A Study

of the

Oxygenation

and

Electro-Deoxygenation

of the

Ti-6Al-4V Alloy

A study of the oxygenation and electro-deoxygenation of the Ti-6Al-4V alloy

A. H. Bridges*

● Abstract

This report describes an investigation of the oxygenation and subsequent electro-deoxygenation of the Ti-6Al-4V alloy. Samples were oxidised at 700°C for various periods and subsequently heat treated at 900°C in inert atmospheres, homogenising the oxygen concentration profiles by diffusion. This homogenisation process is modelled by finite element analysis. Optical analysis, energy-dispersive X-ray analysis scans, X-ray diffraction and microscopic Vickers indentation are used to show the effects of the increased oxygen concentration on the distribution of the α & β phases and the hardness of the alloy. The oxygen concentration of samples is measured by elemental fusion. The relative accuracy and reliability of the different techniques are compared.

Following oxygenation, samples were electro-deoxygenated in molten calcium chloride. The resulting hardness and β profiles are compared with predictions made by an appropriate solution of Fick's 2nd Law.

① Introduction

Although it is the ninth most common element in the Earth's crust, titanium is seldom used in large-scale applications due to its cost[1]. The majority of titanium consumption is by the aerospace industry, where the high performance requirements of a turbine blade necessitate such an expensive alloy. The corrosion resistance of titanium is also exploited in the biomaterials sector, while its lightness and strength make it very useful in military applications and the design of modern automobiles, where it is often used in structural reinforcements and moving parts. A common commercial application of titanium is in high quality golf clubs. The extremely harsh conditions of an offshore drilling platform often call for titanium components.

Techniques for the machining and joining of titanium are constantly developing, such as the friction welding technology developed in The Welding Institute, capable of producing a high-quality join in a fraction of the time required by previous techniques[6]. Titanium has also been considered for use in fuel cells and hydrogen storage.

The high cost of titanium is largely due to the expensive process used to extract the metal from its ores (rutile and ilmenite) [2]. This pyrometallurgical process, developed by Kroll in the 1930s, treats the ore to obtain a $TiCl_4$ vapour that is subsequently reduced by magnesium metal to metallic titanium. A similar process was later developed for the extraction of zirconium. Only recently has a more efficient technique been developed, using direct electrochemical reduction of TiO_2 to Ti in molten calcium chloride. This process works by ionising the oxygen and dissolving it in the molten salt. It is hoped that this process may eventually be adapted for use on other metals[22]. A similar process has been reported for the deoxygenation of oxygen-rich Ti. This will be discussed in section ⑧.

Titanium exhibits α and β phases. The harder α phase is hexagonal close-packed, with $a = 2.92$ Å, $c = 4.67$ Å. The more ductile β phase, which is not observed at room temperature in pure Ti, is body-centred cubic with $a = 3.23$ Å. Alloying elements such as Al (face-centred cubic) and V (body-centred cubic) occupy substitutional sites, stabilising the α and β phases respectively. Fe (present as an impurity) also stabilises the β phase. Ti-6Al-4V is widely used in applications ranging from the medical to aerospace because these concentrations of Al & V give proportions of α & β phases that, together, have a desirable balance between mechanical strength and workability.

Oxygen occupies interstitial sites, inhibiting the motion of dislocations, which increases the mechanical hardness of the alloy locally. It is much more soluble in the α phase than in the β, significantly stabilising the α phase. The reason for this is that hcp α-Ti has larger interstices than bcc β-Ti, even though β-Ti has more interstices per atom. Since oxygen occupies interstitial sites, its diffusion in the titanium is considerably faster than the substitutional diffusion of aluminium and vanadium. The α & β are thought not to form as separate grains; instead they appear as distinct regions within grains, although this is still under debate.

* Department of Materials Science & Metallurgy, University of Cambridge

② Preparation of Samples

An ingot of Ti-6Al-4V alloy was sliced to obtain a plate 40 cm × 10 cm × 1 mm. From this plate, samples approximately 10 mm × 5 mm × 1 mm were cut. Each sample had a hole of diameter 1 mm drilled through it, making the samples easier to suspend in the furnace. (See Fig. 1)

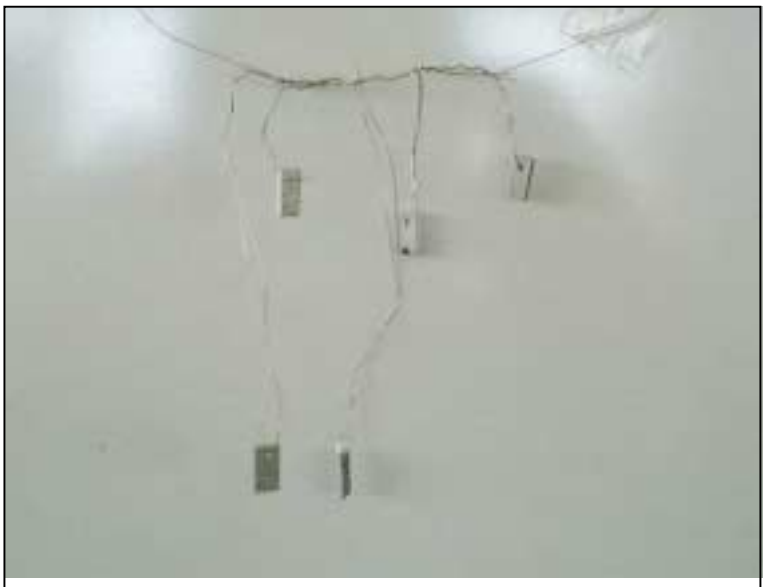

Fig. 1 – 5 samples prepared for suspension in the vertical furnace

The first batch of 11 samples was oxidised in a vertical furnace, programmed to ramp to 700°C at 8°C per minute, followed by 113 hours of dwell time at 700°C, after which the furnace was turned off. The samples cooled slowly in the furnace for the next 12 hours. It was found to be essential that the samples cool slowly, since otherwise the oxide layer would spall off. Further samples were treated similarly, with less or more dwell time at 700°C.

The oxidised samples were wrapped in platinum wire (approximately 8 cm per sample) and sealed individually in small quartz flasks (see Fig. 2). The platinum wire prevented the samples from touching the quartz, since it has been observed that oxygen can diffuse between a sample and a crucible or flask that makes contact with it. Platinum can survive high temperatures and is chemically inert, hence its use here. Unfortunately, Pt wire is very expensive. It would have been better to use Pt foil, thus completely preventing any contact with the samples, but that would have been even more expensive. After initial difficulties, it was found preferable to wrap the Pt around the samples before oxidising them, since the oxide layer was likely to spall off during the wrapping process. The wire did not interfere with the oxidation, and was never observed to adhere to the samples, indicating that the platinum did not diffuse into the samples.

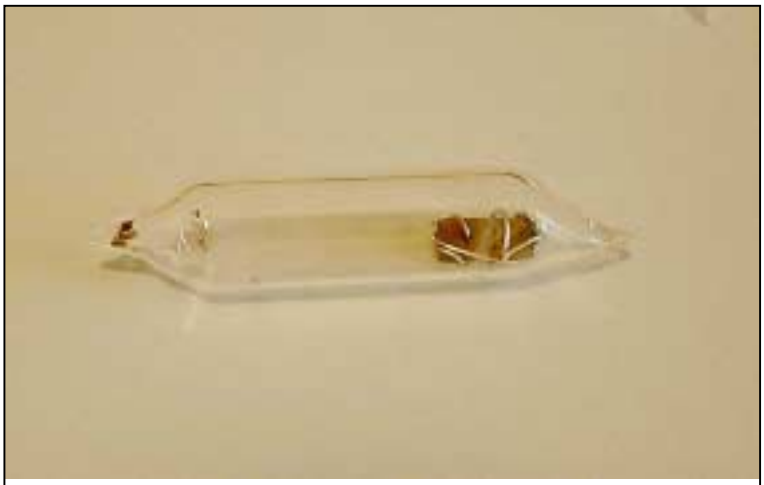

Fig. 2 – An oxidised sample sealed in a quartz flask containing an inert atmosphere

The quartz flasks contained an atmosphere of approximately 0.1 bar argon. They were placed in the furnace, programmed to ramp to 900°C at 8°C per minute, followed by 96 hours of dwell time at 900°C, after which the furnace was turned off. Some of these samples were later returned to the furnace to be homogenised further, since it became apparent that 96 hours at 900°C did not achieve a satisfactory uniformity. Further samples were homogenised for over a week. After homogenisation, most samples were mounted in Bakelite for easier handling in the scanning electron microscope (SEM) and micro-hardness testing machine. In preparation for elemental fusion, some samples had to be removed from the Bakelite and cleaned thoroughly by grinding on silicon carbide paper. Samples intended for electrolysis were not mounted at this point.

The techniques used in the analysis of the samples are described in the subsequent sections. The electro-deoxygenation of oxygenated samples is discussed in section ⑧.

③ Optical Phase Analysis

In scanning electron microscopy the β phase appears lighter in colour than the α, since it contains a higher concentration of the heavier vanadium atoms than is found in α-Ti. (Electronic backscatter is stronger from a heavier atom.) The α phase contains less vanadium, and more aluminium. The relatively less massive Al atoms give it a darker colour in the SEM.

This difference in colour may be used to analyse the phase proportions optically. *Scion Image*, an image-analysis program produced by the Scion Corporation, was used to count the proportions of α and β in selected regions of micrographs, ultimately to obtain the β-fraction profiles of selected samples.

