

In mid-2010 ESO also received the results of a second phase-A study¹² answering the same call, submitted by Laboratoire d’Astrophysique de Grenoble (LAOG). This proposal employed a very different beam-combining technology using integrated optics that they had very successfully developed for interferometric applications,¹³ and proposed also for use in fringe tracking.¹⁴ It might be mentioned that the same group had also explored using integrated optics with a so-called multiaxial beam combination scheme¹⁵ (rather than the pair-wise coaxial interference used in both the submitted POPS proposal as well as the NFT) which they had not submitted to ESO after concluding that its sensitivity would be inferior to their submitted proposal. There continues to be some interest in such multiaxial beam combiners, but no fully developed proposals based on that concept were submitted to ESO.

A more refined version of the NFT proposal was submitted to ESO in mid-2011 at which time ESO was not only still indecisive concerning the choice of technology to be employed in a second generation fringe tracker, but very unclear as to the process that would lead to any such decision. The groups participating in these studies were supplied with each other’s detailed proposals and invited to come together in a meeting in December 2011 where the various merits and concerns regarding the proposals were debated. The NFT had been strongly optimized for limiting sensitivity, and most participants recognized that if it worked as expected its sensitivity would be significantly better than that of the POPS proposal (according to both our comparative analysis and comparison of our respective simulation results, the sensitivity enhancement was approximately a factor of 6). On the other hand POPS relied on a basic hardware design that was already operational, not as a fringe tracker but in the form of PIONIER,¹⁶ an H band visitor instrument at the VLTI which was easily commissioned and has subsequently produced valuable scientific results. So unlike the sort of integrated optics beam combiner that has been successfully demonstrated using starlight, the NFT hardware concept was seen as “risky” by many attending that meeting. The meeting ended in a quandary.

Two of the aspects of the basic NFT proposal which repeatedly raised concern and outright skepticism, were the use of a single phase (rather than quadrature phase) detection of interference (without any OPD scanning), and the lack of photometric monitoring of the individual incoming beams. It was argued by the proponents of the NFT that these features were unneeded, as demonstrated by detailed simulations, and had been avoided in order to use all of the incoming radiation in efficiently producing an error signal for driving the delay lines. However in order to allay the lingering doubts, a simple addition to the NFT hardware is proposed which allows an arbitrary portion of the starlight to be diverted for these very purposes. The principle involved is laid out in sect. 2.3.4 and its proposed implementation in the NFT is described within sect. 3 of this paper. Even though it may well turn out that the original design which lacked this added hardware was already sufficient and achieves the best sensitivity, their inclusion is a good idea even if only for diagnostic purposes. In retrospect, they should have been included in the original proposal.

1.2 Outline of this Paper

A major *qualitative* aspect of the NFT design which is not quantified using any widely recognized performance specification is the achievement of *photometric symmetry*. This is a feature of the beam combination topology, extensively discussed in the related paper in the previous proceedings,¹ which has been dubbed the *polarization-based collimated beam combiner*. This issue is summarily dealt with below in sect. 2. The other design choices and how sensitivity has been maximized are also reviewed in sect. 3 as we walk through the hardware design. However in addition to the hardware described in the earlier paper, we include the hardware enhancements referred to above which permit measurement of the other quadrature phase and of the individual beams' photometries; the principle employed is schematically diagrammed in fig. 3.

Section 4 discusses the high-level algorithms used for control and interpretation of fringe tracking data. In sect. 4.5 we present simulations and performance estimates of the closed loop fringe tracking system using the NFT hardware. These simulations employed such algorithms (thus simulating the entire fringe tracking system) along with the linear control system which simply generates a delay line correction proportional to the interference signal obtained. The algorithms described were not intended as a formal part of the NFT proposal, and their applicability to fringe tracking are not specific to use of the NFT hardware. If they perform in practice as they do in these simulations then they are of value in response to concerns over the high-level realtime control software through which a fringe tracker will control an interferometer's delay lines. In the present case, they have been employed in detailed simulations of the operation of a fringe tracker using the NFT hardware. Qualitative aspects of the tracking algorithm, such as detection and correction of fringe jumps ($\pm N\lambda$ offsets from the intended fringe peak), reliability and confidence determinations, and reporting of status to the science instrument are thereby addressed.

In the previous publication¹ describing the NFT, a laboratory experiment was described which demonstrated operation of the polarization-based collimated beam combiner, and its successful rejection of photometric crosstalk. Since that time, the experiment was updated to include a simulation of visibility fluctuations, an aspect of the experimental setup which had been previously omitted. The augmented apparatus and experimental results are presented in sect. 5.

2. BEAM COMBINER DESIGN AND THE PROBLEM OF PHOTOMETRIC CROSSTALK

The topology used in the NFT has been dubbed the *polarization-based collimated beam combiner*. Not only is it very well suited to producing 6 pair-wise interference signals between 6 telescopes, but it attacks the problem of *photometric crosstalk* which results in performance limitations affecting on-axis beam combiners. For an extensive treatment of this problem and the solution obtained using this beam combiner topology we refer the reader to the 2010 proceedings.¹ A brief review of that discussion follows.

2.1 The Problem of Photometric Crosstalk

Interferometry between two beams can be done in more than one way, however if one wishes to reduce the number of detector elements (or illuminated pixels of a detector array) in order to reduce the effect of detector noise, then one is forced into the basic configuration shown in fig. 1. The beams' amplitudes are superimposed onto a single spatial mode (ideally) with two output arms, each of which can be fed into a single detector. If the beam combining element (a partially transmissive mirror) is lossless, as we would desire, then all of the input power from the two beams, I_A and I_B must wind up as I_1 and I_2 going to one or the other detector. However interference between the beams due to (partial) correlation of their amplitudes will shift some of the power away from one detector and towards the other. Subtracting the detector readings will retrieve the amplitude of interference in a particular phase, that is:

$$I_1 - I_2 = k \cdot \text{Re}\{e^{-j\psi} \tilde{V}\} + I_{\text{xtalk}} \quad (1)$$

Thus the detection is sensitive to the component of the underlying complex visibility $\tilde{V}$ in the phase ψ (and insensitive to the visibility in the quadrature phase $\psi + \frac{\pi}{2}$).

Now ψ is dependent on the exact geometry and can be altered by applying a small OPD and is not a matter of concern. In the ideal case, we would seek a beam splitter whose intensity transmittance and reflectivity are $R = T = \frac{1}{2}$. In that particular case it can be shown that k , the resulting response when $|V| = 1$ at the phase angle ψ is $k = 2\sqrt{I_A I_B}$ and $I_{\text{xtalk}} = 0$. This hardware outcome is perfect.




Figure 1. Single spatial mode beam combiner using a partially transmissive mirror. The output electric fields E_1 and E_2 are a linear combination of the input electric fields E_A and E_B . Subtracting the detected intensities $I_1 = |E_1|^2$ and $I_2 = |E_2|^2$ yields an estimate of the visibility. When the beam combiner lacks photometric symmetry ($\gamma \neq 0$), the estimate will include “photometric crosstalk” due to fluctuations in the incoming beams’ powers $I_A = |E_A|^2$ and $I_B = |E_B|^2$.

However a typical beamsplitter as shown will generally *not* satisfy that exact condition. Rather, even assuming a lossless beamsplitter whose $R + T = 1$, there will be a non-zero γ defined as $\gamma = R - T$. We term this condition ($\gamma \neq 0$) as *photometric asymmetry*. One effect of photometric asymmetry is to reduce k , but this reduction in sensitivity is usually not very great. However of great concern is the introduction of what we shall term *photometric crosstalk*, I_{xtalk} , according to:

$$I_{\text{xtalk}} = \gamma (I_1 - I_2) \quad . \quad (2)$$

In other words, added to our estimate of the visibility at phase ψ is a term whose magnitude may be comparable to that of the measured visibility itself unless $|\gamma| \ll 1$. Since the so-called *photometry** of each beam, I_A and I_B , will fluctuate due to the atmosphere, particularly if spatial filtering is used, I_{xtalk} will be a random quantity; there is no direct way of removing this term from the estimate of visibility. Herein lies the problem.

2.2 Combating Photometric Crosstalk

One obvious solution to avoiding photometric crosstalk is to produce a beam combiner whose $R = T$, and this is often attempted. However making a partially reflective coating with this specification is easier said than done, and moreover any practical coating will have a splitting ratio that is a function of wavelength and polarization. Particularly for a wideband instrument like the NFT, this would be a futile venture.

Another unacceptable solution is that of “photometric monitoring” and correction. Here, the idea is to implement the system shown in fig. 1 but to split off a portion of the input beams and detect those separately in order to measure I_A and I_B along with I_1 and I_2 . With this information it is possible to subtract the contribution of I_A and I_B toward I_{xtalk} according to (2). Now although photometric monitoring is widely employed, particularly for the calibration of visibility (calibration of k), it can at best reduce, not eliminate, photometric crosstalk. As before, R and T will be dependent on wavelength and polarization, as will the beamsplitter used for detection of the photometries I_A and I_B . What’s more, there is a net increase in the error of visibility measurements due to the detection noise of the photometries and due to the weaker interference signal after having removed power in order to measure the photometries.¹

The solution to overcoming the effect of I_{xtalk} that can be made to work and is universally employed, is that of one or another OPD modulation scheme such as “ABCD stepping” or scanning across a fringe packet. This very familiar technique also has performance drawbacks¹ such as requiring a higher detector readout rate in order to reduce the effects of random fluctuations in I_{xtalk} between one detection and the next. However it is widely used precisely because it does work.

