

Techniques for Laser Combining

Roy Clarke – Electro-Optical Systems Engineering

Introduction

Electro-Optic systems often feature a requirement to combine a number of separate laser beams into a single beam. Most commonly, the need is to provide a multi-spectral content but the pursuit of extremely high power levels in industrial lasers and particularly in laser directed energy weapons has led to an interest in “scalable” systems geometries in which an arbitrary number of otherwise identical laser beams can be added together to realise overall power levels in the 10s or even 100s of kW. Needless to say, there are numerous ways of combining laser beams and some are better suited to a given application than others; e.g. in the trivial case of a couple of beams a simple beam-splitter arrangement may suffice but there are limitations in this approach and it rapidly becomes impractical as the number of constituent beams increases. This technical note introduces some of the more novel approaches to the problem, illustrated with example systems that have been developed to deal with the issues posed by the extremes of large beam counts, high powers and wide spectral range.

Over the years we have had many occasions to manipulate and combine laser beams. In recent years we’ve been active in the higher power end of things, principally for defence (DIRCM – directed infrared counter measure, and LDW - laser damage weapons) applications. In particular we’ve been involved in corporate and national programmes to develop very high power coherent beam combining systems and high power multi-wavelength beams.

This material was originally presented as a tutorial lecture in the *Innovations Live* section of the *Photonex 2015* Trade Show.

Basics

In the most elemental form of beam combining we might envisage using some form of semi-transparent mirror (beamsplitter), which transmits one beam and reflects the other, as shown in the left hand side of Figure 1 below:

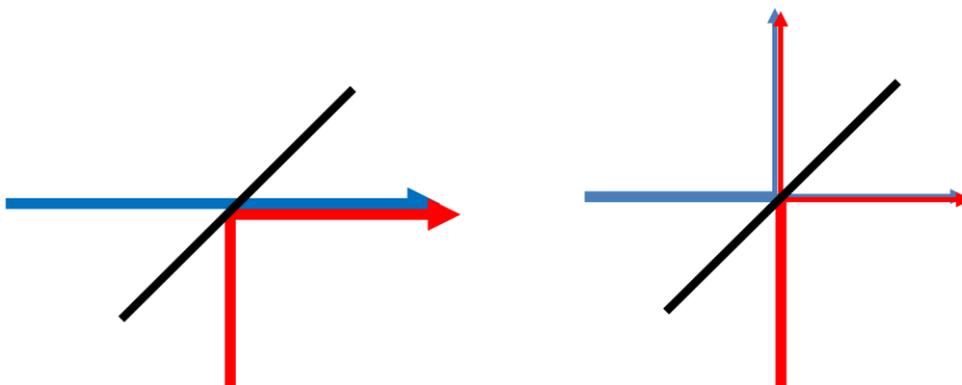


Figure 1: Elemental beam combining concept: ideal (left) and reality (right)

Unfortunately, a moment's reflection (forgive the pun) on this reveals that this simplistic ideal doesn't work (at least not efficiently) because you always lose half the total power that is available (as shown in the right hand side of the figure). For example, a 50:50 beamsplitter would lose 50% from each beam and any other splitting ratio would lose power from the beams proportionally (e.g. a 70:30 beamsplitter would lose 70% of one beam and 30% of the other). Fortunately, as we shall see, there are various means of exploiting the properties of the light to allow the ideal to be realised in practice.

Firstly though, it must be remembered that the beamsplitter is not just a single surface floating in space, but rather is a coating applied to a substrate (Figure 2).

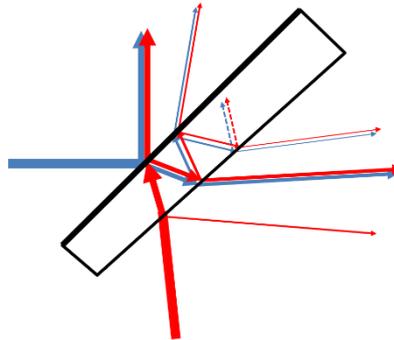


Figure 2: Multiple reflections in a beamsplitter substrate

The properties of this substrate (e.g. its transmission/absorption, optical uniformity, thermal stability, structural strength etc...) clearly have important implications for any beam combining scheme that may be devised. Generally, it will be essential to apply an effective anti-reflection coating to the back of the substrate to frustrate ghost reflections that would otherwise compromise the system.

Dichroic Mirrors

Perhaps the most commonly encountered beam combining element is the well-known dichroic mirror (Figure 3), which transmits one wavelength (or range of wavelengths) whilst reflecting others.

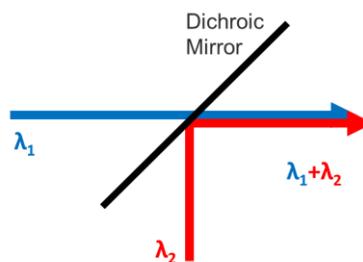


Figure 3: Twin beam combination using a Dichroic Mirror

This has been used effectively for many years and enables beams of different wavelengths to be combined with high efficiency. Modern thin-film coating technology can provide very sharply defined transition edges (i.e. the spectral range between what wavelength is transmitted and what is reflected) such that beams of very similar wavelengths can be combined using this method.

Ultimately, quite complex combination schemes can be realised, as shown in the somewhat arbitrary schematic in Figure 4, below.