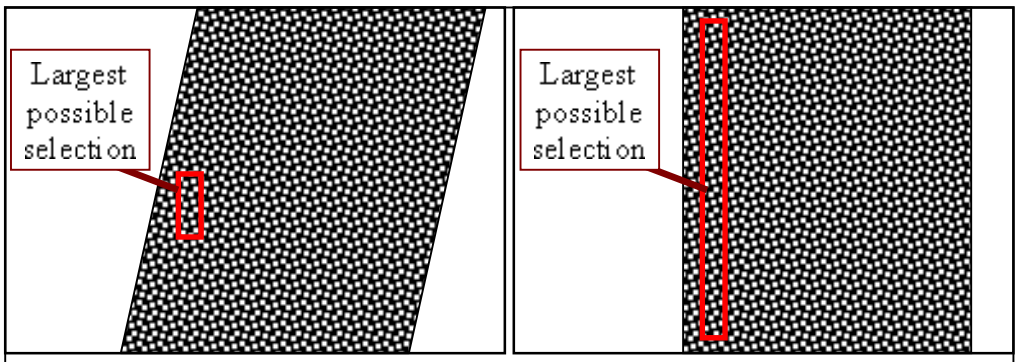

Fig. 3 – The largest possible selection is larger if the sample is oriented parallel to the sides of the micrograph

This software has the capability to display a greyscale image as pure black & white. Furthermore, there is a function which alters the threshold brightness at which a pixel becomes white rather than black. By careful manipulation of this threshold, the α regions may appear black while the β regions appear white. The operator may then select a rectangular region of the image, within which the program will count the black pixels. It also outputs the lengths of the sides of the selected rectangle, from which the total number of pixels may be calculated. The fractions of α and β phase present may easily be calculated from these values.

The statistical validity of a selected region increases with the size of the region. However, to achieve a high-resolution profile, it is necessary to use very thin regions. The micrographs were therefore obtained showing the samples vertically or horizontally on the screen (see Fig. 3), making it possible to select a region that is thin and parallel to the sample edge, yet large.

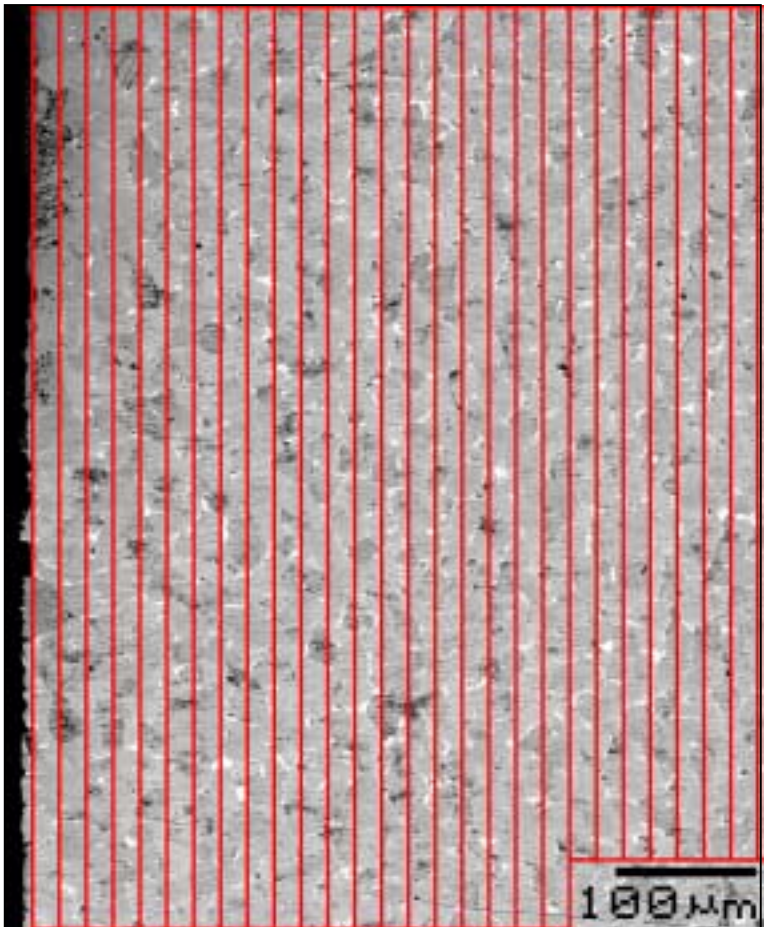

Fig. 4 – Samples were optically analysed as a series of rectangular sections, as represented by these red rectangles

Unfortunately, it was found that if the image is rotated by any graphics software, a compression artefact causes the white regions to become off-white, causing the measured proportion of β-phase to decrease randomly. To avoid this problem, fresh micrographs were obtained in which the sample was oriented appropriately. The selected regions were all 25 pixels wide and as tall as the micrograph (usually approximately 700 pixels). Selected regions were immediately adjacent to each other so that every pixel of the micrograph was counted in exactly one region, as seen in Fig. 4, in which the red rectangles represent selected regions.

The errors involved in this technique are difficult to quantify. During the procedure, it was noticed that the measured fraction of β as a function of threshold level seemed somehow to be discrete rather than continuous. Effectively, when altering the threshold level, the black areas

increased stepwise, not gradually. Small changes in the threshold level therefore did not alter the measured fraction of β. The operator could therefore be reasonably certain, when repeating a measurement, of getting exactly the same result, assuming the threshold level had been well chosen. This would tend to reduce the random error of the technique.

The greatest source of error comes from any topographical features in the micrograph. Although the compositional imaging mode of the SEM eliminates most of these, a scratch would still appear in the micrograph as a dotted line. Despite careful selection of the threshold level, these lines would occasionally be counted as β phase. The edges of a sample also appear white. Measurements taken at an edge are therefore usually more inaccurate than those from the middle. In most cases the β fraction would be zero at the edge, especially in an oxidised sample that had not been homogenised, so the measured value (often over 20%) would be ignored. The separate assumption is made that the observed area represents the whole sample. The error bars plotted on the graphs of the β profiles have consequently been calculated from the spread in the results.

Fig. 5 β profile in an as-received sample

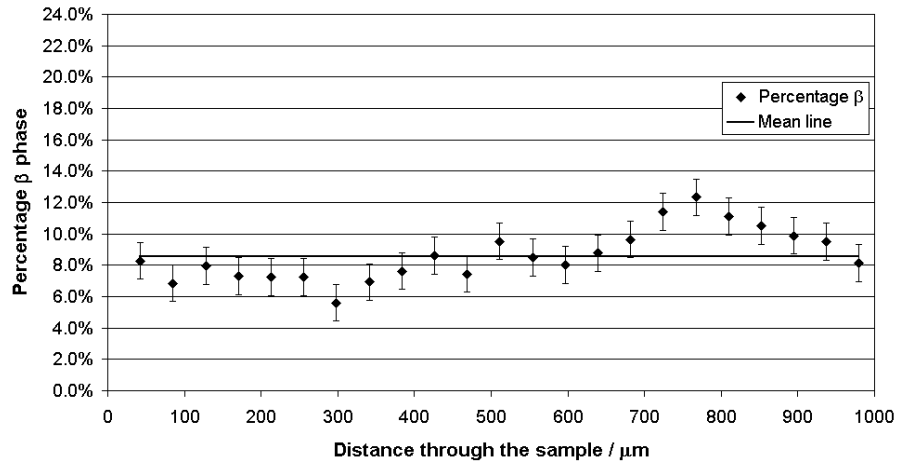

Fig. 5 shows that the proportion of β phase is highest in an as-received sample, and that it is uniformly distributed, within a reasonable margin of error.

Fig. 6 β profile in an oxidised sample

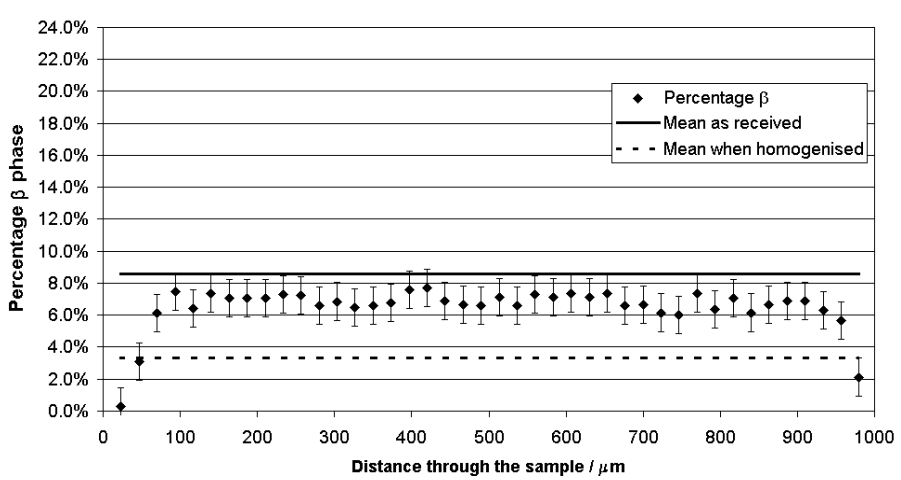

In a sample oxidised but not homogenised (Fig. 6), the central regions have a uniform, slightly lower proportion of β phase. This is probably due to the effects of the heat applied during oxidation rather than any increase in oxygen concentration, since the oxygen profile is unlikely to be so uniform, and oxygen should not be able to reach the centre by diffusion quickly enough.

