④ Interferometers using Two-Frequency Lasers

The interferometers described in the previous sections and found in physics labs (assuming such topics are even taught with hands-on experience!) all use CW lasers and look at the fringe shifts as the relative path length of the two arms is changed. While this works in principle and has been used widely, modern commercial measurement systems based on interferometry often use more sophisticated techniques to reduce susceptibility to signal amplitude changes and noise, and improve measurement accuracy, stability, and convenience. These are called "heterodyne" systems in which the laser beams are in essense carriers for a lower "split" frequency in the MHz range provided by the two-frequency laser. The split frequency is detected optically, but then can be manipulated using straightforward electronics totally in the AC domain. If you're totally confused by now, never fear. There is much more below. ;)

The microchips in virtually all modern electronics (including the CPU and memory inside the PC, MAC, tablet, or Smarkphone you're reading this on) were likely produced on photolithography systems incorporating wafer steppers using two-frequency interferometers for multiple axes of ultra-precise motion control. Based on a scientifically proven metric - the availability of used equipment on eBay :-), heterodyne systems are in much wider use than homodyne systems, by at least an order of magnitude.

Interferometer-based measurements systems typically use some type of low power stabilized helium-neon laser to produce the "yardstick" beam of light. By stabilizing the laser with reference to the neon gain curve, the accuracy of the optical frqeuency/wavelength can easily be known to better than +/-0.1 ppm (parts per million). As noted above, a basic system may use such a laser in a Michelson or similar interferometer, with a quadrature (sin/cos) detector to count fringes representing changes in path length as described above. Problems with such a system are that changes in light intensity will result in measurement errors, alignment is very critical to obtain adequate fringe contrast, and they are more susceptible to noise.

In two-frequency interferometers, a special stabilized HeNe laser is used that produces a beam consisting of two very slightly different frequencies (and corresponding wavelengths) of light simultaneously. This may be achieved by various techniques. HP/Agilent lasers employ a special tube which uses a magnet to perform Zeeman splitting resulting in useful difference frequencies being limited to around 4 MHz due to the Physics. Zygo uses an external acousto-optic modulator to produce a 20 MHz split frequency. As above, both types of lasers are locked in such a way that the optical frequency is very precisely known. A higher split frequency is desirable because it ultimately limits the maximum stage slew rate. But too high a split frequency and subsequent processing for measurement or control becomes complex.

A diagram of the general approach is shown in Interferometer Using Two Frequency HeNe Laser.

U.S. Patent #3,656,853: Interferometer System outlines the overall approach in dry patent legaleze. :) Being a patent, it doesn't really apply directly to any real system, not even the original HP-5500A system. And in this case, doesn't even appear to have one in mind. What's below is more reader-friendly.

The following description applies to the HP/Agilent implementation using Zeeman splitting to create the two frequencies. With Zygo, the method of generating the them differs, but their use in the interferometer is the same.

In the Zeeman split approach, the two-frequency laser consists of a HeNe laser tube surrounded by permanent magnets which produce a constant axial magnetic field. The laser tube is short enough that without a magnetic field, only a single longitudinal mode will normally oscillate if it is located near the center of the neon gain curve. (Those on either side will not see enough gain.) The net result of the magnetic field is that instead of a single longitudinal mode, two modes are produced that differ very slightly in frequency and have right and left circular polarization. The difference between the two frequencies is typically in the 1.5 to 4 MHz range (though some go up to 6 MHz or more), which makes the resulting signals extremely easy to process electronically. The actual difference frequency is determined by the strength of the magnetic field, length of the internal laser cavity, and other physical details, as well as the exact place on the Zeeman-split neon gain curve where the laser has been locked.

To stabilize the laser, there is a piezo element and/or heater to precisely adjust cavity length. A feedback control system is used to adjust the cavity length to maintain the position of the Zeeman-split frequencies - and thus the wavelengths - constant. The feedback is generally based on the simple approach of forcing the orthogonally polarized outputs to be equal, which results in the most stable optical frequency.

The wavelength of the laser is the measurement increment ("yardstick") and will remain essentially unchanged for the life of the instrument. For example, with the doppler broadened gain curve for the HeNe laser being about 1.5 GHz FWHM (1 part in about 300,000 with respect to the 474 THz optical frequency at 633 nm) and a 1 percent accuracy within the gain curve, the absolute wavelength accuracy will then be better than 1 part in 30 million! Not too shabby for what is basically a very simple system. In practice it's even better. :) The laser tube is not much different than the type that was used by the 100s of thousands in grocery store barcode scanners in the 1980s.

Note that the exact value of the difference frequency does not need to be very precisely controlled over the long term. Rather, it is the difference between the reference difference frequency and the measurement difference frequency that matters, and the latter only depends on the motion of the target reflector - and the speed of light. Thus, the exact beat frequency of each laser need not be precisely controlled or even precisely measured and recorded or used anywhere in the calculations.

Since the output of the laser is a beam consisting of a pair of circularly polarized components, a Quarter-Wave Plate (QWP) and Half-Wave Pate (HWP) are used to separate these into two orthogonal linearly polarized components called F1 and F2, and to orient them such that they are parallel to the horizontal or vertical axes.