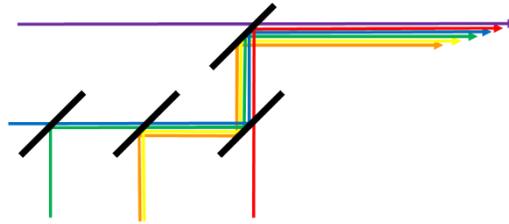


Figure 4: Multiple beam combination using Dichroic Mirrors

The decision as to which wavelength(s) are reflected and which are transmitted in this sort of scheme is usually governed by consideration of the transmission properties of the substrate. As systems become more complex the requirements on the coatings generally become more demanding and an increasing number of dichroic mirrors implies a commensurate increase in the complexity of the alignment process.

Polarizing Mirrors

If the beams are polarized, as with many laser systems, this property can be exploited to provide an efficient combination method using a polarizing beamsplitter. Reflective polarizers reflect s-polarized light whilst transmitting the p-polarizations (Figure 5).

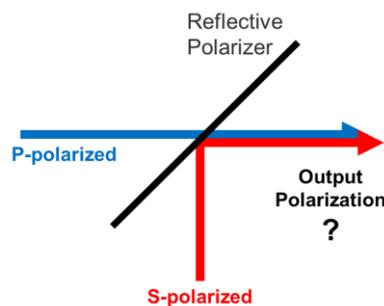


Figure 5: Twin beam combining using a reflective polarizer

If the two beams are polarized, but are otherwise incoherent, then the resulting combined beam is effectively randomly polarised. As such, it cannot be used in any further combination scheme based on this principle – i.e. the scheme is not “scalable”.

If, however, the two beams are coherent then the combined beam will have a resultant identifiable polarization, which can provide a scalable solution. Generally, the two beams will produce an elliptically polarized beam, but if the phase difference between them can be controlled, and fixed at $\pi/2$, then the resultant will be plane polarized (and at 45° to the two input beams, if they are of the same amplitude) and thus can be combined with other polarized beams in further stages.

Coherent Beam Combining with a Mirror

Coherent beams can also be combined courtesy of the π phase change that results from reflection at a higher index interface. By adjusting the relative phase of the two beams it is possible to obtain destructive interference in one of the outputs and constructive interference in the other, as shown in Figure 6 below. This can be a very efficient process and has been used effectively in extremely high power laser weapons system demonstrators.

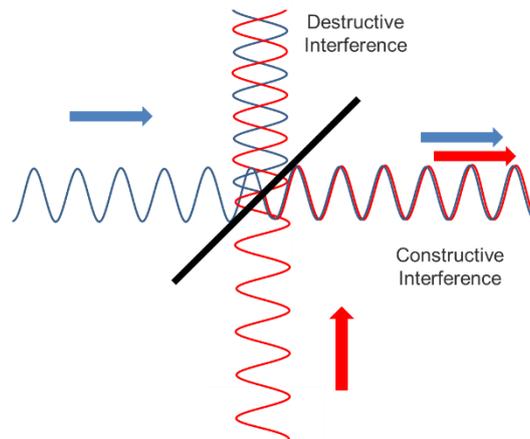


Figure 6: Efficient combining of two phase matched coherent beams

Spectral Beam Combining using dispersive optical elements

Now we'll come to more specialist methods that allow for greater scalability or more simplified optical set-ups; firstly using dispersive optical elements – on the basis that “if you split light going one way you can recombine it going the other”. These methods are obviously only useful for combining beams of different wavelengths (though, in practice, the difference can be quite small).

Those of us of a sufficient vintage might remember the rock group Pink Floyd’s 1973 album “Dark Side of the Moon”, which featured a motif similar to that depicted in Figure 7.

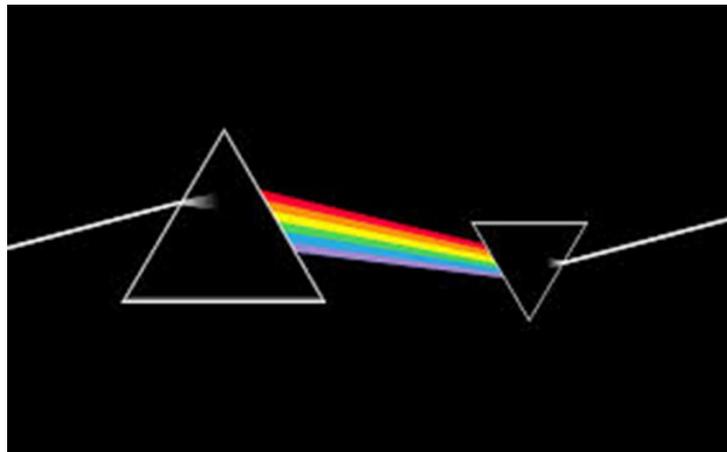
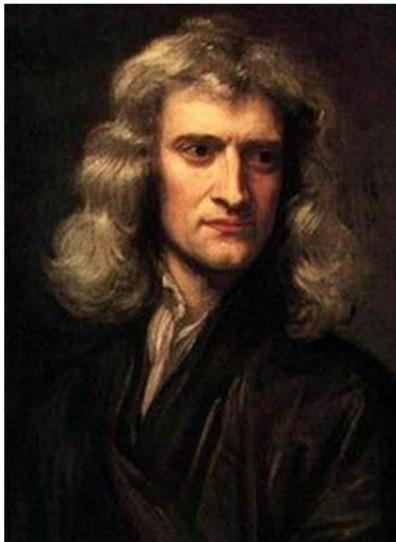