Fig. 7 β profile in a homogenised sample

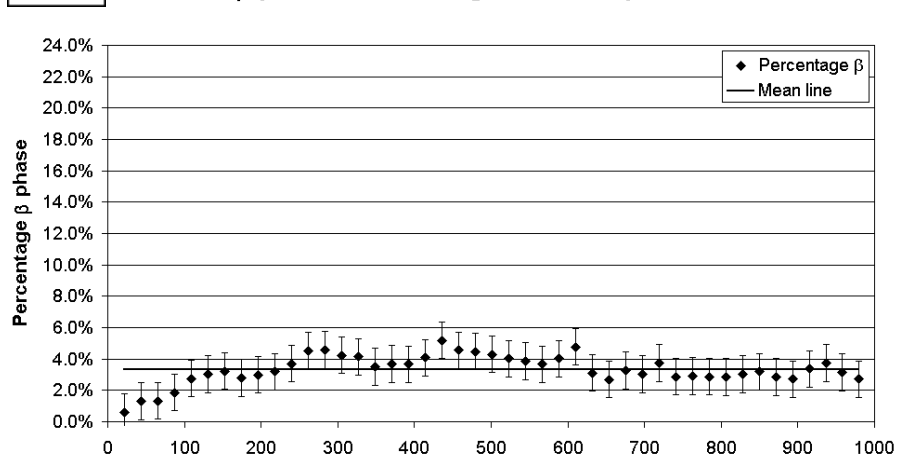

There is considerably less β phase near the edges where the oxygen concentration is high. The oxygen has stabilised α-Ti, causing this decrease in β phase.

The mean β fractions for the as-received and homogenised samples are displayed on the graph of the oxidised sample (Fig. 6) to facilitate comparison.

The homogenised sample has a fairly uniform concentration distribution (Fig. 7), which is much lower than in the as-received sample.

This demonstrates that oxygen stabilises and thermodynamically favours the α phase, since these profiles match the expected forms of the oxygen concentration profiles of the samples. When compared with the hardness profiles, the same match is found, as will be shown in section ④.

The exact oxygen concentration at a point cannot be measured directly, for reasons discussed in section ⑤. However, by elemental fusion, the bulk oxygen concentration of a sample may be found and compared with the mean β fraction. A series of six samples was oxidised for different times, ranging from a day to a week, ensuring a wide range of oxygen concentrations. One further sample was not oxygenated at all. All seven samples were homogenised as described in section ② for 184 hours at 900°C, 115 hours at 700°C and 105 hours at 900°C, including the un-oxidised sample. It was important to homogenise the un-oxidised sample in the same way as the others, since this ensured that all factors would be identical in each sample apart from the oxygen content. During oxidation, the samples were heated; this could cause differences in their microstructure that would not be directly relevant to the oxygen content. The extremely long homogenisation period was intended to render any such differences negligible.

Fig. 8 β profiles in samples oxidised for different periods and subsequently homogenised

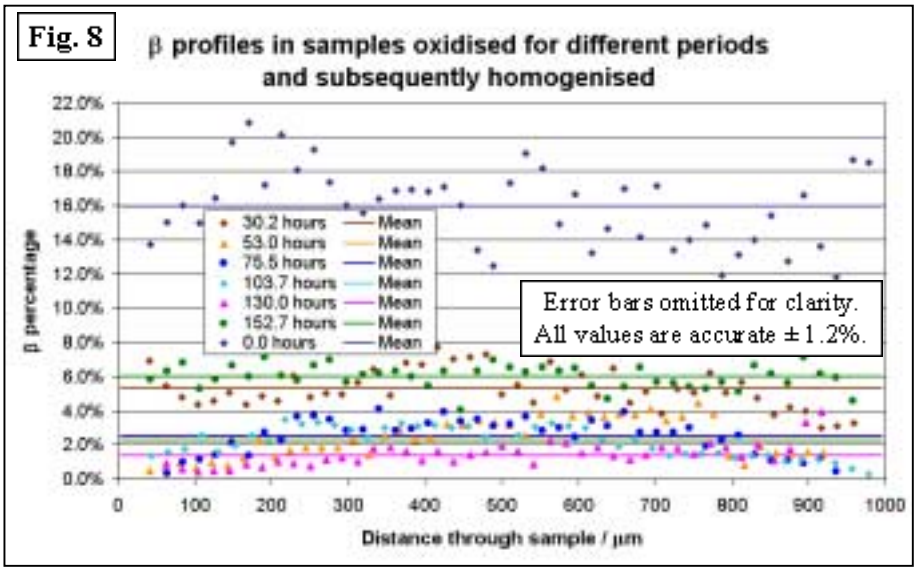

When these samples were ready, they were mounted in Bakelite. Their hardness and β fraction profiles were measured (Fig. 8), after which they were un-mounted by shattering the Bakelite in a vice. Each sample was then tested in an Eltra OHN2000 oxygen analyser. This analyser uses an elemental fusion method to find the weight percentages of oxygen, nitrogen and hydrogen in a sample. Samples are wrapped in baskets made from some appropriate metal (for Ti samples, a Ni basket is recommended) and placed in a graphite crucible in a miniature furnace. As the samples are incinerated, the machine measures any electric currents or potentials through the system, calculating from these the mass of oxygen, nitrogen or hydrogen present. This is output as a weight percentage. In the process, the sample is destroyed. Unfortunately, it was not possible to make multiple identical samples, which would allow repetition of tests, since this would have taken far too much time. Consequently, when the machine failed to give an output for a certain sample, there was no way to try again.

The errors associated with this machine are variable. Although the machine was carefully calibrated before use, its accuracy may be decreasing unnoticed. To quantify the errors, measurements were made of several identical samples in their as-received state. From seven such measurements, the mean value was calculated as 0.1310 ± 0.02 wt-%. This error range is henceforth used as the error range for all data obtained from the Eltra OHN2000. (See Table 1.)

Sample number	Oxidation time / hrs	[O] wt-%	β vol-%
11	30⅙	0.99 ± 0.02%	5.3 ± 1.2%
12	53.0	1.15 ± 0.02%	2.4 ± 1.2%
13	75½	1.48 ± 0.02%	2.5 ± 1.2%
14	103⅔	1.54 ± 0.02%	2.2 ± 1.2%
15	130.0	[Failed]	1.5 ± 1.2%
16	152⅔	1.08 ± 0.02%	6.0 ± 1.2%
18	0.0	0.28 ± 0.02%	15.9 ± 1.2%

Table 1 – Oxygen contents and β fractions of samples

It is intuitively expected that the oxygen content should be increased for a sample oxidised for a longer time. For sample 16, this is not observed. However, micrographs of sample 16 showed that it was only 950 ± 10 μm thick, rather than 970 ± 10 μm as observed in most other samples. It is reasonable to suggest that the oxide layer developing on sample 16 spalled off, taking most of the accumulated oxygen with it and reducing the sample's thickness. The tests performed on sample 16 all give results consistent with an oxygen content only slightly increased from its as-received level (as discussed below, see also section ④).

Although sample 18 should have the same oxygen content as the as-received samples (0.1310 ± 0.02%), it does not. This could be due to an erroneous result from the Eltra OHN2000, but it seems more likely that the quartz flask in which the sample was sealed contained a small quantity of oxygen that entered the sample. It would not necessarily form an oxide film; the very low partial pressure of oxygen in the flask would lead to an oxidation rate sufficiently low to allow the oxygen to diffuse into the metal rather than accumulating in a surface layer. This would also account for the apparently greater hardness and lower β fraction observed at the surface of each homogenised sample.

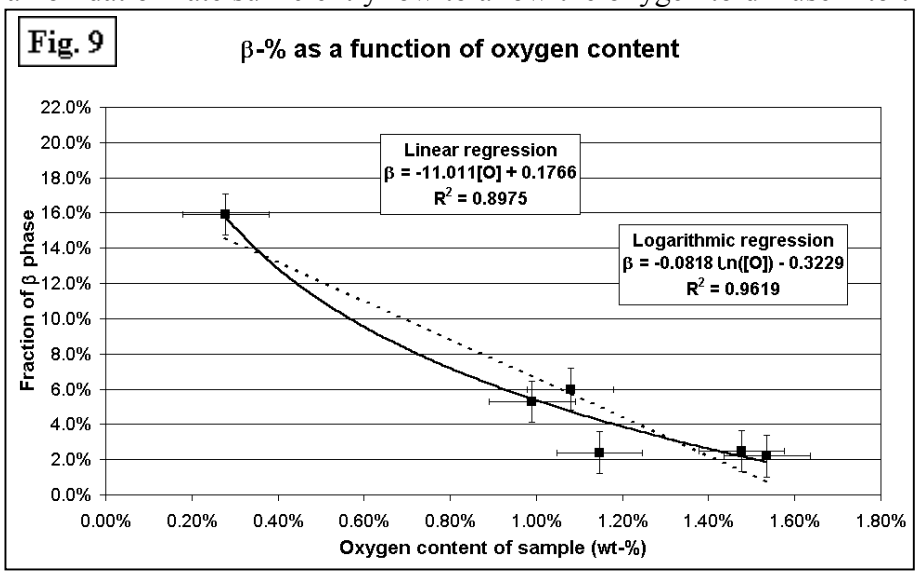

Fig. 9 shows the relationship between the volume fraction of β and the oxygen concentration. The data fits a logarithmic trend-line extremely well. However, it is important to note that such a trend-line predicts $\beta \to \infty$ as $[O] \to 0$. This is, of course, ridiculous. The linear trend-line is also a good fit and does not fail so catastrophically at [O] = 0, although it may not be accurate for low β fractions. Comparison with a phase diagram for Ti-6Al-4V – O would probably reveal that the relationship should be linear within a certain range of [O].

The linear relationship, $\beta_{vol-\%} = -11.011[O]_{wt-\%} + 0.1766$, may be used to predict the oxygen content of sample 15, for which the elemental fusion technique failed. From the average β fraction (1.5 ± 0.5 vol-%), the predicted oxygen content is 1.47 ± 0.05 wt-%. From the same datum, the logarithmic trend-line ($\beta_{vol-\%} = -0.0818 \ln([O]_{wt-\%}) - 0.3229$) would predict an oxygen content of 1.61 ± 0.08 wt-%. Sample 15, having been oxidised for 26⅓ hours longer than sample 14, might be expected to have a significantly higher oxygen content, as predicted by the logarithmic function. The prediction of the linear function would only fit well with the other data if it is assumed that some of the oxide layer was lost before homogenisation took place, causing a decrease in the oxygen content.