We now examine the proposed solution to eliminating photometric crosstalk without any compromise in performance.

*“Photometry” has become a jargon in the field of stellar interferometry in referencing the total power of an optical signal regardless of its correlation with any other optical signal.




Figure 2. The essential optical components of the polarization-based collimated beam combiner. Input beams from a number of telescopes are combined pairwise, such as the horizontal polarizations of beams from telescopes 2 and 3. The beam from telescope 2 (and other even beams) has its polarizations reversed (1) so that its horizontal polarization is now vertically (*s*) polarized entering the polarization recombining stage (2). The resulting “polarization-recombined beam” contains waves from telescopes 2 (*s* polarized) and telescopes 3 (*p* polarized) in the same spatial mode, and may be operated upon by any series of optical elements so long as those components have flat polarization characteristics. Finally, the polarization analyzer (3) separates the polarizations as seen at +45° and -45°, sending them to separate detectors (or pixels of a detector array). Subtracting these detector readings yields an estimate of visibility (at a particular phase) unaffected by photometric crosstalk.

2.3 The Polarization-Based Collimated Beam Combiner

This proposed topology for beam combination implements the intended function of fig. 1, that is interfering two input beams to get two output beams whose intensities can be subtracted as in (1) to get an estimate of visibility in a particular phase. However this time we can actually achieve a very small $|\gamma|$ over a wide wavelength range. This does *not* rely on any precision coating in an optical element such as the beamsplitter in fig. 1 for instance. Rather the achievement of $\gamma \approx 0$ hinges on a particular geometry, and canceling any residual γ involves only the angle of a mechanical rotation, and a gain adjustment to the signals received from a photodetector.

The polarization-based collimated beam combiner produces interference between one polarization of the light coming from two telescopes, thus we are content to divide the light of each telescope by polarization in order that it can interfere with two other telescopes, as is required for a 6 beam fringe tracker. Its operation is dependent on three essential stages, as is schematically outlined in fig. 2:

1. Alternate incoming beams are subject to a polarization reversal.
2. In the *Polarization Recombination Stage* (PRS) the light from two input beams is superimposed onto the same spatial mode.
3. A polarization analyzer (perhaps after any number of intervening optical elements) splits that light into polarization components at +45° and -45° which are each detected as I_1 and I_2 in fig. 1 with (1) applying as before, but achieving $I_{\text{xtalk}} \approx 0$.

2.3.1 Polarization reversal

The same polarization on the sky obtained by one telescope must be interfered with that same polarization as received by a second telescope. In order to properly feed the next stage (PRS), that requires that every alternate beam have its polarization reversed with respect to its neighbors (see fig. 2). This rather trivial operation can be implemented using a system of mirror reflections, but a preferable solution for the NFT is use of achromatic half wave plates whose axes are oriented at 45° for the beams whose polarizations we wish to reverse. Similar plates oriented at 0° are included in the other channels for dispersion compensation. However as these plates can all be rotated, it is possible to reverse the role of channels having or not having their polarizations reversed, among other possibilities.

2.3.2 Polarization Recombination Stage

In the *Polarization Recombination Stage* (PRS) a polarizing beam splitter reflects the *s* polarization of beam N, while transmitting the *p* polarization of beam N+1. The beam alignment is adjusted such that these wind up in the same *spatial* mode (but a different polarization mode, which is why it is possible to do this with the full power of each beam intact). The implementation of this for 6 input beams in order to measure 6 visibilities is diagrammed in fig. 4 which involves 12 polarizing beam splitting surfaces.

We term each output beam from the PRS a *polarization-recombined beam*, which has some special properties. In particular, such a beam may be operated upon by a wide range of optical elements *as long as* they are polarization-flat, that is, their response is not a function of polarization[†]. And most importantly, such a beam path after the PRS will *not* introduce any differential OPD effects, so vibrations and mechanical stability, along with the surface quality of the optics, will not degrade either the phase or amplitude of the resulting interference that will finally be measured!

2.3.3 Polarization analyzer and detector

Finally, the polarization-recombined beam, is split into two polarization components, each of which is detected. However the analyzer is oriented so as to select the polarization components in the +45° and -45° directions, whereas the two telescope input beams, E_A and E_B were at 90° and 0° respectively. It can be seen that the beams sent to the detectors, E_1 and E_2 , from the analyzer separating the polarizations at +45° and -45° respectively, will have amplitudes given by:

$$E_1 = \frac{\sqrt{2}}{2} (E_A + E_B) \quad \text{and} \quad (3)$$

$$E_2 = \frac{\sqrt{2}}{2} (E_A - E_B) \quad . \quad (4)$$

This achieves photometric symmetry ($\gamma = 0$).

2.3.4 Implementations for quadrature detection and photometric monitoring

The original concept for the NFT requires only detection in a single quadrature phase as is achieved by the configuration shown in fig. 2. However there is some interest in adding a capability to send a greater or lesser portion of the light into a channel measuring the other quadrature phase or one which senses the individual photometries of the two interfering beams (the total photometry from those beams can be inferred by summing the two detector readings, but it cannot be differentiated as to which beam it had come from). And these capabilities would be of particular interest in use of the topology of the polarization-based collimated beam combiner in a science instrument (to accurately estimate visibility) for which it is also well suited. The configuration schematically shown in fig. 3 adds these capabilities to the basic configuration of fig. 2 as follows:

- Splitting off part of the polarization-recombined beam and passing it through a quarter wave plate oriented with its axes in the horizontal and vertical directions, followed by a polarization analyzer at 45° and detection as before, simultaneously permits measurement of the quadrature phase of interference. In this case, the precision of the 90° phase relation between the two measured phases is exactly dependent on the retardation of the waveplate, which will not generally be 90° at all wavelengths, but at least will be fixed so that a true quadrature result can be obtained through a correction involving a linear combination of the two measurements.
- Splitting off part of the light to be directly analyzed in the horizontal and vertical polarizations (not shown) will clearly measure the power in those two polarizations received from the two telescopes which had been launched into those polarizations.
- The same individual photometry of the two telescopes' incoming beams can be accomplished while still using a polarization analyzer at 45° (which is convenient for the NFT optical configuration) by a polarization transformation of the incoming horizontal and vertical polarizations to the monitored polarizations at $\pm 45^\circ$. This can be accomplished as shown in fig. 3 through the use of a half wave plate at 22.5°, and this additional capability is proposed to be included in the NFT design.




Figure 3. For quadrature detection, the polarization-recombined beam may be split, or its wavefront divided; the division need not be equal. En route to the analyzer, one part passes through a quarter wave plate oriented horizontally (delaying the p wave). When run into a polarization analyzer at $\pm 45^\circ$, the resulting interference observed in its two outputs is in phase quadrature with respect to the other beam. Likewise photometric monitoring of the individual input beams as affected by the common pinhole or spatial filter is possible through sampling the vertical and horizontal polarizations separately. This can also be done using the same polarization analyzer at $\pm 45^\circ$ by inserting a half wave plate at 22.5° , which is convenient in the context of the NFT optical configuration.

An important point to note regarding these configurations is that any operations on the polarization-recombined beams exiting the PRS such as spatial filtering in particular, will apply identically to the light analyzed in the other quadrature phase or for individual photometry as described above. This is of particular value in a science instrument where the use of a *single* spatial filter (rather than one for each incoming beam) eliminates any need to align the overlap of the beams' spatial filters. However this does not fully hold in a design in which more than one optical mode is transmitted (rather than true spatial filtering) and in which the splitting of the beams for the different detection systems is performed using wavefront division, as there may still be some structure in those wavefronts which entail a difference in the portions sampled; this imprecision applies to the NFT configuration.

3. NFT HARDWARE DESIGN

One suggested layout for the NFT optical train is shown in fig. 4. The layout is planar with the length of the optical table shown being about 1.5 meters using folded beams as shown. The red lines plot the beam centers. The beams are nominally 18mm in diameter as received from the VLTI pick-off mirrors or dichroics (not shown), and remain afocal up until the M4 off-axis paraboloids which focus them onto the pinholes in the mask assembly. Following that intermediate focus, the beams slowly diverge to about 9mm in diameter when they finally reach the camera lens. We shall now briefly describe each stage in the optical system, as shown in fig. 4.

Dispersion compensator In order to allow for broadband interference, unequal longitudinal dispersion between the beams due especially to the air in the delay line and other beam paths must be eliminated. This dispersion can be due to the dry air in those paths but also to the water vapor component which has a different dispersion signature over wide bandwidths. What's more, in order for the NFT hardware to operate through measurement of a single interference phase (rather than quadrature phase), its most sensitive mode, the requirements on dispersion compensation are much more stringent. Not only must so-called group velocity dispersion be attacked in order for all wavelengths to be simultaneously received with full *coherence*, but the correction has to be such that all wavelengths can be simultaneously received at *zero phase*. This will involve adjustable dual wedges of certain combinations of dispersive substances that have been identified which can approximately achieve this phase correction over the near IR bands to correct the dispersion of dry air and of water vapor respectively. Thus two degrees of adjustment are required on all (or all but one) of the 6 beams. These are not shown in the diagram as their placement is not yet determined, however one option is to place them in the main beams in the VLTI lab. That way the applied correction, determined from a closed loop system involving the NFT detection and nulling of the residual dispersion, will similarly affect the light going to any near IR science instrument.