Figure 7: Spectral beam combining (as envisaged by Pink Floyd)

Superficially, this suggests that if you can split colours with a prism then you can recombine them with one too. This is indeed the case, though it will be appreciated that the way this is depicted in the artwork isn't quite right – the colours diverge out of one prism yet converge into the other. I remember, as a physics undergraduate, thinking about this at the time and was helped to the conclusion that it was OK by the discovery that a certain Isaac Newton had already demonstrated the fact some 300 years previously! As a historical aside, although his “Opticks” wasn't published until 1704 Newton had demonstrated the separation and subsequent recombination of colours by a prism

during his 1666 *Annus Mirabilis* (see Figure 8). He had retired to the country to evade the plague, which freed up his time to develop calculus, experiment with dynamics and optics, and devise his theory of gravitation (amongst other things).



"In the year 1666

... I procured me a triangular Glaſs-Prifm, to try therewith the celebrated Phaenomena of Colours.

... all the Colours of the Prifm being made to converge, and thereby, to be again mixed, as they were in the Light before it was incident upon the Prifm, reproduced Light, entirely and perfectly White, ...

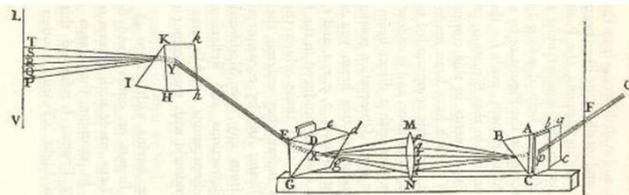
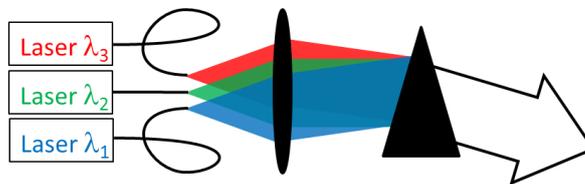


Figure 8: The first reported instance of beam combining?

Returning now to the practicalities of beam combining with dispersive elements, Figure 9 illustrates the basics of using refractive (prism) and diffractive (grating) elements.

Refractive:

- Simple & Robust
- Needs well separated λ s
- Relation between $\delta\lambda$ and $\Delta\lambda$
- Limited scalability



Diffractive:

- More complex
- Small $\delta\lambda$ possible (limits $\Delta\lambda$)
- Usually in reflection
- Surface relief gratings can be cooled for high powers
- Bragg gratings allow tailored λ s, extremely small $\delta\lambda$, low scatter, and robust
- Reasonably scalable

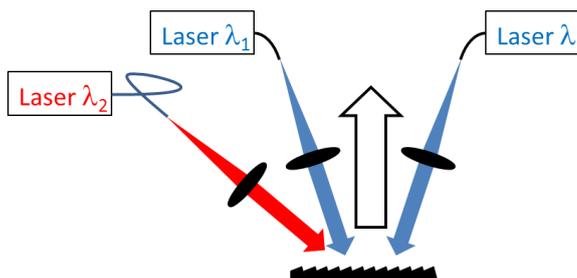


Figure 9: Essential principles of beam combining with refractive (top) and diffractive (bottom) elements

Listed in the figure are a number of issues that should be appreciated with these two basic schemes. In the case of the prism the dispersion (ability to separate wavelengths) is such that there needs to be a reasonable separation between the wavelengths being combined (otherwise the system becomes overly large as the beams need space in which to effectively separate). This also places restrictions on the range of wavelengths covered and prism systems are thus best suited to the combination of a small number of spectrally disparate beams. That being said, they can furnish simple, robust and efficient designs – as we shall see shortly in a practical example.

Diffraction gratings offer more flexibility in design over the prism, in that their dispersion is dependent on dimensional parameters (groove spacing) rather than the fundamental optical constants of the substrate material. This provides a means for combining closely separated wavelengths. The use of reflection gratings allows the substrate to be cooled, which facilitates their use in very high power laser systems. Both these attributes have been combined successfully in US laser weapon demonstrator systems.

Practical example: Ultra-wideband spectral beam combining prismatic demonstrator

A practical example serves to illustrate some of the issues that influence dispersive beam combining designs. Some time ago we undertook a proof of principle investigation to assess the practicality of combining a number of moderately powerful beams (typically, several Watt), distributed throughout the spectrum from the UV through to the mid-wave IR (as per Table 1):

Table 1: Specified laser wavelengths

Spectral Band	Specific Wavelength (nm)
UV	355
Visible	532
Near IR	1064
Short Wave IR	1908
Medium Wave IR	4075

Since these various wavelengths are quite well separated it was decided to investigate a prismatic design, which would provide a robust, simple and compact solution. Recalling Figure 9, the design concept is fundamentally very simple – a lens and a prism; the former needs to generate collimated beams whose angular separation matches the dispersive properties of the latter. The design variables are: the quality (i.e. complexity) of the collimator; the prism material; the latter’s shape & orientation; and the source positions.