Comparing the R^2 values of the trend-lines and their predictions for high & low β fractions, it may be concluded that the logarithmic relationship is more accurate across this range of values, but that it is inaccurate for a low oxygen content, over which range ([O] < 0.4 wt-%) the linear relationship is more reliable. This could more effectively be tested in a similar experiment involving many more samples, achieving improved resolution. Samples with lower oxygen contents than the as-received state could be made by electro-deoxygenation, if it may be assumed that the same relationship would be valid for an electrolysed sample. Electro-deoxygenation is investigated in section ®.

④ Micro-hardness Tests

As has been explained in section ①, dissolved oxygen strongly affects the hardness of the alloy since it inhibits the movement of dislocations. The hardness profile may therefore be used as a guide to the oxygen concentration profile, although the relationship between the two is not linear (this will be discussed later). α-Ti is generally much harder than the ductile β-phase.

It is believed that the concentration of oxygen in β-Ti is low enough not to affect its hardness significantly. In α-Ti, the higher solubility of oxygen allows more oxygen to dissolve in the alloy, so the hardness is affected by it. The overall hardness of the sample will be a function of the hardness of α-Ti, the hardness of β-Ti and the proportions of both (probably not a linear function). The β-Ti regions are up to 10 μm in size (in a homogenised sample – they are much smaller before homogenisation), as are the indents. This makes it fairly unlikely that the measured hardness would refer to pure β-Ti.

Fig. 10

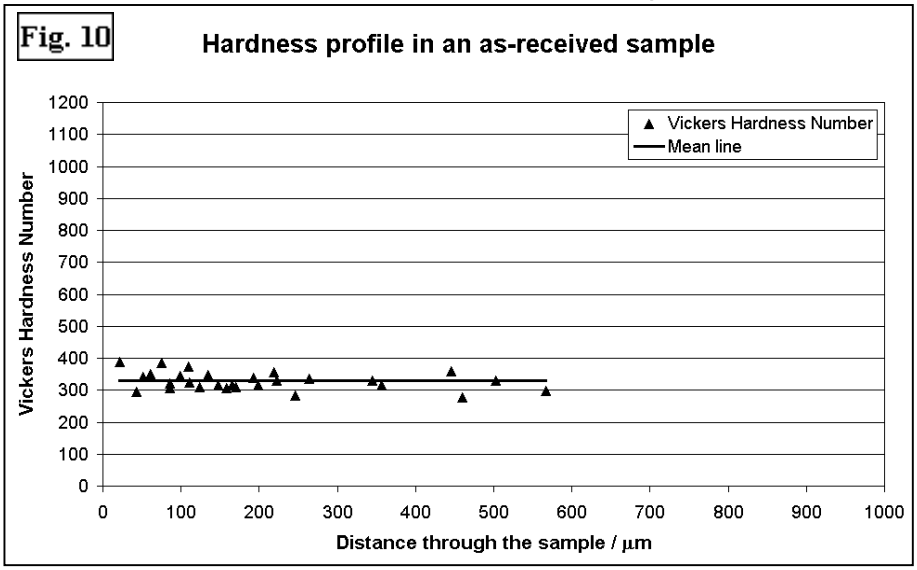

Fig. 11

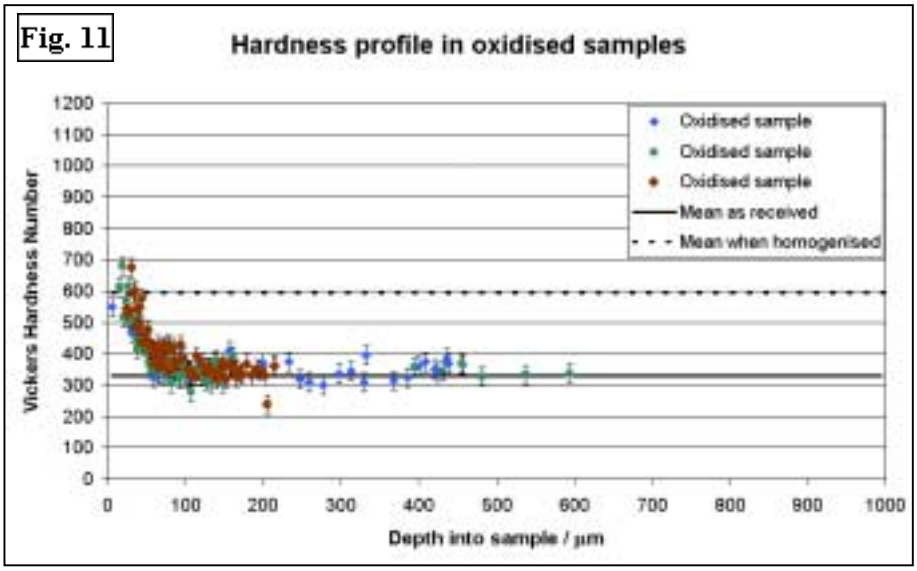

A Mitutoyo HK3 Micro-hardness Tester was used to measure the hardness of the samples at intervals of a few μm, using a load of 50g. Although greater loads generally lead to improved accuracy, it was found that the hardness profile in oxidised samples was so steep that a large indent would be deformed by the non-uniform hardness even within such a small area. This problem was minimised by use of the lightest available load. (The Mitutoyo HK3 offers lower loads, but these were found to be defective.)

Fig. 10 shows that the hardness has a uniform value of approximately 329 across the as-received sample. This was verified by placing the original large plate of Ti-6Al-4V in a full-size hardness tester. This gave readings of 346, 330, 325 and 325 at four random positions. These micro-hardness results are, therefore, well supported.

Fig. 12

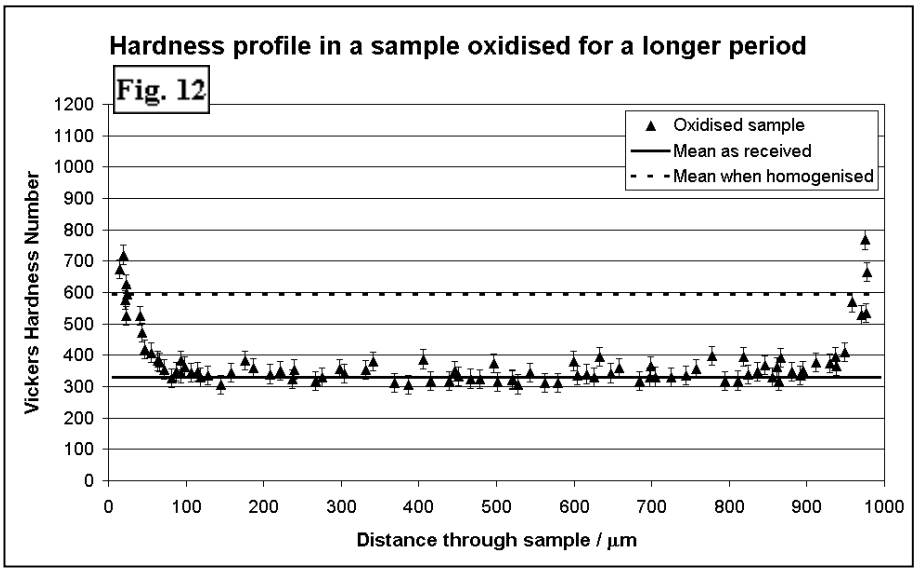

Fig. 11 displays the hardness profiles of three samples that were oxidised for 113 hours. It shows decreasing hardness across the first 120 μm, beyond which depth their hardness is comparable to the as-received state. It may be concluded that the oxygen has not diffused much deeper than 120 μm.

A separate sample was oxidised for 153 hours, producing the profile shown in Fig. 12. Its hardness follows an identical form across the first 120 μm to the other oxidised samples. Although the maximum hardness observed (on the surface) is greater than in the other samples, there is no evidence that the oxygen has penetrated any deeper during the extra 40 hours of oxidation time.