[†] And even then there are some exceptions permitted. For instance, the folding mirrors in fig. 4 are permitted so long as the reflections are within the plane, so that the mirrors' birefringence applies along the horizontal and vertical polarization directions which are the very directions of that beam's polarized components coming from the two telescopes.




Figure 4. One proposed layout of the NFT (for illustration, not to scale) accepting up to 6 nominally 18mm diameter beams reflected from 45° (dichroic) mirrors (not shown) staggered every 120mm along the VLT laboratory's beam paths (240mm separation between beams) in order to provide for OPD matching at the PRS. See sect. 3 for explanation of the optical train shown here.

Front end A set of pick-off mirrors or dichroic mirrors, not shown, intercept two or more beams in the VLT lab sending them toward the NFT hardware shown in fig. 4. Each is redirected toward the PRS by the local switchyard which may be implemented using 6 mirrors each on a long translation stage. The possible positions of those mirrors are shown in dotted lines, and by positioning the 6 mirrors at a set of those locations any of the incoming VLT beams can be routed to any of the NFT input channels. Included on each VLT pickoff mirror and switchyard mirror is a motorized fine alignment system so that the pointing of the telescope beam, and also the shift of the pupil, can be automatically adjusted during the nightly instrument setup. The switchyard mirrors are also used for OPD adjustment during setup (through small motions along their linear stages) and the mirrors are also affected by a fast piezo transducer (OPD range $\approx 50\mu\text{m}$) with internal metrology which may be used for fine OPD alignment and also for local phase locking[‡] in “tracking mode 2.”

Polarization reversers Under the principle of the polarization-based collimated beam combiner the horizontal and vertical polarization of every even numbered input beam must refer to the opposite polarization on the sky with respect to the odd telescope beams. The polarization reversal can be accomplished using achromatic half-wave plates at 45°, with similar half-wave plates, but at 0°, included on the odd input beams in order to introduce identical longitudinal dispersion. In practice, all of the half-wave plates are rotated using motors so that there is an ability to switch from the even telescope beams having their polarizations reversed to the odd beams, among other possibilities. Since these are transmissive optics with flat surfaces, the rotation can even be done during an observation without breaking lock or affecting alignment.

Polarization Recombining Stage (PRS) This is the component most resembling a classical “beam combiner,” however it should be remembered that the actual detection of interference also requires the polarization analyzer (Wollaston prism) prior to the camera. For the (up to) 6 beams, the PRS comprises a bank of 12 polarizing beam splitters which separate the polarizations of the input beams, and superimpose the *s* and *p* polarizations from adjacent input beams into the same spatial mode of each of its 6 output beams. The “end around channel” requires an additional alignment mechanism (held constant during operation) adjusting an external flat mirror which also contains an additional piezo OPD actuator used for

[‡]In tracking mode 1, the NFT controls only the VLT delay lines in order to null the offset OPD. In tracking mode 2, the OPD is nulled immediately using the local piezo with the VLT delay lines allowed to catch up at their somewhat slower rate. Tracking mode 3 involves *scanning* the OPD using the piezo, intended mainly for initial fringe acquisition and technical purposes.

setting the “closure OPD.” The other 5 channels are totally implemented within a cemented block of optics comprising 11 polarizing beam splitting surfaces.

Beam steering and intermediate focusing This consists of mirrors M4 which are off-axis paraboloids in conjunction with the tiny M5 mirror surfaces (not individually visible in fig. 4) which steer the converging beams toward the intermediate focus mask with their individual axes almost parallel but actually meeting at the camera lens. The alignment of the M4 mirrors are adjusted using precision motors during nightly instrument setup to center the stellar images on the mask’s pinholes. The tiny flat M5 mirrors might be also adjustable for pupil control, or it might suffice for them to be fixed surfaces ground onto a block.

Intermediate focus mask This is an opaque mask with a small hole for each beam, at which the star is focused. It is proposed that there be a motorized sliding mechanism or turret so that different levels of (partial) spatial filtering can be selected according to conditions[§] Even when no spatial filtering at all is desired, some field limitation is needed in order to limit the effect of background light near the detector’s long wavelength cutoff. Since thermal radiation becomes a concern at these longer wavelengths, the blocked portions of the mask, as seen from the camera side, must appear cold, so this element is blackened and refrigerated[¶].

Spectrometer This consists of a zero-deviation dispersive prism that separates the wavelength components of each of the 6 beams as they are propagating from the intermediate focus mask toward the camera. It is considered useful for the amount of spectral dispersion to be somewhat variable, so that the instrument can be easily reconfigured to map the entire wavelength range ($1.2 - 2.4\mu m$) over only 4 or up to 7 (or more) pixels on the detector. One recommended implementation involves a coarse switching between a higher and lower dispersion prism, and for finer adjustments to the spectral resolution to be effected through longitudinal adjustment of the prism position in the path toward the camera over a certain range as is suggested in fig. 4 (please note that the prism’s active surfaces are so as to effect spectral dispersion in the vertical direction, not literally as depicted in fig. 4). Use of the fine adjustment over a very wide range is to be avoided as it also induces a slight pupil shift of the various colors as they reach the camera lens. Note that the camera focuses the beams with a PSF that is much smaller than a detector pixel, so the wavelength cutoff of each spectral channel is relatively sharp with only a rather small wavelength overlap region between spectral channels.

Insertable waveplates The primary mode envisaged for fringe tracking with the NFT is for the stars’ power to be interfered in a single phase such that the measured visibility at that phase is zero (at each spectral channel) when the OPD is at the intended tracking point. Nor is there any light diverted for measuring the “photometry” of individual beams. However there is an addition which will add this option to the original NFT design. Two mechanisms will be present which can partially insert a quarter wave plate with its axis in the horizontal direction, and a half wave plate oriented at 22.5° , respectively. This implements the principle discussed in sect. 2.3.4 and depicted in fig. 3. These waveplates include a slight *wedge* which deflects the light affected by them onto a different set of pixels than the other portions of the wavefront. The increased diffraction (in the horizontal direction) due to the subapertures thus formed will not be problematic given that detector pixels are already much larger than the focused PSF using the proposed fast camera lens. It is not certain whether these capabilities will be employed only for alignment and testing or during normal operations. However it is possible for either waveplate to be inserted into the beam to a greater or lesser degree and retracted all during an observation without breaking lock or affecting the alignment.

[§] It was shown in simulations that different simulation conditions, especially in regards to the atmosphere, resulted in different sized pinholes obtaining the best performance. Having different sizes available will also be of use for testing and technical purposes.

[¶] A temperature of $-15^\circ C$ would be needed in order to prevent a substantial increase in net detection noise, but in order to operate with detector exposures as long as 10 ms a mask as cold as $-40^\circ C$ would be called for. Use of a detector with very low readout noise, such as the proposed SELEX device, could exacerbate this requirement requiring a temperature as low as $-65^\circ C$ along with a cold hood around the optics following the mask to address stray reflections.

Polarization analyzer This involves a Wollaston prism which deflects the polarization components in slightly different horizontal directions, causing the focused spots (or actually spectra) to fall on different pixels of the detector. Unlike a common Wollaston prism, this is a custom device with the birefringent material (such as crystalline quartz) cut so that its optic axis is at $\pm 45^\circ$ to the horizontal (thus implementing the principle of the polarization-based collimated beam combiner but deflecting the polarizations onto horizontally adjacent pixels).

Camera and detection system A relatively fast ($\approx f/2$) focusing lens (or mirror) supporting a clear aperture of only 10 mm is located inside the cryostat. The low resolution spectra are formed upon a relatively small portion of a PICNIC HgCdTe detector^{||} array. With this fast lens (low magnification), a diffraction limited PSF would contain most of its power in a spot much smaller than the $40\mu\text{m}$ pixel width even at the longest wavelength. Thus in conjunction with the camera's small required FOV, very limited demands are placed on the focusing lens. With 6 interference channels there will be 12 spectra formed for the main interference signal (but up to 36 if detecting both quadrature phases and photometric monitoring all at the same time), each up to 7 pixels long for the higher spectral resolution. The detector may be read out at the nominal frame rate of 1kHz (or faster) but for reaching the best sensitivity the detector integration times can be extended to beyond $\tau_0/4$, thus perhaps 10 ms under good seeing conditions.

4. FRINGE TRACKING CONTROL ALGORITHM AND PERFORMANCE SIMULATIONS

Laboratory results are described in sect. 5.4 which, when scaled to the VLTI, demonstrate that then there may be continuous tracking on a single fringe peak for minutes (barring the effect of real-world "glitches" or other imperfect hardware) when fringe tracking using fluctuating wavefronts corresponding to atmospheric conditions with a ratio of r_0 to the telescope diameter D of approximately .5 (assuming perfect tip-tilt tracking, but no AO). With a worse atmosphere, and in particular with a dim star, phase tracking will much more frequently encounter a fringe jump, which is generally unacceptable particularly if the fringe tracker is employed for coherent integration at a wavelength different from that used by the fringe-tracker. When either the atmosphere or the signal strength becomes *too* unfavorable, then the mean duration of uninterrupted phase tracking becomes so short that using that fringe tracker is pointless. Between these extremes, however, are the typical situations that a fringe-tracker must deal with, so that tracking on the intended fringe peak may proceed for seconds at a time but will be frequently punctuated by fringe jumps.