From an optical design perspective this system is all about controlling the chromatic aberration; ordinarily, something to be avoided but here an attribute to be exploited. In essence, lateral chromatic aberration in the collimator equates to an angular shift in the output beams, whilst longitudinal chromatic aberration equates to a defocus. The physical size of source (e.g. fibre diameter) equates to a finite waveband, and hence is also an implied aberration.

The eventual design comprised a simple achromatic doublet (sapphire and calcium fluoride) with a sapphire prism (Figure 10). The chromatic aberrations were “managed” to allow the various laser beams to be delivered by optical fibres that were positioned at various lateral and longitudinal separations in the collimator object plane. The fibres were held in place using an etched “V-groove” assembly, shown in Figure 11. The left hand picture of this figure is a close up of the assembly, clearly showing the various “staggered” positions of the fibre tips; the right hand image shows the assembly mounted on an alignment stage during testing.

The final picture shows the prototype design being evaluated in the laboratory. The fibres can clearly be seen entering the v-groove block in front of the collimator lens, whose mount is obscuring the prism. Somewhat out of focus, in the distance, is a long focal length parabolic mirror which was used to focus the beams down to facilitate measurement of their collimation quality and relative angular alignment.

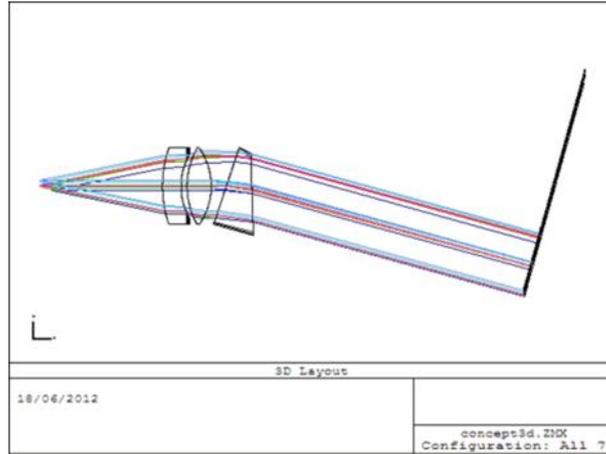


Figure 10: Optical design layout of prismatic wideband beam combiner demonstrator

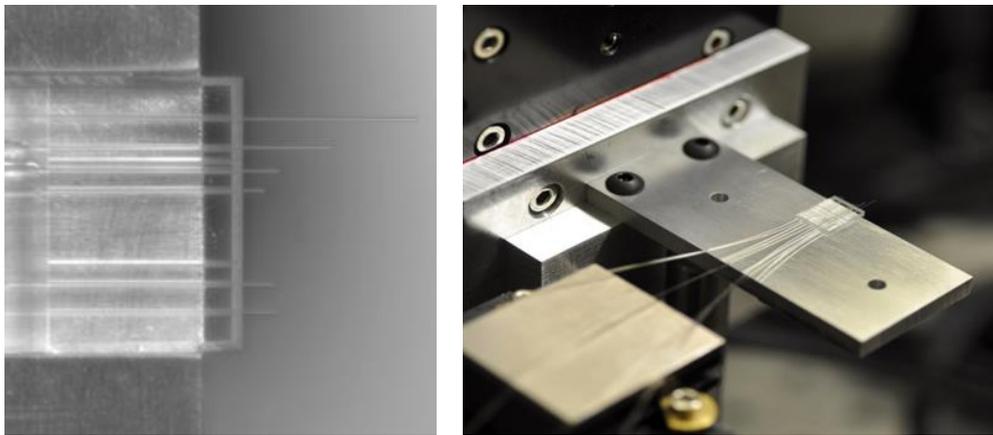


Figure 11: Detail of "V-groove" laser delivery fibre assembly

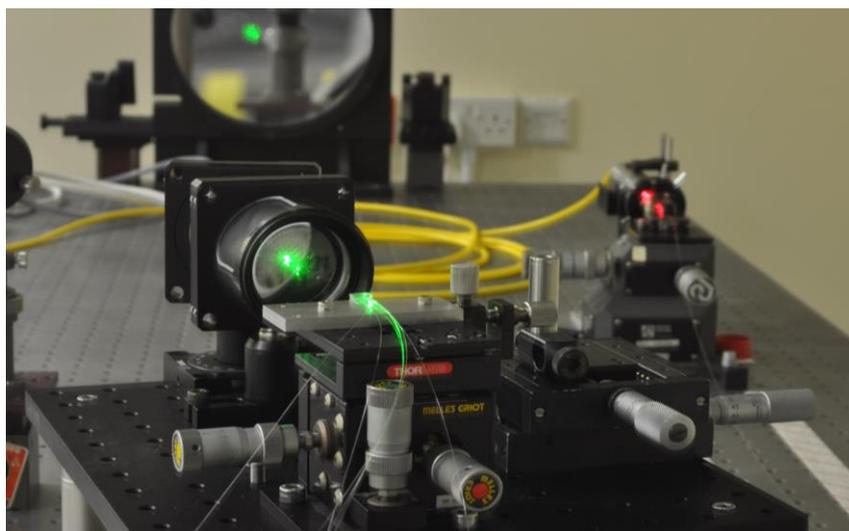


Figure 12: Prismatic beam combiner components during testing

Ultimately, this exercise successfully demonstrated that the exacting design goals could be met with a simple, robust optical system.