Fig. 13

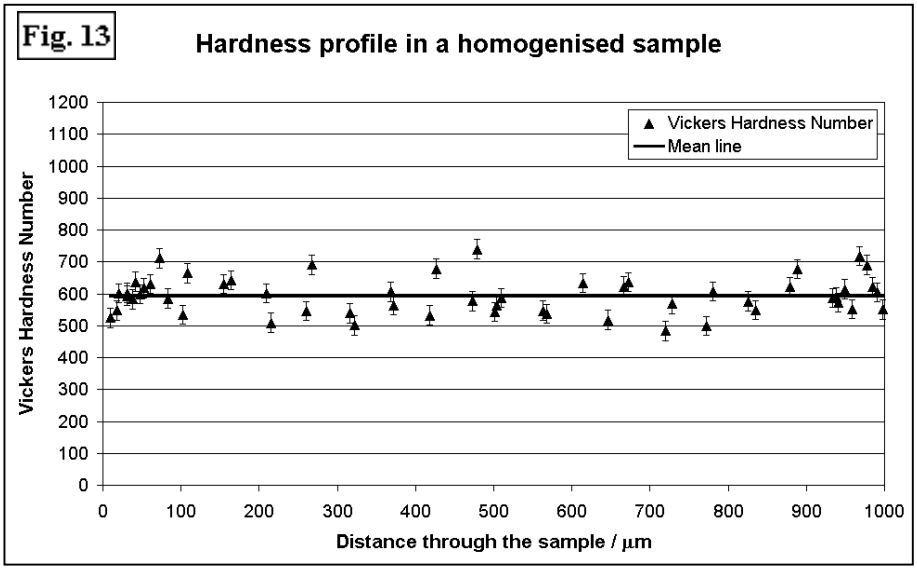

Fig. 14

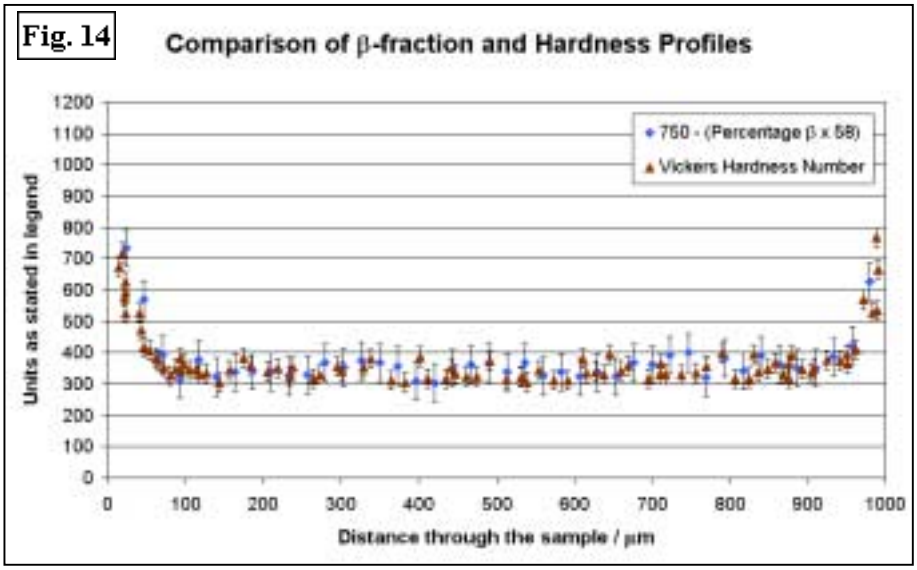

The homogenised sample has a lot of scatter in its results (Fig. 13), but shows an average hardness of 594, considerably harder than in the as-received state (Fig. 10). The hardness is uniform across the sample, to within a reasonable margin of error. During the measuring of the hardness profiles, it was observed that the samples sometimes developed microscopic cracks locally when indented. Such indents were, of course, omitted from the profiles, but the cracking shows that the samples were becoming very brittle after oxygenation. When removed from the quartz flasks, some homogenised samples were found to have fallen apart, possibly by brittle fracture induced by some small stress. This would be rather difficult to study.

These findings compare very well with the optical analyses of the samples. The trends observed are exactly the opposite of each other. Fig. 14 shows a β-profile inverted, superimposed on the hardness profile for the same sample. The match is exact, well within the error range.

From samples 11-16 & 18, described in section ③, a relationship may be found between the sample's hardness and its oxygen concentration. This relationship has already been described for commercially pure Ti, but it may well be different for this alloy, since the α / β distribution is completely different. Fig. 15 shows the variation in hardness of the samples with different oxygen contents.

The oxygen concentrations of these samples were measured by elemental fusion, as described in section ③, and are compared with their Vickers Hardness numbers in Table 2 and Fig. 16. Sample 16 has an average hardness much lower that would be expected from its oxidation time, corroborating the theory that its oxide layer has spalled off, reducing the eventual oxygen concentration (discussed in section ③).

Fig. 15

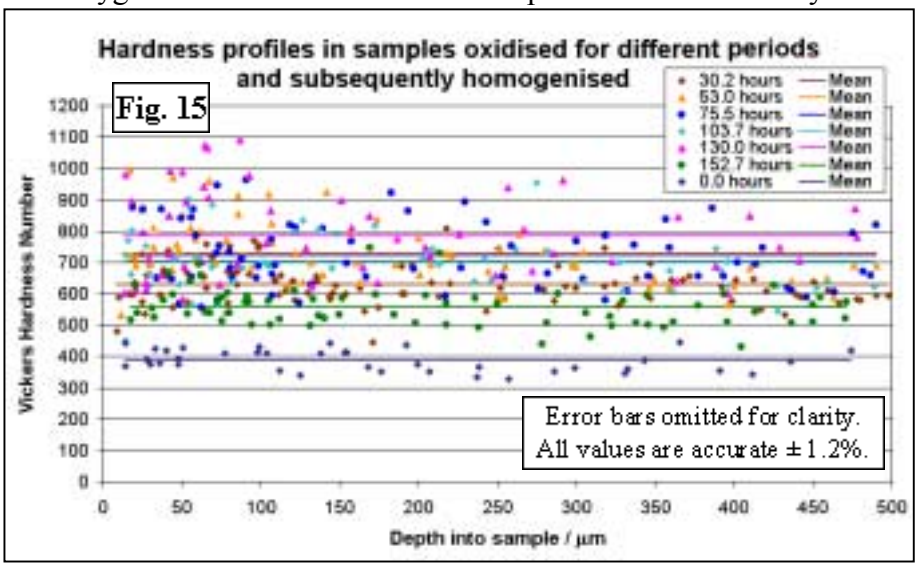

The hardness / oxygen concentration relationship found for CP-Ti[1] was approximately linear, with the power of [O] slightly greater than 1. For the Ti-6Al-4V alloy, a similar power-law

fits the data well, shown in Fig. 16. Such a trend-line incorrectly predicts VHN = 0 at [O] = 0. The linear trend-line is not quite such a good fit across most of the range investigated, but is more accurate as [O] → 0.

The power-law shown in Fig. 16 uses $[O]^{0.37}$. This is a lower index than has been found for CP-Ti. It is not unexpected that the hardness' dependence on the oxygen concentration is much stronger, since the softer β phase is present in the Ti-6Al-4V alloy but not in CP-Ti (at room temperature), and the oxygen affects the phase fractions, as discussed in section ①.

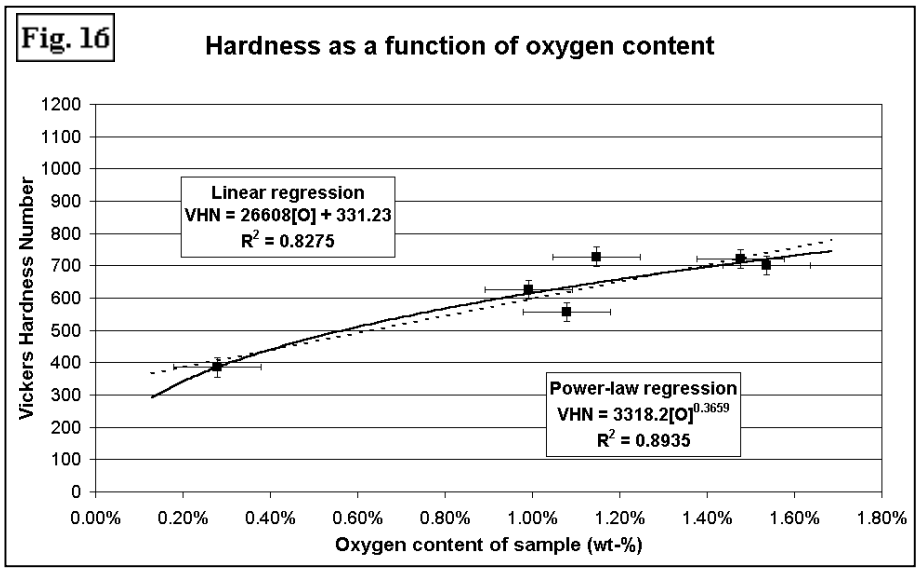

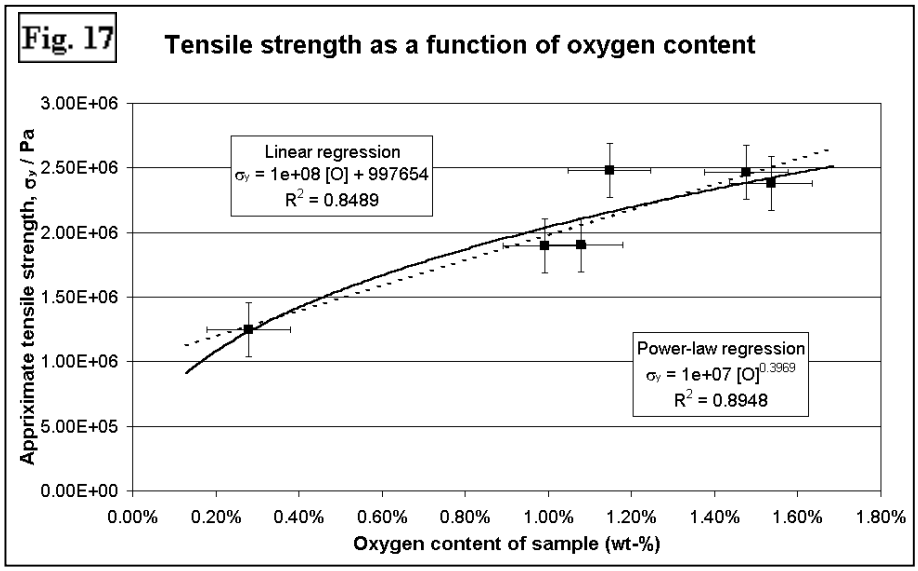

Since the tensile strength, σ_y, may be estimated from the hardness, this data may be used to relate it to the oxygen concentration. The effect of the oxygen concentration on the tensile strength is shown in Fig. 17.

The results show that the oxygen concentration can have a very significant effect on the strength of the alloy. If carefully oxygenated, components made of oxygenated Ti-6Al-4V alloy could ideally be used in situations where the tensile strength is required to be high. The oxygenated alloy is, unfortunately, much more brittle, which makes it rather less useful.

⑤ Energy-dispersive X-ray analysis scans

The energy-dispersive X-ray analysis (EDAX) function of a JSM-5800LV electron microscope was utilised to analyse the chemical composition of the samples. X-rays are emitted when the electron beam electrons excite an inner-shell electron in the sample; the resulting ion is unstable, and emits an X-ray of a characteristic wavelength when it decays. Analysis of the energies of these X-rays gives the chemical composition of the sample at the point where it is scanned.