4.1 Fringe Jumps

Fringe jumps must be detected by the higher-level control algorithm, concentrating on the multi-spectral aspect of the interferometer's data. In particular, suppose we are tracking by nulling the average phase delay obtained from interferometric channels at optical frequencies (or wave numbers) ν_1 and ν_2 (so that the phase at $\nu_{eff} = \frac{1}{2}(\nu_1 + \nu_2)$ is driven to zero). Then we measure the average phase difference $\Delta\phi_{21} = \phi_2 - \phi_1$ between those two spectral channels whose center optical frequencies are separated by $\Delta\nu = \nu_2 - \nu_1$. If a fringe jump of N fringes has occurred, then the average $\Delta\phi_{21}$ observed over many detector frames will tend toward the mean given by:

$$E\{\Delta\phi_{21}\} = (2\pi N + \phi_0) \frac{\Delta\nu}{\nu_{eff}} \quad (5)$$

Such use of spectral interferometric detection is used in PRIMA for instance, and is the basis for fringe jump detection in the NFT, as employed in the control software and simulations that we shall describe below (sect. 4.5).

Of course this is not only needed in marginal conditions. For even in the case of a strong signal and quiet atmosphere, there may still be occasional fringe jumps which must be dealt with. And at the beginning of fringe tracking the system will need to identify the central fringe peak, or the so-called group-delay point (not a concise quantity, but which for the purposes of this discussion can be taken as the OPD corresponding to the N , no longer required to be an integer, needed to make $\Delta\phi_{21}$ in (5) equal to zero) around which one would normally designate a fringe peak on which the tracker will attempt to remain.

^{||} There is an interest in taking an advantage of emerging detector technologies which promise a lower read noise such as the SELEX detector which ESO has been exploring. Replacement of the proposed PICNIC detector with such a device would be straight forward and would increase the limiting sensitivity by an estimated factor of 5 when tracking on an unresolved star.

4.2 Dispersion Tracking

Another complication is that although $\Delta\phi_{21}$ will abruptly shift by $\pm 2\pi \frac{\Delta\nu}{\nu_{eff}}$ when there is a fringe jump, there will also be more gradual changes in ϕ_0 due to longitudinal dispersion fluctuations. Part of this will be due to changes in the air path of the delay lines dependent on the sidereal motion of the star, but this component is deterministic and therefore not problematic. However there will also be significant *random* fluctuations in ϕ_0 over periods of seconds due to so-called *water vapor seeing*, that is, varying amounts of differential water vapor in the atmosphere above the telescopes,^{17,18} and this varying quantity must be tracked,^{19,20} both in order to calculate a correct phase reference to apply to an instrument at a different wavelength, and in the case of the NFT in order to determine a fine adjustment to the dispersion compensator which is required not merely to cancel $\Delta\phi_{21}$ as measured between the two spectral channels in the K band, but to insure that other wavelengths detected by the NFT (H band in particular) will still be near zero phase when ϕ_1 and ϕ_2 are near zero.

There is an unfortunate degeneracy between changes in the measured $\Delta\phi_{21}$ given by (5) due to water vapor dispersion fluctuations affecting ϕ_0 , and due to an integer change in N due to a fringe jump. In most practical cases ambiguities due to this degeneracy can be resolved, but it becomes more difficult in cases of a low SNR and bad water vapor seeing. First, these are distinguishable by water vapor dispersion fluctuations affecting ϕ_0 being very gradual whereas fringe jumps causing $N = \pm 1$ are sudden. But because a $2\pi \frac{\Delta\nu}{\nu_{eff}}$ change in $\Delta\phi_{21}$ cannot be detected very quickly under a poor SNR, the possibility of such a change in ϕ_0 over the period required for that detection cannot always be ruled out. Secondly, under typical conditions of water vapor seeing, the range of random fluctuations in ϕ_0 stays within a range less than 2π wide. That means that if there is an undetected fringe jump which is misattributed to a 2π change in ϕ_0 , the consequent unlikelihood of such a large deviation in ϕ_0 will be detected either right away or after some period of time in which such an unlikely shift in the mean value of ϕ_0 will become apparent, thus identifying the error in retrospect. However under conditions of very poor water vapor seeing the range of actual ϕ_0 fluctuations might not be limited to a 2π range.

And finally, with the spectral capability of the NFT, there is not only detection of ϕ_1 and ϕ_2 (or actually $V \sin(\phi_1)$ and $V \sin(\phi_2)$) at two K band wavelengths, but a third channel in the H band will detect $V \sin(\phi_3)$. This quantity has a very different response between a fringe jump ($N = \pm 1$) at K band and a $\pm 2\pi$ change in ϕ_0 at K band due to water vapor fluctuations, providing yet another means for distinguishing between these two sorts of events. But unfortunately, at a marginal tracking SNR, the measurement of average phase offsets such as $\Delta\phi_{21}$ at the required level of precision (dictated by the small relative bandwidth of the K band $\frac{\Delta\nu}{\nu_{eff}}$) requires averaging hundreds or thousands of detection frames, so even using these strategies an error of this sort can remain undetected for several seconds. At a very marginal SNR the mean tracking interval without a fringe jump may be shorter than the time thereby required to detect the last fringe jump and resolve such ambiguities, in which case the system essentially fails and cannot ever claim confidence in tracking on the intended fringe peak with a valid measurement of ϕ_0 .

4.3 Detection of Fringe Loss

One objection raised to the NFT configuration was that because the NFT hardware nominally measures only one quadrature phase of interference $V \sin(\phi)$ (but not cosine), and seeks to null that value by adjustment of the OPD (forcing ϕ to zero), there is no positive signal present to confirm that it actually is tracking on the fringe. For if there were no fringe at all (either because of a pointing error or because the delay line had moved completely outside the coherence range of the interference) then $V = 0$ and this is indistinguishable from $\sin(\phi) = 0$, perfect tracking!

In response, a few solutions were proposed. One that was simulated with success involved superimposing a *small* square wave OPD modulation on the NFT's local delay line (piezo) at a frequency of one cycle per every two detector readouts. With an amplitude of less than .2 radians peak-to-peak, the feedback from this modulation in the detection of $V \sin(\phi)$ was typically well below the noise level when tracking on a dim star. But its presence could be confirmed through correlation over some period of time, and thus within about 1/2 second (typically) a catastrophic loss of fringe tracking could be ascertained, even at the lowest SNR which permits reasonable tracking.

Another scheme that appeared to be similarly sensitive in detecting loss of signal, was to look at the correlation of the specific outputs from two of the spectral channels, $(V \sin(2\pi\nu_1\tau) + noise_1)$ and $(V \sin(2\pi\nu_2\tau) + noise_2)$. Due to fluctuations in the uncorrected OPD offset τ , these should be positively correlated while tracking, whereas if there is a complete fringe loss ($V=0$) then the noise alone between these independent channels will not have any degree of correlation. There were other schemes considered, but this problem was judged less critical and easier to solve than the more involved detection of and correction for $\pm N\lambda$ fringe jumps.

4.4 Validation of Tracking Intervals

A fringe tracker will typically be used to control the interferometer's delay lines in order to cancel atmospheric OPD fluctuations in realtime. This allows the science instrument to perform long detector integrations and/or combine successive detector frames coherently. It is also possible for the fringe tracker to only supply an *estimate* of the atmospheric OPD for use by the science instrument in post-processing;⁸ this requires that the science instrument still take short exposures (well shorter than τ_0) but when interpreting those data the results can be combined coherently after applying the phase correction obtained from the fringe tracker's concurrent measurements. However in either case there is an additional indicator which needs to be supplied by the fringe tracker in order to interpret the science data, the importance of which has been largely unappreciated in the relevant literature. Specifically, in addition to the fringe tracker supplying (in hardware or offline) its estimate of atmospheric OPD, it must attach an indicator of *confidence* to its tracking at those times.

This cannot generally be done in realtime, but it doesn't have to be. Rather, the control system will have indications that it is actually tracking at time t (sect. 4.3) and will be looking for signs that there has been a fringe jump (sect. 4.1). Under certain conditions (SNR, τ_0 , etc.) it may be that any fringe jump will almost always be detected after 2 seconds, for instance. Then at time $t+2$, if there has *not* been a detection of a fringe jump, and the indicators of tracking are still present, the software can declare data taken at time t (2 seconds ago) to be "validated." On the other hand, when a fringe jump *has been* detected through integrating the $\Delta\phi_{21}$ indicator, it also goes back and finds the earliest point at which the fringe jump it has just detected could have occurred. Let's say that at time t it has detected a positive trend in $\Delta\phi_{21}$ indicating a fringe jump in the positive direction $N = +1$, but after looking back at the recent history of $\Delta\phi_{21}$ it can be determined that between $t-3$ and $t-2$ the average value of $\Delta\phi_{21}$ had been negative or near zero such that, considering the noise level of $\Delta\phi_{21}$, $N = 1$ is ruled out during that earlier period. Then the algorithm will *invalidate* the period between then and time t (the present), since there had been a fringe jump at some unknown point during that period continuing until time t . Since $N = 1$ has been (likely) detected, the control program orders the delay lines to jump by $-\lambda$ to compensate and records that action. A few seconds later, at time $t+3$ say, when the control loop has sufficient reason to conclude that it has been tracking on the intended fringe (with the average of $\Delta\phi_{21}$ during that period close to zero) it will feel reassured that the correction jump ordered at time t indeed placed the tracking back on the intended fringe, and will thus validate the period from t to $t+1$ (not until $t+3$ because it never validates a period until at least 2 seconds have passed). And so on.