This work was reported in: **CD Stacey, C Stace, RG Clarke; “Ultra-broadband Spectral Beam Combiner Spanning over Three Octaves”; Applied Optics, Vol 52, No.29, pp 7200-7205 (2013)**

Tiled Aperture Beam Combining

Thus far we have been discussing what are usually referred to as “**full aperture**” systems, in which the need is to ensure that each beam is not only co-aligned (i.e. going in the same direction) with the others but is also co-axial (i.e. concentric) with them. The remainder of this article will deal with what are referred to as “**tiled aperture**” systems, which are important in long range applications (such as optical communications and laser weapons). Here the beams are co-aligned but there is no attempt to make them co-axial; instead the assumption is that at long range the natural diffraction spreading of the individual beams will ensure that they overlap (to all intents completely), as shown in Figure 13. The tiled approach provides increased flexibility, since the co-axiality condition is not required, and facilitates readily scalable systems. It is also not restricted to spectral combining, and is used in numerous laser weapons programmes.

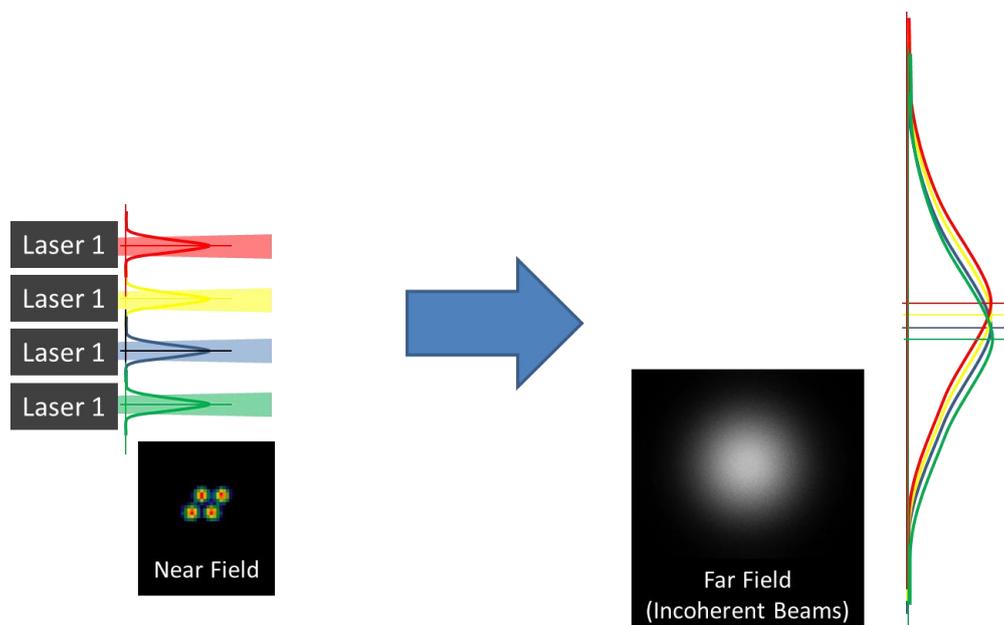


Figure 13: Principle of tiled aperture beam combining, showing far-field overlap of constituent beams

As with the previously discussed full aperture schemes, there are different approaches to tiled aperture beam combining, depending on whether or not the beams are coherent. Referring now to Figure 14, in the general case of incoherent beams (whether the beams are of the same wavelength or not) the far field intensity pattern is simply the sum of the individual intensity patterns (as per the previous illustration). In this case the beam size is fundamentally limited by the diffraction spreading characteristic of the individual apertures. If however the beams are coherent, and their relative phase can be controlled such that the wavefronts effectively form a continuous single larger wavefront (ignore the joints for the moment), as shown in the lower part of Figure 14, then the far field intensity pattern is set by the diffraction limit of this physically larger wavefront. The consequence is that the intensity of the beam increases, because the same power is contained within the reduced diffraction spread of the larger wavefront. Overall, for “*n*” beams the intensity increases as n^2 for the coherent

case, whereas it scales only as n for the incoherent case. It will be appreciated that for a large number of beams the advantages of coherent combining can be profound.