The beam consisting of F1 and F2 is split into two parts with a non-polarizing beam-splitter: One part goes through a polarizer at 45 degrees (to recover a signal with both F1 and F2 linearly polarized in the same direction) to a photodiode which generates a local copy of the reference frequency (REF, the difference between F1 and F2) for the measurement electronics; the second is the measurement beam which exits the laser. The return beam is called MEAS.

The purpose of the remainder of the interferometer is essentially to measure the path length change between two points called displacement. In a typical installation, the beam consisting of F1 and F2 is sent through a polarizing beamsplitter. F1 goes to a cube-corner (retro-reflector) on the tool whose position is being measured and F2 goes to a cube-corner fixed with respect to the beamsplitter and laser. However, differential measurements could be made as well using F2 in some other manner. Various "widgets" are available for making measurements of rotary position, monitoring multi-axis machine tools, etc. But they all ultimately result in the same sort of change in basic displacement readings - but perhaps with different scale factors or more complex calculations.

The return from the object corner reflector is F1+ΔF1 which is recombined with F2 and sent to an "optical receiver" module - a photodiode behind a polarizer at 45 degrees and preamp which generates a new difference frequency, F2-(F1+ΔF1). This signal, called "MEAS" is compared with REF to produce an output which is then simply ΔF1. The "Signal Processing" block might be implemented with digital logic like counters and subtractors, a fast microprocessor, or combination of the two. A change in the position of the object by 316 nm (1/2 the laser wavelength) results in ΔF1 going through a whole cycle. By simply keeping track of the number of complete cycles of ΔF1, this provides measurements of object position down to a resolution of a few hundred nm with an accuracy of +/-0.02 ppm! And the typical implementation will either multiply the REF and MEAS frequencies by 16 or 32 or more using a pair of phase-locked loops, or perform interpolation using sub-cycle phase compison of the REF and MEAS signals to produce a corresponding improvement in resolution down to a few nanometers or better!

The primary disadvantage of heterodyne systems is that the maximum velocity is limited in the direction that would reduce MEAS since going through 0 Hz would be confusing at best. So, one of the key specifications for these lasers is the (minimum) split frequency. For example, the HP-5517B has a split frequency range of 1.9 to 2.4 MHz with typical samples being 2.20 MHz. But the minimum is the critical value and for 1.9 MHz, the maximum velocity will be around 0.5 m/s using the simplest (linear) interferometer. Zygo lasers have a 20 MHz split frequency so the velocity can be over 10 times higher. But even 0.5 m/s is adequate for many motion control systems.

More information on the two-frequency HeNe laser can be found in the sections: Hewlett-Packard/Agilent Stabilized HeNe Lasers and Two Frequency HeNe Lasers Based on Zeeman Splitting. Searching on the Agilent Web site will yield product specific information and application notes on two frequency interferometers. A comprehesive but not too hairy description of the two frequency approach can be found in the Hewlett-Packard`Journal, February, 1976. Yes, this is an old technique (actually much older)! Searching at HP Archive for "Interferometer" and similar terms will turn up many more interesting articles. For an excellent introduction written by Agilent insiders :), see "A Tutorial on Laser Interferometry for Precision Measurements", Russell Longhridge and Daniel Y. Abramovitch, 2013 American Control Conference (ACC), Washington, DC, June 17-19, 2013. While this is an IEEE conference paper, an on-line version may be found by using the search string: a tutorial on laser interferometry for precision measurements agilent.

Zygo, another manufacturer of interferometer measurement systems using two-frequency lasers had an excellent tutorial called "A Primer on Displacement Measuring Interferometers" but it seems to have disappeared from their Web site. But never fear, I archived it at Sam's Copy of Zygo's Primer on Displacement Measuring Interferometers. And, an even more extensive introduction can now be found in the PowerPoint presentation at Introduction to Displacement Measuring Interferometry and Sam's Copy of Introduction to Displacement Measuring Interferometry. (The actual title found by Google is the same as the previous one, "A Primer on Displacement Measuring Interferometers".)

And if you find info on a laser claiming a 618 to 658 MHz difference frequency, run the other way. ;-) As of 2019, REO (Research Electro-Optics) is offering directly and on Newport's Web site what they dub "OPIS" or "Orthogonally Polarized Interferometer System" which appears to be largely MARKETING HYPE. OPIS appears to be simply a bog-standard (for any other manufacturer) 3 mW 633 nm unstabilized random polarized HeNe laser head with power supply. But the list price is about twice that of an almost identical 3 mW laser without the OPIS label. This is no doubt partly due to the MARKETING HYPE, but also because REO has to use a more complex expensive tube to do what comes naturally to other companies. More info may be found in the section: Research Electro-Optics Stabilized HeNe Lasers. I am not aware of any practical use for such a laser that would make it advantageous for heterodyne interferometry. The primary reason is that it is not stabilized, as all metrology lasers must be to have decent wavelength accuracy. It would be 50 to 100 times *worse* than a typical HP/Agilent/Keysight laser. And at times, 3 longitudinal modes may be present, and due to mode pulling, the difference frequencies between them will vary slightly representing a further complication. And performing the digital processing obtain the displacement information at 600+ MHz would be a serious challenge even if a custom chip could be designed, and probably impossible or at least highly impractical without one. Ironically, REO's stabilized HeNe (having different Marketing hype) would be suitable if (1) the polarizer that normally blocks the second mode at its output were removed and (2) the custom chip could be designed. ;-)