Now the science instrument has been taking data the whole time and cannot be told to change what it has already done. But what is important is that the fringe tracker has stored a determination of which intervals were validated (or invalidated). In the subsequent processing of the science instrument's data, this validation record is consulted and *only science data taken during validated intervals* is employed in the visibility estimator. If this were *not* done and all of the data from the science instrument were accepted equally, then the visibility estimation would be corrupted by the inclusion of bad data during the periods in which the fringe tracker was not tracking on the intended fringe. For instance, if the fringe tracker had not been properly tracking 20% of the time, then the readings of the science instrument at those times might be at a more or less random phase and will not add coherently to the estimate of visibility. This might cause the estimated visibility to be reduced by 20%. This problem is solved by only integrating data acquired during periods which the fringe tracker has validated.

Such an algorithm to validate intervals where the fringe tracker has a high confidence of having been correctly tracking, has been implemented in conjunction with a simulation program which models the atmosphere and operation of the NFT. Various algorithms and parameter values were tried in order to obtain not only the best realtime tracking, but the best detection of fringe jumps (through examination of $\Delta\phi_{21}$) and determination of which intervals to retrospectively "validate" as just outlined. These algorithms were not developed to the point of perfection, but were employed in order to place a *lower limit* on the performance attainable from the simulated hardware using some control algorithm. The important performance statistics (obtained through long-term simulations) are the percentage of the time that was validated (or *tracking rate*) and the portion of time where it validated the tracking when in fact it had *not* been tracking on the intended fringe as claimed; this is termed the *tracking error rate*. By adjusting criteria used for determining validation or invalidation, it was possible to trade-off one against the other. Thus by using a tighter criterion for validation, the tracking error rate can be reduced, but at the expense of a lower tracking rate as much of the time it actually *had* been properly tracking, the criteria for validation were not quite met. Setting a lax criterion for validation, on the other hand, increases the tracking rate, but also allows more invalid periods to be inadvertently validated, contributing to a higher tracking error rate. In the simulations performed, the criteria were adjusted to obtain a reasonable compromise between these outcomes.




Figure 5. Using a frame time of 5ms, simulated tracking is successful at a K band photon flux of only 70,000 per second with a typical tracking rate over 80%. The tracking (red) follows the actual atmospheric OPD (white) in the top graphs, with their difference, the actual tracking error shown in the lower plots. Periods when the tracker was locked onto a side fringe (three are shown) are detected by the tracking algorithm (based on the phase differences between the 3 spectral channels) after typically 1/2 second, leading to a correction applied to the delay lines and invalidation of the intervals where it may not have been correctly tracking (“validated” intervals are marked in green). This simulation as run over a long term reveals that less than 1% of “validated” intervals are actually off of the intended fringe peak. OPD is measured in femtoseconds. The left and right plots show results obtained from two different detector noise realizations when tracking the same atmospheric OPD and incident wavefronts.

4.5 Simulations of NFT Tracking

A simulator was developed which modeled atmospheric wavefronts and OPD independently at two telescope apertures. The wavefronts were spatially generated according to the Kolmogorov model with a magnitude determined by the requested r_0 , and were swept across the telescope apertures according to the Taylor model in order to simulate the requested τ_0 . Images (including phase) from both wavefronts were formed at the telescope focus and truncated by a pinhole (or acted upon by a perfect spatial filter, as requested) and the results of the two telescopes interfered; this was done for each detection wavelength. The “detection” imposed a quantum efficiency and added random noise based on the detector’s readout noise level and photon noise. The simulated detector signals $V \sin(2\pi\nu_i\tau) + noise_i$ at each wavelength were run into the linear tracking algorithm (with an adjustable gain) to control the simulated delay line. Meanwhile, the high-level control algorithm as described in sect. 4.1 and 4.4 attended to the non-linear aspects of the tracking including generating corrections when a fringe jump was detected, and determining which periods would be *validated* as described in sect. 4.4.

Two typical simulations are shown in fig. 5 for the case of a marginal signal level below which reliable tracking fails. The two results shown, over a simulated period of 7.5 seconds, have employed interference using the exact same random realization of atmospheric OPD (top graph, white) and wavefronts, but in which only a new realization of detection noise has resulted in dissimilar tracking errors and fringe jumps in particular. In the left example, there were two random fringe jumps of $-\lambda$ and $+\lambda$ respectively, as can be seen in the lower graph where we magnify the difference between the tracking (red in the upper graph) and the actual atmospheric OPD. Using a smoothed version of $\Delta\phi_{21}$ (not shown) and a suitable criterion, the control algorithm detects the first fringe jump after about .4 seconds, invalidating the previous .6 seconds. The second fringe jump took about .6 seconds to detect with about .7 seconds invalidated. In both cases a correction jump was generated returning the tracking to the actual atmospheric OPD, immediately after which a new validation would be declared, but only after another second or two during which $\Delta\phi_{21}$ was observed to be sufficiently close to zero.

In the example on the right, a more interesting fringe jump of -3λ takes place. A detection using $\Delta\phi_{21}$ invokes a $+\lambda$ correction, but only after two more such detections and $+\lambda$ corrections has the tracking returned to the actual OPD. The validation routine correctly identifies the 1.2 seconds during which the tracker was not on the intended fringe. The “validated” intervals are those where the horizontal bar is colored green.




Figure 6. Key performance statistics from long-term simulations employing the tracking algorithm discussed in the text, with the NFT hardware operating at a detector frame time of 1 ms using light from the K and H bands. Assumed atmospheric parameters are $r_0 = .8m$ and $\tau_0 = 21ms$ at $2.22\mu m$ (135 THz) using the 1.8m apertures of the AT telescopes with perfect tip-tilt correction (no AO). The horizontal axis is the number of K band photons per second per telescope present at the input of the NFT. Between signal levels of 100,000 and 200,000 photons/sec the *qualitative* aspects of fringe tracking go from unacceptable to almost perfect.

4.6 Limiting Sensitivity

The above simulation shown in fig. 5 employed an extended detector exposure time of 5 ms in order to increase the fringe tracker's sensitivity; this is at the expense of a further increase in actual tracking error due to the less frequent application of corrections to the delay lines.

In this simulation (and the other quoted simulation results) we have assumed an atmosphere with $r_0 = .8m$ and $\tau_0 = 21ms$ which are better than average but not unusual for the Paranal observatory. We assumed use of the 1.8 meter AT telescopes and assumed perfect tip-tilt correction (no AO) on wavefronts characterized according to this r_0 . The above simulation, about at the limit of the fringe tracker's sensitivity, required 70,000 photons per second per telescope beam over the K band arriving at the input of the NFT optics, plus the H band radiation that would also be received (these figures are weakly dependent on the assumption of the spectrum for a K0III star). Taking into account the transmission loss of the VLT infrastructure, this flux level corresponds to use of an unresolved star with a K magnitude of 11.

Most of the simulations were performed using the faster 1 kHz frame rate requested by ESO; this lowers the average delay since the latest update to the delay line position. Now about twice as much stellar flux was required for reliable tracking. Extensive simulations obtained tracking statistics for light levels between 100,000 and 300,000 photons per second over the K band which are plotted in fig. 6. For each signal level the optimum sized pinhole was identified (through trial and error) and a simulation over a long period was performed in order to accumulate useful statistics. The well known thresholding behaviour for fringe tracking is displayed in this plot. At a flux of 140,000 photons per second, we see that the tracking rate has reached 80%, the error rate is less than 1%, and the rms tracking error (during the periods where it was properly on the intended fringe, thus 79% of the time) was 210 nm.

Since detection noise at near IR under low light conditions is due to detector readout noise, there is great interest in the development of detectors with lower readout noise or with internal gain which lowers the effect of readout noise. ESO has been exploring the SELEX detector for instance. Simulations assuming the much lower noise level of that device predict similar tracking using the light of stars 5 times dimmer than reported above.

4.7 Smoothing of Atmospheric OPD Estimates

The fringe tracker is required to correct the physical OPD offset τ through application of the error signal received from each detector channel i (each representing $V_i \sin(2\pi\nu_i\tau)$) properly weighted, and multiplied by a gain factor. According to simulations, and as might be expected, the gain which supplies the best tracking error in the linear sense is close to the gain that results in the lowest frequency of fringe jumps. Selecting the gain according to the latter criterion is preferred. The resulting rms OPD offset σ_ϕ describes the precision of the tracking enjoyed by the science instrument. Near the sensitivity limit for reasonably reliable tracking, we find $\sigma_\phi \approx .7$ radians at the K band.