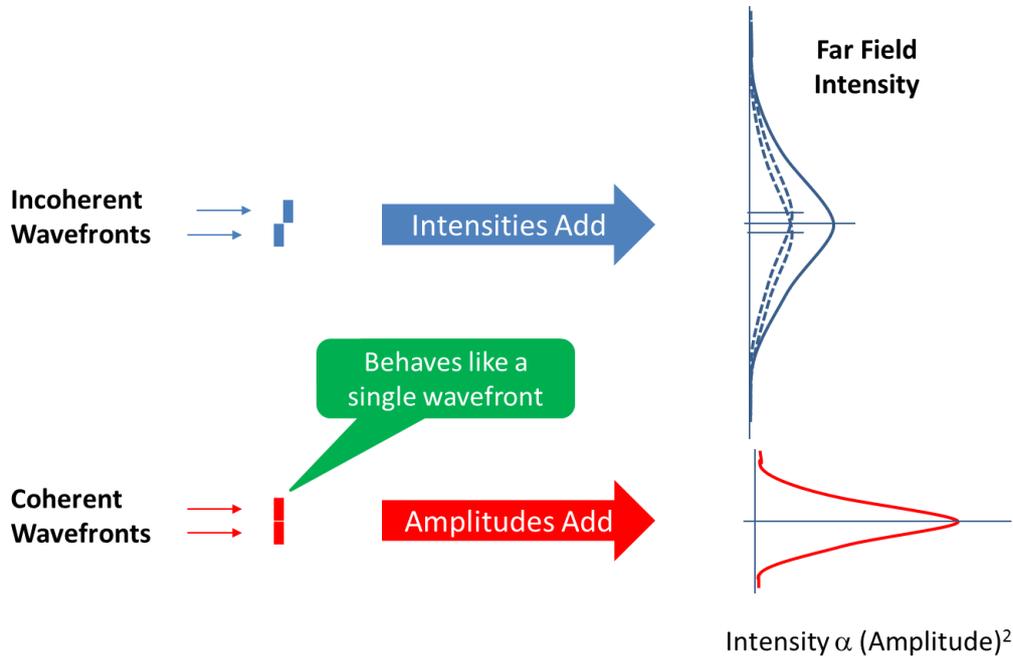


Figure 14: Difference between coherent and incoherent tiled aperture beam combining

In reality of course the wavefront is not continuous but has discontinuities, which may be significant depending on the optical set-up. The effect of these discontinuities is to superimpose “ringing” structures on the far field pattern (which is the Fourier Transform of the near field pattern), which “smears” some of the available power out into side-lobes of the main beam, but these effects are generally small compared with the central beam.

It should also be appreciated that this simple description of a single flat wavefront being synthesised out of a number of smaller ones does assume that the intensity profile of each beam is flat. This condition can be achieved by prior manipulation of the wavefront of each channel prior to the phase alignment but, in general the beams will more or less Gaussian in cross-section and the resultant wavefront will be non-uniform. Without going into the detailed mathematics it can be shown that there is an optimal spacing between the beams (corresponding to the situation where they overlap at their $1/e$ points), which results in a maximum intensity in the far field (and we shall see how this can be achieved later).

Figure 15 (on page 10) introduces some of the practicalities of tiled aperture coherent beam combining. Conventionally, generating mutually coherent beams is achieved by driving a number of power amplifier stages from a single master oscillator, as shown. However, this mutual coherence is not in itself sufficient to provide the constructive interference that will yield the n^2 intensity increase; the beams must also be brought into phase alignment. In the upper part of the figure the relative phase of the beams is measured (we’ll discuss how this can be achieved shortly) and used in a feedback loop to control phase shifting elements that will bring the beams into phase alignment.

In the example shown here 4 beams, in a 2×2 square array, produce the far field pattern shown. This exhibits the side-lobe ringing alluded to earlier, but it should be noted that the intensities of these side-lobes are much less than that of the central peak. The lower part of Figure 15 illustrates the effect

of switching-off the phase controlling feedback; the far field pattern now averages out to the beam size determined by the individual apertures.

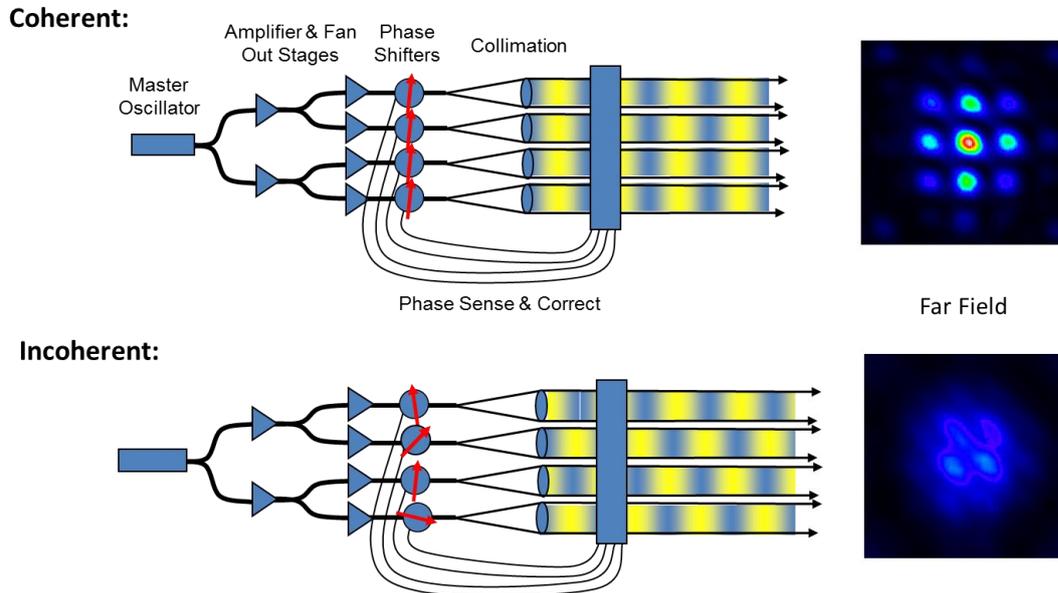


Figure 15: Typical derivation of phase-locked coherent beams in tiled aperture combining

Practical example: Tiled aperture, 19 channel, high power coherent beam combiner

In this section I'll introduce some of the practicalities of tiled aperture coherent beam combining, with reference to an actual system that we have worked on.