Initial EDAX tests of a homogenised sample gave spectra similar to those shown in Fig. 18. Analysis of these spectra revealed the composition of the phases (see Table 3).

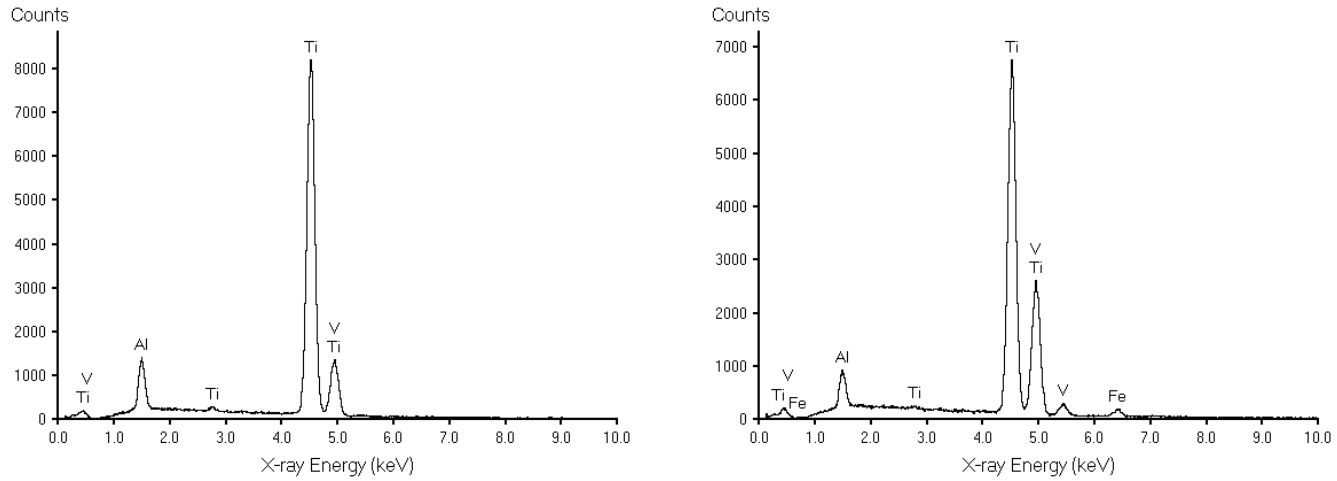

Fig. 18 – EDAX spectra of the α phase (left) and the β phase (right)

	α-Ti			β-Ti					
	wt-% #1	wt-% #2	**Mean wt-%**	wt-% #1	wt-% #2	wt-% #3	wt-% #4	wt-% #5	**Mean wt-%**
Ti-K	91.09 ± 0.53	90.77 ± 0.53	**90.93 ± 0.23**	71.43 ± 0.49	72.05 ± 0.48	72.05 ± 0.48	74.42 ± 0.50	71.42 ± 0.48	**72.27 ± 1.24**
Al-K	5.44 ± 0.10	5.47 ± 0.10	**5.46 ± 0.02**	3.31 ± 0.05	3.33 ± 0.05	3.33 ± 0.05	3.68 ± 0.09	3.24 ± 0.05	**3.38 ± 0.17**
V-K	3.33 ± 0.18	3.59 ± 0.18	**3.46 ± 0.18**	21.89 ± 0.38	22.02 ± 0.25	22.02 ± 0.25	19.56 ± 0.37	22.45 ± 0.25	**21.59 ± 1.15**
O-K	0.00 ± 0.01	0.00 ± 0.01	**0.00 ± 0.00**	0.00 ± 0.01	0.00 ± 0.01	0.00 ± 0.01	0.00 ± 0.01	0.00 ± 0.01	**0.00 ± 0.00**
Fe-K	0.15 ± 0.13	0.17 ± 0.13	**0.16 ± 0.15**	3.37 ± 0.17	2.59 ± 0.30	2.59 ± 0.30	2.34 ± 0.30	2.89± 0.31	**2.76 ± 0.39**

Table 3 – Composition of the phases on Ti-6Al-4V after homogenisation

The presence of Fe was not expected; it is an impurity. Both the k-peak and l-peak of Fe were observed clearly, so there is little doubt about the element's identity. In the α-phase regions, the quantity of Fe was so low as to be of the same order of magnitude as the error-range (e.g. 0.16 ± 0.15 wt-%), so it is not possible to be sure how much Fe is present in the α-phase. Since Fe at room temperature is body-centred cubic, like V and β-Ti, it is likely that the Fe is stabilising the β-phase. Before or during the homogenisation process, the Fe has diffused into the β regions where it can alloy substitutionally with less inherent strain from lattice mismatch.

The major disadvantage of this technique is that the L-peak of vanadium almost exactly corresponds to the K-peak of oxygen. It is therefore impossible, with the energy resolution available, to distinguish between oxygen and vanadium. Ideally, a line would be scanned through the sample, recording the oxygen concentration at each point, getting an oxygen profile instantaneously. This is not possible, hence the use of the elemental fusion technique in sections ③ & ④.

Line-scans were obtained of homogenised, oxidised and as-received samples. They show the concentration profiles of Ti, V, Al and Fe.

The profile of a homogenised sample (Fig. 19), shows most clearly that there are distinct regions of low Al-content (the β phase). In these regions, there are significantly greater proportions of V and Fe - the Fe is found mostly in the β phase, as discussed above. This phase also contains significantly less Ti.

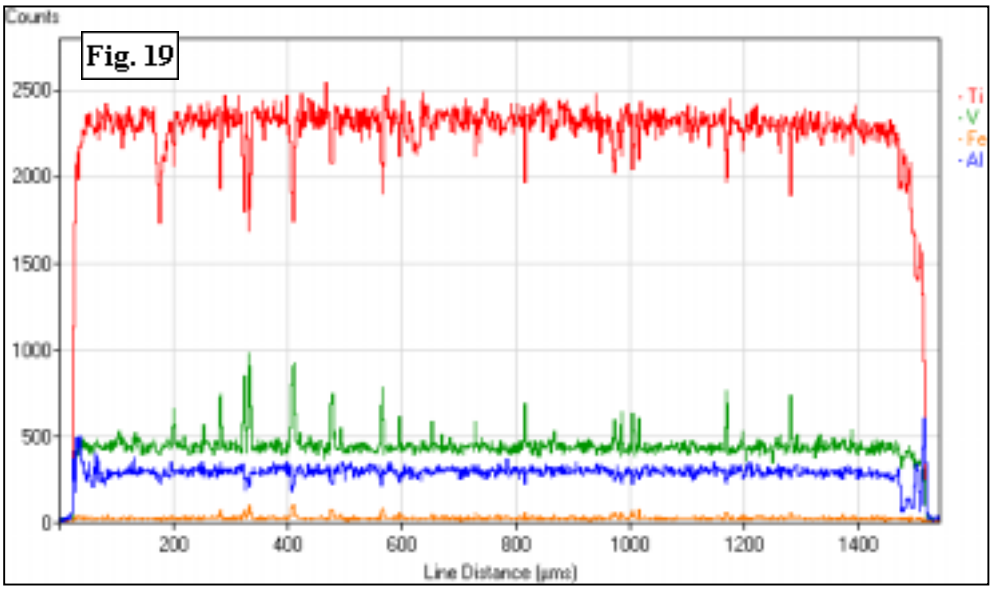

The oxidised sample (Fig. 20) shows similar features, but the phases are much smaller. (The micrographs, Figs. 23 & 24, demonstrate this visually.) This shows that the homogenisation process tends to encourage the phases to separate, forming larger regions. The overall composition, of course, cannot be significantly changed by the homogenisation since this process is conducted in a sealed flask containing an inert atmosphere. The edges may be slightly affected by any traces of oxygen in the flasks, as discussed in section ③.

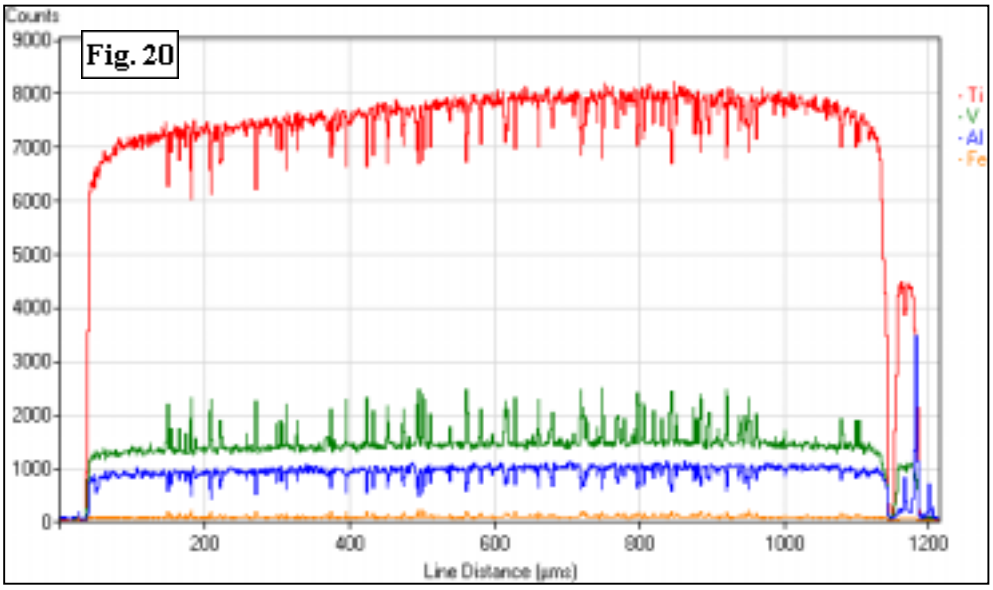

In the as-received sample (Fig. 21), the features appear again, but the composition difference between α and β is much less significant. The increased oxygen content of the treated samples has decreased the proportion of β-Ti in favour of α-Ti (see section ③), so the vanadium and iron are more concentrated into the remaining β. The treatment also allows the elements to diffuse to their preferred phases, emphasising the difference between the α and β. It appears that the proportion of Ti is decreasing throughout the scan. On closer examination, all the elements are seen to show this trend, relative to their average levels, so the cause is probably a time-dependent decrease in the total number of X-rays produced.