We have mentioned that if the science instrument does not rely on long detector exposures then it is possible to perform coherent integration in software using an *offline* phase reference.⁸ Even if the science instrument had been using light affected by the delay lines which were optimally controlled by the fringe tracker, any better knowledge of what the actual OPD offset *had been* can be employed offline in order to improve the coherently integrated result. This can help to remove the portion of tracking error due to the latency of the main delay lines affecting the light going to the science instrument; moreover one can gain from using an optimally smoothed estimate of what the OPD offset had been. Note that the realtime estimate supplied by the fringe tracker has already been smoothed by using a fringe tracking loop gain intentionally smaller than 1, however an improved estimate of the atmospheric OPD at time t can be based not only on data taken by the fringe tracker at (or before) time t but also on *future* data which had not yet been obtained for use in realtime tracking.²¹

Thus a *smoothing filter* can be applied to the previously estimated atmospheric OPD function. The following smoothing filter function was found in terms of the atmospheric coherence time parameter τ_0 and the realtime tracking error per frame σ_ϕ (typically .5 to .7 radians for dimmer stars). The smoothed result for time t is the weighted mean of samples of the realtime estimations at $t + \Delta t$ using the following weighting function $w(\Delta t)$.

$$w(\Delta t) = \frac{1}{1 + \frac{1}{\sigma_\phi^2} \left(1.5 \frac{|\Delta t|}{\tau_0}\right)^{5/3}} \quad (6)$$

The filter is implemented as follows. Suppose the delay line affecting the light going to the fringe tracker is according to $\tau_{dl}(t)$ and the raw fringe tracker estimates of the offset OPD of the fringe (thus with respect to that delay) are $\Delta\tau(t)$. That means that the *raw* estimate of the atmospheric delay is $\hat{\tau}(t) = \tau_{dl}(t) + \Delta\tau(t)$. Then the *smoothed* estimate of what the atmospheric delay had been applies the above filter to $\hat{\tau}(t)$:

$$\hat{\tau}(t)_{\text{smoothed}} = A \sum_{n=-N}^N w(nT) \hat{\tau}(t + nT) \quad (7)$$

with the $w(\Delta t)$ given as above. T represents the detector frame time and the filter is applied over a range of $\pm N$ (thus $2N+1$ samples) where $w(\Delta t)$ is not close to zero. A is a normalizing factor: $A = 1 / \sum_{n=-N}^N w(nT)$.

In simulations with a detector frame time of 1 ms and $\tau_0 = 21$ ms, use of such a filter was found to approximately cut the OPD estimation error in half^{**}. With the realtime estimate having an error of typically 200nm rms at marginal signal levels (see fig. 6), an offline estimate can thus reduce the fringe tracking error to about 100nm rms, lowering the coherence loss affecting the coherent integration of the science data when this correction can be applied in post-processing.

The improved estimate of the atmospheric delay $\tau(t)$ can be applied to data from the science instrument as follows. Suppose the science instrument had received light subject to a delay line of $\tau'_{dl}(t)$ (which may or may not be identical to the hardware correction applied to the light received by the fringe tracker $\tau_{dl}(t)$) Then a correction is required given by $\Delta\tau(t)_{\text{corr}} = \hat{\tau}(t)_{\text{smoothed}} - \tau'_{dl}(t)$. If the science data's raw (complex) visibility estimates at optical frequency ν_i at time t is called $\hat{V}_i(t)$ then it can be corrected as follows:

$$\hat{V}_i(t)_{\text{corr}} = \hat{V}_i(t) e^{-j 2\pi\nu_i \Delta\tau(t)_{\text{corr}}} \quad (8)$$

and those results integrated coherently instead of the original $\hat{V}_i(t)$.

^{**}There is less advantage realized in cases using longer detector integration times (relative to τ_0).




Figure 7. Experimental setup as described in sect. 5.1. In addition to the experiment previously reported,¹ the augmented experimental setup (sect. 5.4) shown here includes a beam expander followed by a beam compressor between which an additional hair dryer perturbs the wavefront so that the amount of interferometric power passing through the pinhole at the analyzer (subassembly on right) randomly fluctuates. The effect of this first hair dryer does not alter the differential OPD while the second hair dryer has little effect on the wavefront and correlated power, so the magnitude of these two “atmospheric effects” can be controlled independently.

5. LABORATORY DEMONSTRATION

A demonstration of an apparatus implementing the polarization-based collimated beam combiner, the principle of which the NFT is based on, was performed at a laboratory at TU Delft (one party to the NOVA consortium). A discussion of results obtained from that experiment and their significance was included in section 5 of the previous proceedings¹ which we summarize in sect. 5.3.

Since that time, an augmentation of the experimental setup was undertaken in order to answer concerns over the ability of the single quadrature interferometric detection used by the NFT to reliably fringe track in the presence of wavefront fluctuations affecting the power passed through a spatial filter (or visibility fluctuations when not using a spatial filter). The results of this enhanced experiment are reported in some detail in sect. 5.4.

5.1 Apparatus

Figure 7 is a diagram of the experimental apparatus including those enhancements. The portion of the apparatus to the left of polarizing beam splitters PBS1 and PBS2 can be regarded as a simulation of starlight producing two (partially) mutually coherent beams, affected by simulated atmospheric turbulence and photometric variations. Those polarizing beam splitters implement the PRS (sect. 2.3.2) while the following assembly, which includes a pinhole, polarizing beamsplitter PBS3, and two silicon photodiodes, represents the remainder of the NFT (sect. 2.3.3).

The correlated component of the simulated starlight is supplied by a polarized He-Ne laser, whose axis is tilted at approximately 45° so that it will contain both a vertically and horizontally polarized component (which are 100% correlated). With the enhanced experimental setup, that laser beam first encounters a region in which its wavefront and pointing stability are degraded when a hair dryer is turned on in order to affect the air in that path; the effect is enhanced by first expanding (demagnifying) the afocal beam followed by reduction back to a pencil beam. In conjunction with the pinhole prior to the analyzer and detector assembly, these wavefront disturbances are translated into fluctuations of the signal levels, the same effect of the atmosphere as received over a telescope with an aperture of diameter D which is not small compared to the *Fried Parameter* r_0 . Although the warm turbulent air also affects the optical path length traversed by the beam at this point, that is of no consequence since it the light has not yet been split. Therefore the first hair dryer has the sole effect of




Figure 8. Waveforms downloaded from a digital oscilloscope. **A:** The raw detector readings and **B:** the sum (red) and difference (white) of those detector values (please ignore noise introduced by the 8 bit quantization). **C:** The balanced interferometric output (white) and the piezo control voltage (green) with a fixed piezo voltage and **D:** with the piezo tracking the fluctuating OPD.

creating a visibility amplitude $V(t)$ fluctuating randomly between 0 and its maximum value, so that the *detected* visibility (in the measured phase) is now given by $V(t) \sin(\phi(t))$ where ϕ is the phase due to the residual OPD being measured.

Plane mirrors M1 and M2 are used only for adjustment of beam alignment.

The laser beam is split by BS1, the transmitted light of which contains a portion of the laser's vertical polarization and simulates the correlated component of starlight received by telescope 1. The reflected light from BS1 includes a component which contains a portion of the laser's horizontally polarized light, simulating the correlated component of starlight that would have been received by telescope 2 after having passed through a polarization reverser; this is then reflected by plane mirror M3. To each of these beams is mixed radiation from two red laser diodes pulsating at two different frequencies (approximately 150 Hz and 600 Hz in the example shown). These diode lasers' light is obviously incoherent relative to the He-Ne laser and to each other, and is used to simulate a radical "photometric fluctuation" in each beam. It also effectively lowers the fringe visibility (mutual correlation) between the two resulting beams through dilution of the correlated component with incoherent radiation.

These two beams representing the starlight received by two telescopes (one with a polarization reversal) are incident on polarizing beam splitters PBS1 which reflects the vertically polarized component of the top beam, and PBS2 which transmits the horizontally polarized component of the lower beam, while reflecting the beam coming from PBS1. This implements the polarization recombining stage (PRS) (sect. 2.3.2) outputting a "polarization-recombined beam" toward the simulated backend of the NFT. That consists of an assembly which rotates about the incoming beam axis on a precision rotary mount. After transmission through a pinhole which could represent partial spatial filtering, the light is split by polarizing beam splitter PBS3, sending the components polarized at $+45^\circ$ and -45° with respect to the horizontal to their respective silicon photodiodes.

Now, the lower "starlight" beam has also passed through a region in which a second hair dryer has created turbulent warm air, but because the beam is narrow at this point, effects on the wavefront are minor. However unlike the first hair dryer this path affects only one beam, so differential atmospheric OPD variations are simulated when this hair dryer is turned on. Additionally, the mirror M3, affecting only one arm of the interferometer, is attached to the end of a piezo stack which implements a delay line that the fringe tracker can control (or which can be used to scan the OPD).

5.2 Processing of Detector Signals

The outputs of the photodiodes are amplified and subtracted (analog electronics were exclusively employed) following a differential gain control whose fine adjustment (along with a fine adjustment of the angle of PBS3 about its incoming optical axis) is used to null any remaining "photometric crosstalk" (whose signature is a component fluctuating at the frequency of either pulsating laser diode). The two detector signals are also *added* in order to obtain a measure of the combined photometric levels, with the interferometric component nulled out. With the piezo scanning, dividing the peak difference signal by the sum of the detector readings, we obtain the visibility contrast of the simulated system. With the laser diodes turned off completely, this indicates the instrumental visibility obtained by the interferometer and which, with proper beam alignment, attains a contrast exceeding 95%. Typical raw detector waveforms with the piezo in scanning mode are shown in fig. 8A which are added in fig. 8B to obtain the *total* photometric contribution (red) which is dominated

by the pulsating radiation from the two laser diodes, but when subtracted reveal the much smaller interferometric signal (white) due to the scanning of the piezo during the same period.