The first requirement of these systems is that the beams must be pointing in the same direction to a very high degree of accuracy. In a system intended to generate a beam size of only a few centimetres at a range of several kilometres the need is to ensure that the beams align to an accuracy of the order of just a few microradians. Figure 16 illustrates one means by which this can be achieved, in this case for a series of Gaussian profile beams delivered by monomode optical fibres. We previously alluded to the fact that for beams of this type there is an optimal spacing such that they overlap at their $1/e$ points; but how can this be achieved since the supporting structure of the collimating lenses of any one beam will obviously encroach on the other beams?

This example shows a component that we developed to align a number (up to 19) of Gaussian beams (each of which delivered several hundred Watts). In essence it comprises an array of fused silica plano-convex collimating lenses and a corresponding array of optical delivery fibres, each at the focal point of its collimating lens. The lenses are optically "wring" to a fused silica optical flat, such that they are effectively held in place without any supporting structure (the bond is formed by Van der Waals forces of attraction, there is no adhesive of any kind in the optical path, and fused silica has negligible absorption at the wavelength in question). Each lens therefore collimates its own beam, within the $1/e$ radius, and the remaining part of the beam is collected by the surrounding lenses and directed into a surrounding heat-sink. In order to achieve the required pointing accuracy the fibres in this

component are individually positioned to an accuracy of 1 micron and are glued in place with an ultra-low shrinkage cement.

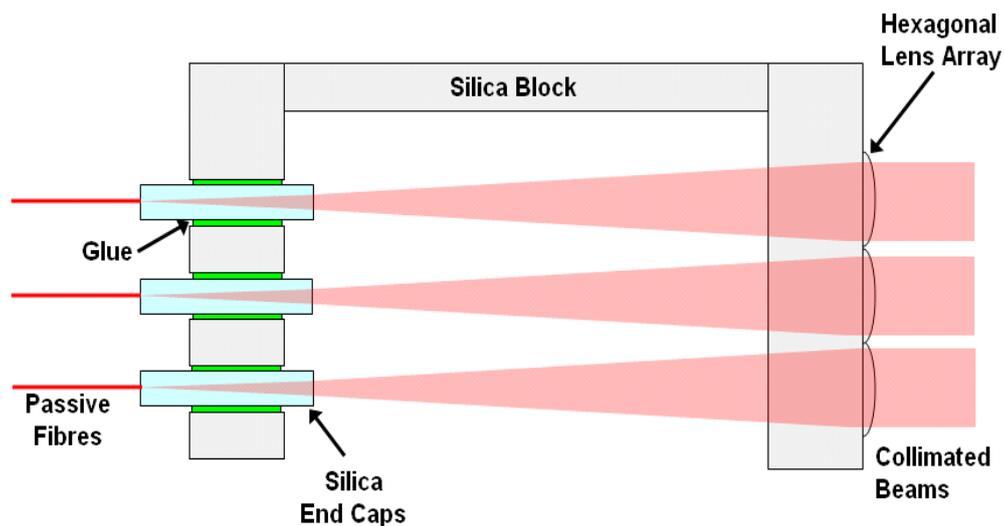


Figure 16: Pure Silica monoblock lens assembly for multiple high power beam alignment

Figure 17 illustrates the scheme which was developed (in association with collaborators at QinetiQ) to sense and correct the phases of the various beams. In essence it involves bleeding off a small amount of the high-power output beam via a high reflectivity beamsplitter and comparing this with a reference wavefront, derived from the master oscillator, in what is basically a Michelson interferometric arrangement. The phase error is encoded as a fringe pattern which is detected and fed back to phase shifting elements in each channel prior to the main amplification stages.

This work was reported in: ***“A multi-channel phase locked fibre bundle laser”***, David C Jones ; Andrew J Turner ; Andrew M Scott ; Steven M Stone ; Roy G Clarke ; Christopher Stace ; Craig D Stacey, *Proc SPIE 7580, Fiber Lasers VII: Technology, Systems, and Applications (February 17, 2010)*.