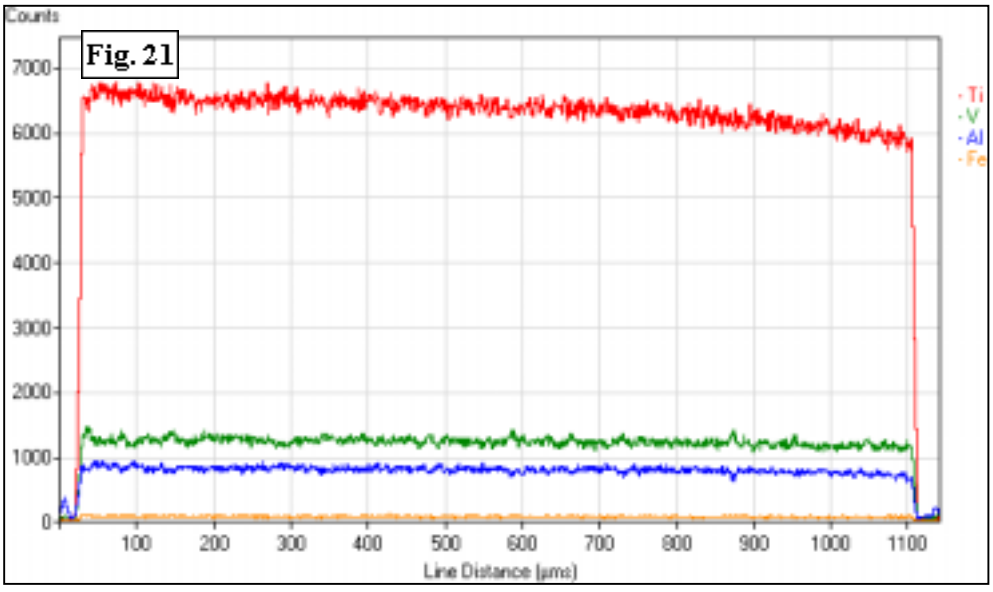

The investigations detailed in sections ③ & ④ required that a sample be homogenised without first being oxidised. Fig. 22 shows its EDAX spectrum. The composition differences between the α & β phases are significantly more pronounced than in the as-received sample, showing that, as discussed briefly in sections ③ & ④, the homogenisation allows the Al & V to diffuse into their preferred phase regions.

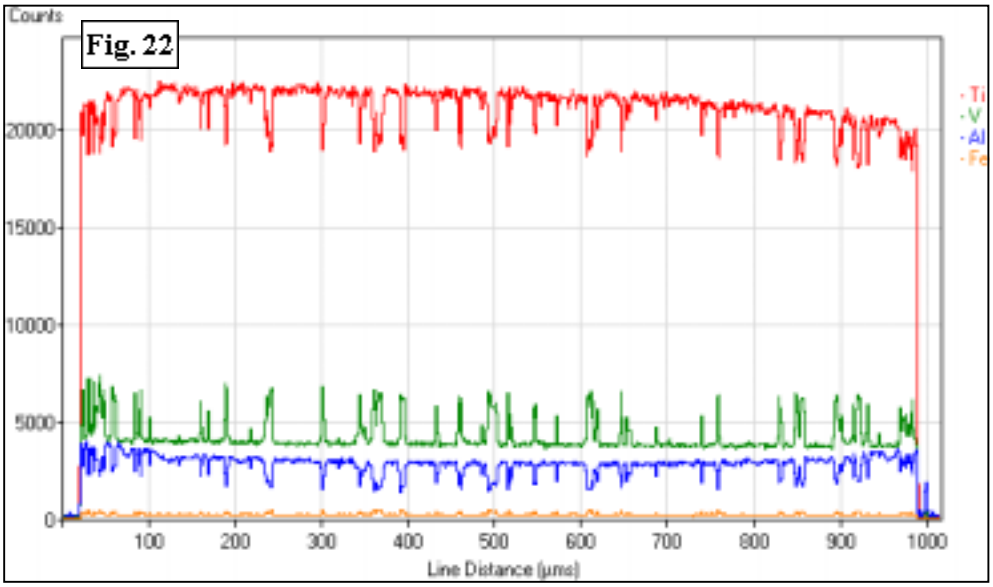

Fig. 22

These micrographs (Figs. 23-26) show the different microstructure scales of the samples.

The β phase regions are clearly much larger in the homogenised sample, of the order of 10 µm. In the as-received and oxidised samples, they are less than half this size. This is caused more by the heat treatment, which allows annealing and phase aggregation, than by the increase in oxygen concentration, which does not directly influence the size of the phase regions.

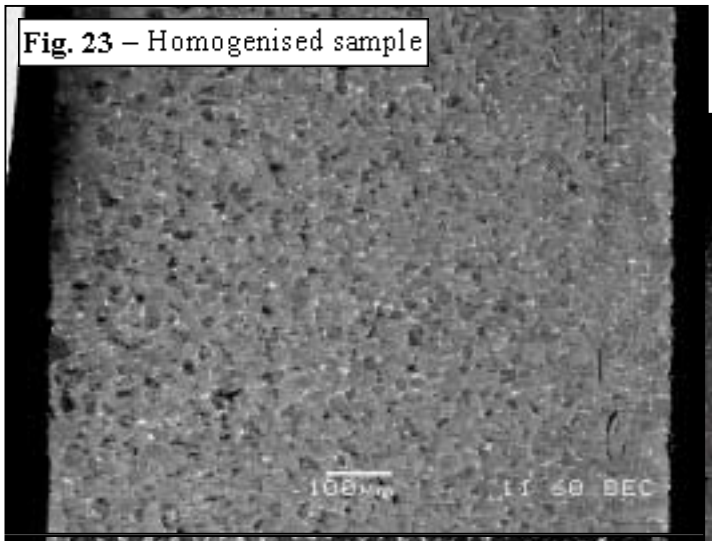
Fig. 23 – Homogenised sample

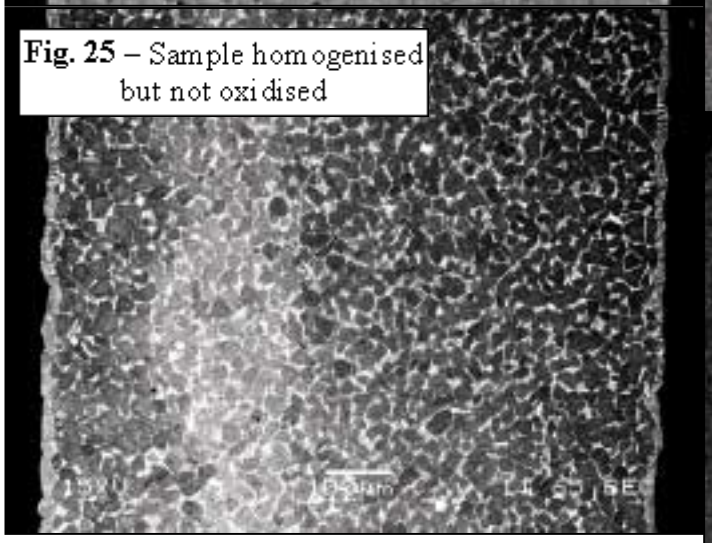
Fig. 25 – Sample homogenised but not oxidised

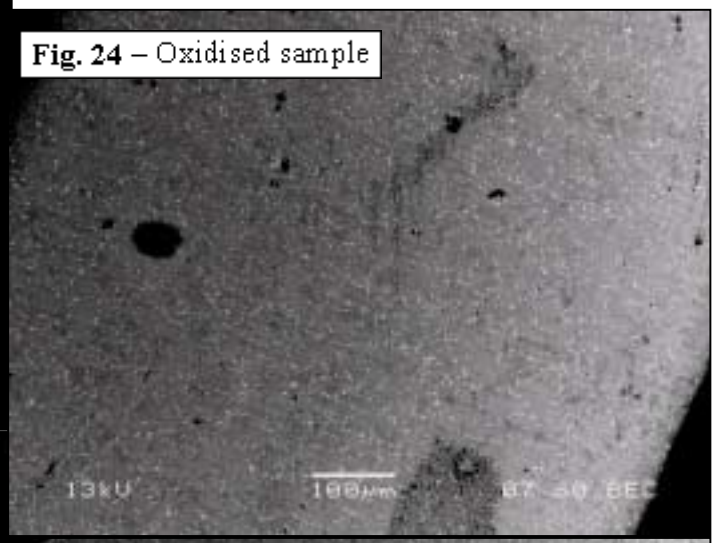

Fig. 24 – Oxidised sample

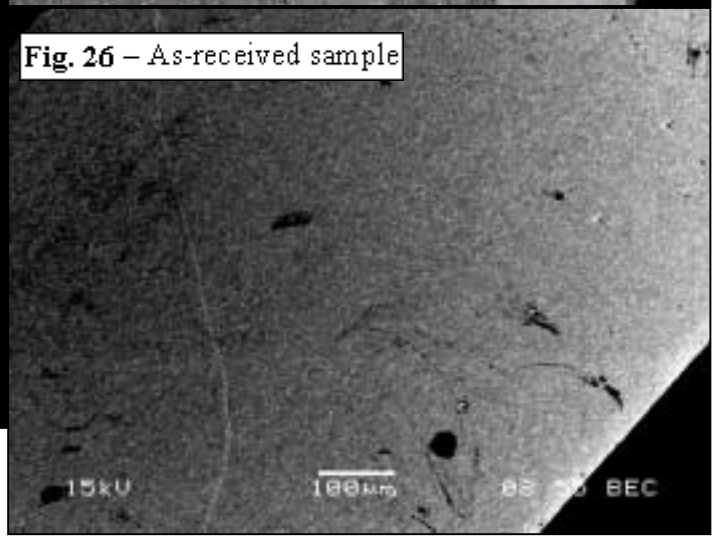

Fig. 26 – As-received sample

The effect of the heat treatment on the distribution of the Al and V between the α and β phases will be discussed in section ⑥.