Rather than scanning, as used for testing and inspection, the piezo could also be controlled by the interferometric signal in order to null the resulting phase difference between the beams entering the beam combiner (thus “fringe tracking”). This employed a first order (analog) control loop filter with separate controls for the gain of the feedback signal and a damping term. Thus actual tracking of the simulated atmospheric OPD fluctuations due to the second hair dryer is shown in fig. 8D. with the tracking signal sent to the piezo in green. The residual optical phase imbalance ϕ can be seen in white, having an rms amplitude of .05 radians rms or 5nm of red light. The full $\pm V$ range of $V \sin(\phi)$, on the other hand, is seen in fig. 8C where the voltage going to the piezo is held fixed while ϕ fluctuates randomly over a few wavelengths.

5.3 Experimental Conclusions

The experiment performed as expected, demonstrating several key aspects of the polarization-based collimated beam combiner, that is, the beam combiner topology proposed for the NFT. Summarizing the previously reported¹ results:

- Adjustment of the analyzer angle and balance of the differential amplifier readily achieved rejection of the introduced “photometric fluctuations,” as expected, to a fairly high degree.
- The photometric rejection was fairly robust and stable at a level much better than 100:1 over timescales of many weeks and involving various disturbances to the optical system as long as these two critical adjustments were not disturbed.
- The interferometric response achieved a measured visibility sometimes exceeding 95% following a careful optical alignment.
- Phase tracking as proposed in the NFT design (but employing analog electronics) was achieved as shown in fig. 8D.
- Application of a square wave perturbation into the control loop indicated a loop response time shorter than .5 milliseconds with the control loop parameters adjusted optimally.
- There was no noticeable effect due to large “photometric fluctuations” induced by the pulsating laser diodes, even when the light from the laser diodes was 50 times brighter than the “starlight” from the He-Ne laser, thus simulating an observation with a visibility as low as .02. Nor was there any deleterious effect on fringe tracking when one of the two incoming interferometric beams was attenuated by a factor of 30 relative to the other one using a neutral density filter.
- In one further demonstration, the light in between the PRS and the analyzer was corrupted by passing it through both sides of a wine glass whose position and orientation were manually manipulated. This very unacceptable optical element, which would have had a catastrophic effect on the interference had it been placed in one or the other arm of the interferometer, caused no noticeable effect on the instruments’ tracking and after turning off the hair dryer the piezo voltage returned to its previous value, thus verifying that no fringe jump had occurred during this interval.
- Not present in the previously reported results, with the augmented experimental setup tracking of the fluctuating OPD was successful even with the addition of the first hair dryer to simulate atmospheric wavefront effects resulting in fluctuations of the visibility magnitude $V(t)$, as we shall now present in detail.

5.4 Results Including Simulation of Wavefront Fluctuations

Following the earlier publication¹ of these results, the experimental setup was augmented to include the first hair dryer (fig. 7) which, unlike the second hair dryer, has no effect on the differential OPD, but does cause the wavefront and pointing of the beams to fluctuate. As those beams eventually reach the pinhole prior to the polarization analyzer and detectors, that results in fluctuations in the interfered beams’ power and consequently fluctuations in the strength of the interference amplitude.

With both hair dryers running, not only is there a fluctuating differential optical path delay simulating that of the atmosphere $\tau_{atm}(t)$, but the magnitude of interference $V(t)$ fluctuates as well. Now subtracting the two detector reading supplies only the signal $V(t) \sin(\phi(t))$ which is applied to the control loop (where the uncorrected phase $\phi(t) =$




Figure 9. **Left:** Running the apparatus shown in figure 7, we record (top) a voltage reflecting the fluctuating light level ($A+B$) caused by wavefront fluctuations induced by the first hair dryer. The OPD fluctuations caused by the second hair dryer are canceled by movement of the piezo whose applied voltage (magenta trace, scale on right calibrated approximately in wavelengths) is shown. After every 10 seconds of such operation, the stream from the second hair dryer is momentarily blocked returning the induced OPD to a reference position; changes in the piezo voltage between those periods are evidence of a fringe jump during the previous 10 seconds. **Right:** Histogram of the fluctuating light level shown on the left. If this had actually been due to atmospheric wavefront perturbations affecting the portion of light passing through a subsequent spatial filter, then the histogram displayed would correspond to a particular ratio of the *Fried Parameter* r_0 to the telescope diameter; in this case it corresponds to $r_0 \approx \frac{1}{2} D$ assuming that perfect tip-tilt correction had been applied.

$2\pi \nu(\tau_{atm}(t) - \tau_{piezo}(t))$. One cannot then disentangle V and ϕ , so any reduction in V has the effect of randomly reducing the effective gain of the fringe tracking control loop. These gain fluctuations are expected in practice and are not in danger of causing delay line corrections in the wrong direction, since the sign of $V \sin(\phi)$ is always the same as that of $\sin(\phi)$. The tracking system will still operate with only a modest reduction in instantaneous tracking performance except in the case of severe fluctuations causing V to drop to near zero for a significant amount of time (dropouts of duration much smaller than the atmospheric coherence time τ_0 are not likely to lose lock), or causing spurious “phase loops” due to multipath cancellation. Simulating the effect of atmospheric wavefront fluctuations in the laboratory demonstrated the continuity of tracking under these less than ideal (but realistic) circumstances.

With both hair dryers running simultaneously (and the laser diodes pulsating) we simultaneously simulate atmospheric OPD fluctuations, strehl or visibility fluctuations, as well as photometric fluctuations. As expected, fringe jumps become frequent when either hair dryer is run at too high a power. The powers of both hair dryers were reduced just to the point that fringe jumps became infrequent. Using those settings we recorded both the OPD measured (according to the voltage applied to the piezo, canceling the simulated atmospheric OPD due to the second hair dryer) and simultaneously recorded the fluctuating power admitted through the pinhole (due to wavefront fluctuations induced by the first hair dryer). In order to perform this measurement it was necessary to turn off the laser diodes whose pulsating power was far in excess of the detected power from the He-Ne laser, so now the detector $A+B$ signal is indicative of the power accepted by the pinhole and the fluctuations in $V(t)$ which are modulating the gain of the feedback loop. However there was no difference in the apparent performance (and in particular, no increase in the incidence of fringe jumps) when the “photometric fluctuations” simulated by the laser diodes were also present, as these were efficiently canceled as explained in sect. 2 and verified in the original experimental setup.

Figure 9 shows a plot of the piezo voltage (lower trace in magenta), over a period of 110 seconds with the scale (right) calibrated approximately in wavelengths. The top trace (black) plots the total optical power ($A+B$) received, which we have caused to fluctuate. In order to assess the incidence of fringe jumps, the experiment was run as described over approximately 10 second intervals, in between which the second hair dryer (responsible for differential OPD fluctuations) was physically blocked for about 5 seconds. That gives us an opportunity to observe whether there had been a loss of lock (and consequent fringe jump) during the previous 10 seconds. It can be seen that a fringe jump occurred during 2 out of 6 such intervals, from which we can infer that fringe jumps occur at a rate of very roughly .04 per second. Examination of

the structure function of the piezo voltage (not shown) indicated that the simulated atmosphere has a $\tau_0 \approx 4$ ms, compared to about 20 ms for the atmosphere above Paranal at K band. Taking that into account, the experimental result corresponds to about .008 fringe jumps per second at the VLTI, or a mean tracking interval of about 2 minutes.

Of course while simulating the effects of the atmosphere using both hair dryers we are still using a laser as a source and thus obtaining a very high signal to noise ratio not available when using starlight. Therefore the only fringe jumps expected would be due to a failure of the feedback loop (implemented with analog electronics) to follow the OPD fluctuations using the error signal from the optical phase detector^{††}. In order to further identify the observing parameters which have been simulated by these experimental conditions, a histogram is plotted on the right for the fluctuations in the interferometric power (black trace, left plot) normalized by dividing by .037 volts, the value observed with both hair dryers turned off (shown in the $t < 0$ portion of that graph). The resulting histogram shown on the right, very roughly corresponds to the distribution of visibility fluctuations obtained in simulations using a model in which D/r_0 is around 2.0, assuming perfect tip-tilt correction. Using the AT telescopes whose $D = 1.8$ m, that would correspond to an atmosphere with $r_0 = .9$ m, somewhat better than average for Paranal in the K band.

6. CONCLUSION

The Nova Fringe Tracker (NFT) addresses the need for a new fringe tracking facility at the VLTI capable of cophasing up to 6 telescopes to support new planned and not yet planned science instruments that will come online. The hardware's sensitivity has been enhanced using unconventional choices such as interferometric detection only in the optical phase relevant for fringe OPD detection, and ultra-wideband operation over a 2:1 wavelength range. With every design option taken in order to maximize limiting sensitivity, the instrument is predicted to track reliably using the 1.8 meter AT telescopes on an unresolved star of K magnitude 10.