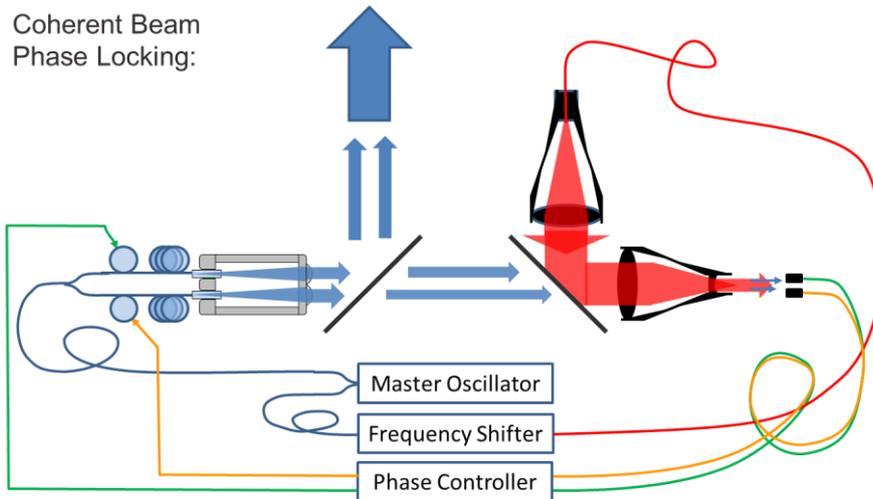


Figure 17: Interferometric method of producing phase-locked coherent beams

The final photograph, Figure 18, shows the practical implementation of the scheme just described in a demonstrator system constructed largely out of COTS components. The laser delivery fibres are seen entering the fused silica monoblock beam combiner & collimator on the left. Next to this is the input for the reference fibre (not connected in this picture). The bundle of yellow fibres delivers the optical error signals from the interferometer out to a series of detectors for processing into the electrical signals for controlling the phase shifting elements in the fibre laser amplifier chains. The three grey boxes in this set-up are cameras that monitor the near and far-field of the beams to provide diagnostic information and are otherwise not functionally important.

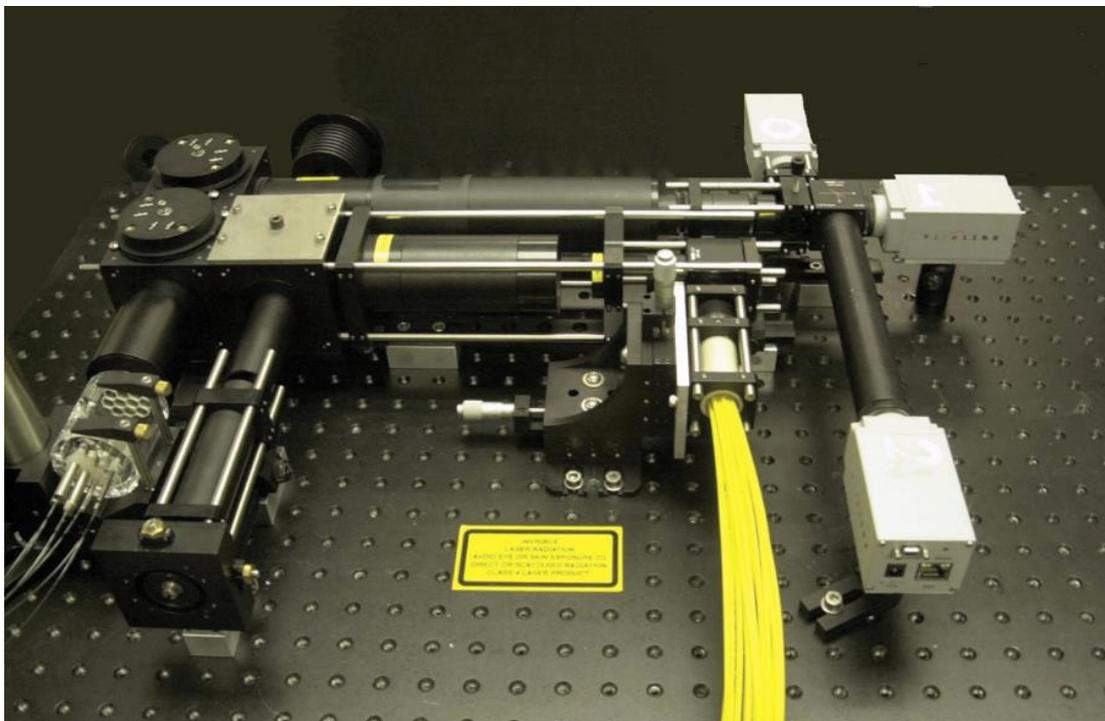


Figure 18: COTS realisation of the high power multi-channel coherent beam combining system

In Conclusion

So, which technique do you use for any given application? It is the usual story of choosing horses for courses; there are some general “rules” that should be considered, summarised in Table 2, but there are likely to be various ways of achieving the aim; and the optimal solution will depend upon the specific requirements of performance, cost, waveband, power, etc. Certainly, it is advisable to look at all the options at the outset.

Table 2: Summary of beam combination methods available for spectral and power scaling

<p>Spectral Scaling:</p> <ul style="list-style-type: none"> • Dichroic <ul style="list-style-type: none"> - simple (for low scaling) - efficient with modern coatings - moderate scaling (complexity!) • Polarizer <ul style="list-style-type: none"> - other options better • Prism <ul style="list-style-type: none"> - simple, robust, efficient - restricted $\delta\lambda$ - good $\Delta\lambda$ - moderate scaling • Grating <ul style="list-style-type: none"> - good $\delta\lambda$ possible (limits $\Delta\lambda$) - multiplexed Bragg gratings - allow tailored λs (and robust) - reasonably good scaling 	<p>Power Scaling:</p> <ul style="list-style-type: none"> • Use Spectral Scaling • Incoherent (but same λ) <ul style="list-style-type: none"> - polarizer (but only x2) - Tiled Aperture • Coherent: <ul style="list-style-type: none"> - polarizers (moderate scaling) - Tiled Aperture (ultimate intensity) <p>N.B. All the above have been used in US LDW demonstrator systems</p>
--	---

*This article is based on material originally presented as a tutorial lecture in the **Innovations Live** section of the **Photonex 2015 Trade Show**.*

If you found this article interesting and would like to know more about what we do, then contact us on info@lumoptica.com.