In addition to addressing the specific requirements posed by ESO to groups proposing a second-generation fringe tracker, the design of the NFT addresses further qualitative issues, particularly in relation to eliminating the detrimental effects of *photometric crosstalk* (sect. 2.1) by employing the *polarization-based collimated beam combiner* topology which achieves photometric symmetry. As demonstrated through a laboratory experiment (sect. 5.3), the consequential elimination of photometric crosstalk will not only address ESO's requirement for fringe tracking using both the smaller AT and 8 meter UT telescopes together (involving beams with a 20:1 intensity ratio), but will allow for tracking on very low visibility objects (such as stars resolved past the first or second visibility null). This design utilizes standard sorts of optical components, involving no technology risks. The optical train (sect. 3) is relatively simple, confined to a plane, and includes essentially no moving parts except those involved in aligning and setting up an operational mode. The inclusion of capabilities for quadrature and photometric detection (sect. 2.3.4), along with the local piezo OPD actuators, results in a great deal of operational flexibility and potential to use the hardware in currently unforeseen modes.

Using low resolution spectral detection, the instrument will not only detect fringe jumps (sect. 4.1), but be able to track longitudinal dispersion fluctuations (sect. 4.2) which are used as feedback to a hardware dispersion compensator. A control algorithm has been demonstrated in simulations (sect. 4.5) which efficiently detects and orders corrections for fringe jumps. Additionally that algorithm generates a report of which tracking intervals it has "validated," (sect. 4.4) so that a science instrument can base its results only on data where there was a high confidence of the fringe tracker having followed the intended fringe peak.

The sensitivity and performance of the VLTI, particularly with the use of new science instruments combining the light of 4 or more telescopes, will be dependent on the performance of a fringe tracking facility. The NFT hardware design (sect. 4.5) has been optimized to reach the best limiting sensitivity, and the control and evaluation algorithms such as those presented in sect. 4 will be needed in order for the science instruments to take maximum advantage of that fringe tracking ability. Perceived risks regarding these proposals have been addressed in laboratory experiments as reported in sect. 5, however we are encouraging ESO to undertake further testing of the concept using starlight perhaps through temporarily installing a stripped-down version of the hardware at the VLTI. With a greater acceptance of the NFT concept, the next step of assembling the full hardware diagrammed in fig. 4 and testing of new control algorithms, will be a great advantage to the operation of second generation science instruments coming online over the next years.

^{††}The feedback loop gain and dampening had been adjusted empirically in order to achieve the most rapid response subject to stability. The achieved fringe tracking loop had a response time of approximately .5ms.

REFERENCES

- [1] Meisner, J. A., Jaffe, W. J., Le Poole, R. S., Pereira, S. F., Quirrenbach, A., Raban, D., and Vosteen, A., "The polarization-based collimated beam combiner and the proposed NOVA fringe tracker (NFT) for the VLTI," *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **7734** (July 2010).
- [2] Gai, M., Menardi, S., Cesare, S., Bauvir, B., Bonino, D., Corcione, L., Dimmler, M., Massone, G., Reynaud, F., and Wallander, A., "The VLTI fringe sensors: FINITO and PRIMA FSU," *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **5491**, 528 (Oct. 2004).
- [3] Sahlmann, J., Ménardi, S., Abuter, R., Accardo, M., Mottini, S., and Delplancke, F., "The PRIMA fringe sensor unit," *Astronomy and Astrophysics* **507**, 1739–1757 (Dec. 2009).
- [4] Lopez, B., Lagarde, S., Jaffe, W., and others, "MATISSE: perspective of imaging in the mid-infrared at the VLTI," *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series (these proceedings)* **8445** (July 2012).
- [5] Colavita, M. M., "Phase referencing for stellar interferometry," in [*High-Resolution Imaging by Interferometry II, ESO Conference and Workshop Proceedings No. 39, Edited by J. M. Beckers and F. Merkle*], 845–851 (Mar. 1992).
- [6] Quirrenbach, A., Mozurkewich, D., Buscher, D. F., Hummel, C. A., and Armstrong, J. T., "Phase-referenced visibility averaging in optical long-baseline interferometry," *Astronomy and Astrophysics* **286**, 1019–1027 (June 1994).
- [7] Meisner, J. A., "Coherent integration of fringe visibility: a generalized approach," in [*Interferometry in Optical Astronomy, Pierre J. Lena; Andreas Quirrenbach; Eds.*], *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **4006**, 1068–1082 (July 2000).
- [8] Jorgensen, A. M. and Mozurkewich, D., "Coherent integration: to real time or not to real time? That is the question.," *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **7734** (July 2010).
- [9] Mozurkewich, D., Johnston, K. J., Simon, R. S., Hutter, D. J., Colavita, M. M., Shao, M., Hershey, J. L., Hughes, J. A., and Kaplan, G. H., "Phase Referenced Averaging as a Method for Decreasing the Variance of Visibility Measurements," in [*NOAO-ESO Conference on High-Resolution Imaging by Interferometry: Ground-Based Interferometry at Visible and Infrared Wavelengths, Garching bei München, Germany, Mar. 15-18, 1988. Edited by F. Merkle, ESO Conference and Workshop Proceedings No. 29, p.851, 1988*], 851–854 (1988).
- [10] Delplancke, F., Leveque, S. A., Kervella, P., Glindemann, A., and D'Arcio, L., "Phase-referenced imaging and micro-arcsecond astrometry with the VLTI," in [*Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*], Léna, P. and Quirrenbach, A., eds., *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **4006**, 365–376 (July 2000).
- [11] Shao, M. and Staelin, D. H., "First fringe measurements with a phase-tracking stellar interferometer," *Applied Optics* **19**, 1519–1522 (May 1980).
- [12] Blind, N., Le Bouquin, J.-B., Absil, O., Alamir, M., Berger, J.-P., Defrère, D., Feautrier, P., Hénault, F., Jocou, L., Kern, P., Laurent, T., Malbet, F., Mourard, D., Rousselet-Perraut, K., Sarlette, A., Surdej, J., Tarmoul, N., Tatulli, E., and Vincent, L., "The planar optics phase sensor: a study for the VLTI 2nd generation fringe tracker," *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **7734** (July 2010).
- [13] Benisty, M., Berger, J.-P., Jocou, L., Labeye, P., Malbet, F., Perraut, K., and Kern, P., "An integrated optics beam combiner for the second generation VLTI instruments," *Astronomy and Astrophysics* **498**, 601–613 (May 2009).
- [14] Blind, N., Absil, O., Le Bouquin, J.-B., Berger, J.-P., and Chelli, A., "Optimized fringe sensors for the VLTI next generation instruments," *Astronomy and Astrophysics* **530**, A121 (June 2011).
- [15] Tarmoul, N., Hénault, F., Mourard, D., Le Bouquin, J.-B., Jocou, L., Kern, P., Berger, J.-P., and Absil, O., "Multi-axial integrated optics solution for POPS, a 2nd-generation VLTI fringe tracker," *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **7734** (July 2010).
- [16] Le Bouquin, J.-B., Berger, J.-P., Lazareff, B., Zins, G., Haguenauer, P., Jocou, L., Kern, P., Millan-Gabet, R., Traub, W., Absil, O., Augereau, J.-C., Benisty, M., Blind, N., Bonfils, X., Bourget, P., Delboulbe, A., Feautrier, P., Germain, M., Gitton, P., Gillier, D., Kiekebusch, M., Kluska, J., Knudstrup, J., Labeye, P., Lizon, J.-L., Monin, J.-L., Magnard, Y., Malbet, F., Maurel, D., Ménard, F., Micallef, M., Michaud, L., Montagnier, G., Morel, S., Moulin, T., Perraut, K., Popovic, D., Rabou, P., Rochat, S., Rojas, C., Roussel, F., Roux, A., Stadler, E., Steff, S., Tatulli, E., and Ventura, N., "PIONIER: a 4-telescope visitor instrument at VLTI," *Astronomy and Astrophysics* **535**, A67 (Nov. 2011).
- [17] Colavita, M. M., Swain, M. R., Akeson, R. L., Koresko, C. D., and Hill, R. J., "Effects of Atmospheric Water Vapor on Infrared Interferometry," *Publications of the Astronomical Society of the Pacific* **116**, 876–885 (Sept. 2004).

- [18] Meisner, J. A. and Le Poole, R. S., "Dispersion affecting the VLTI and 10 micron interferometry using MIDI," in [*Interferometry for Optical Astronomy II. Edited by Wesley A. Traub. Proceedings of the SPIE, Volume 4838, pp. 609-624 (2003).*], 609–624 (Feb. 2003).
- [19] Meisner, J. A. and Tubbs, R. N. and Jaffe, W., "Coherent integration of complex fringe visibility employing dispersion tracking," *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **5491**, 725–740 (2004).
- [20] Colavita, M. M., Booth, A. J., Garcia-Gathright, J. I., Vasisht, G., Johnson, R. L., and Summers, K. R., "Fringe Measurement and Control for the Keck Interferometer," *Publications of the Astronomical Society of the Pacific* **122**, 795–807 (July 2010).
- [21] Meisner, J. A., "Estimation and Tracking of Atmospheric Delay Noise in a Long-Baseline Optical Stellar Interferometer and Determination of the Expected Estimation Error," *Ph.D. Thesis* (1995).