⑥ X-ray Diffraction

The data obtained by optical phase analysis (see section ③) demonstrates the decrease in the fraction of β-Ti when a sample is oxygenated. This may also be observed by the standard technique of X-ray diffraction (XRD), independently verifying the results of the optical analysis.

	As-received sample	Homogenised sample
Diffractometer	Vertical diffractometer	Vertical diffractometer
Generator settings	10 kV, 10 mA	10 kV, 10 mA
Divergence slit / °	0.50	0.50
Receiving slit / mm	0.20	0.20
Radiation	Cu-Kα	Cu-Kα
Scan range (2θ/ °)	20.000 – 140.015	20.000 – 140.000
Step size (2θ/ °)	0.035	0.030
Scan time per step / s	10.30	12.00

Table 4 – Scan parameters of X-ray diffraction

The samples were mounted (on separate occasions) in a vertical diffractometer, affixed to silicon crystals by a minimal quantity of Blu-Tak™. The silicon crystals were cut with their crystallographic axes oriented so as not to interfere with the data at all. The other parameters of the scans are detailed in Table 4. The scans themselves are shown in Figs. 27 & 28.

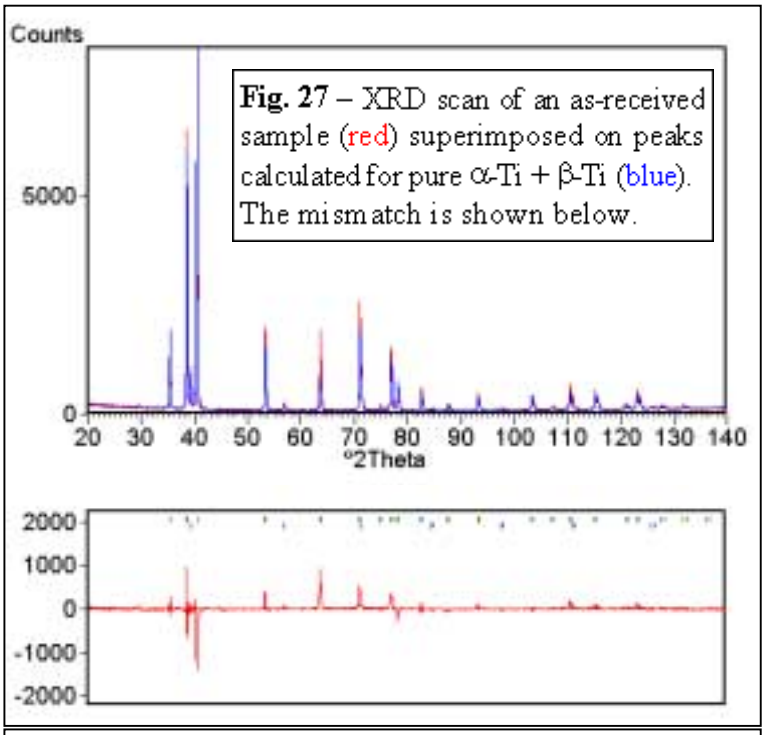

Fig. 27 – XRD scan of an as-received sample (red) superimposed on peaks calculated for pure α-Ti + β-Ti (blue). The mismatch is shown below.

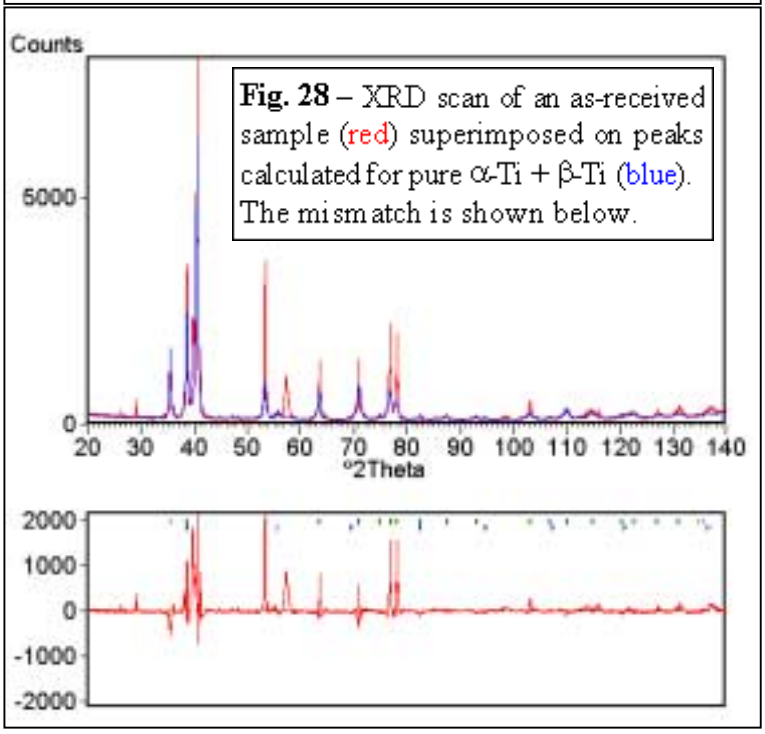

Fig. 28 – XRD scan of an as-received sample (red) superimposed on peaks calculated for pure α-Ti + β-Ti (blue). The mismatch is shown below.

The data was compared with α-Ti[10] & β-Ti[13] structures found in a database. Using Rietveld refinement, the data from the as-received sample showed the β fraction to be 6.8 ± 0.3 wt-% and the α fraction to be 93.2 ± 0.6 wt-%. The calculated diffraction data for this fraction matched the observed data very well (Goodness-of-fit = 6.533). The same technique was used on a homogenised sample, revealing a β fraction of 0.4900 ± 0.0005 wt-% and an α fraction of 99.51 ± 0.06 wt-% (Goodness-of-fit = 27.604). The fraction of β-Ti has clearly decreased significantly during the oxygenation processes.

The optical phase analysis gives volume-fractions rather than weight-fractions. To compare its results with those from the XRD (see Table 5), it is necessary to convert all the XRD results to volume-fractions. The Rietveld refinement gives estimated densities for the observed phases.

The results from the two techniques are in significant disagreement, being outside of the standard error ranges. Since the optical technique is prone to inaccuracy from several sources (such as topographical features on the samples; for a full discussion see section ③), and the postulated composition fits the XRD data well, the XRD technique is probably the more accurate of the two.

Although the optical phase analysis technique is marginally less reliable, it is relatively easy to use, requiring only a clear compositional micrograph and some software. In a typical laboratory, where X-ray diffractometers are generally in high demand, such an alternative may be highly desirable.

	As-received sample		Homogenised sample	
	α-Ti	β-Ti	α-Ti	β-Ti
Density, ρ / kg m^{-3} (from XRD)	4610	4700	4530	4340
Weight-%(from XRD)	93.2 ± 0.6	6.8 ± 0.3	99.51 ± 0.06	0.49 ± 0.001
Volume-%(calculated from ρ)	93.3 ± 0.6	6.7 ± 0.3	99.49 ± 0.06	0.51 ± 0.001
Volume-%(from optical analysis)	91.5 ± 0.02	8.5 ± 0.02	96.66 ± 0.02	3.34 ± 0.02

Table 5 – Comparison of β fractions measured by different techniques

The XRD data contains information about the lattice parameters of the crystals present. In theory, the composition of a two-element system may be calculated from its lattice parameters by assuming linear variation of the unit cell parameters between their values in pure crystals of the elements. Modelling β-Ti as pure Ti-V, the atomic fraction of V in the β phase of an as-received sample may be calculated in this way.

A review of literature[7-21] on β-Ti and V revealed the following unit cell parameters for them:

V: 3.0282 ± 0.0002 Å (at 303 K), 3.0278 ± 0.0001Å, 3.0281 Å, 3.036 ± 0.002 Å, 3.0236 Å, 3.0241 Å (at 298 K), 3.026 ± 0.001 Å and 3.0258 Å.

β-Ti: 3.3111 Å (at 1181 K) and 3.283 ± 0.002 Å (at 298 K).

Taking a(β-Ti) = 3.283 Å and a(V) = 3.024 Å, and the lattice parameter of β-Ti in an as-received sample of Ti-6Al-4V as 3.2349 ± 0.0003 Å, the composition of the β-Ti is estimated as 18.6 at-% V (and 81.4 at-% Ti).

The EDAX tests on a homogenised sample found that the β phase contained 21.59 ± 1.15 wt-% V, which is 20.11 ± 1.12 at-% V, fairly close to the value calculated above. Unfortunately it was not possible to perform a similar EDAX test on an as-received sample, since the electron microscope that had the EDAX facility became in need of a new detector for backscattered electrons, without which it was impossible to select points in either the α or the β phases for energy-dispersive X-ray analysis. It could be assumed that the compositions of α and β in a homogenised sample are similar to those in an as-received sample. However, as discussed in section ⑤, the distribution of V and Al (and Fe) between the α and β phases changes significantly. This is shown by the peaks in the EDAX line-scan of the homogenised sample, which are higher than those in the as-received sample. Therefore it is not reliable to compare the results of XRD from an as-received sample with the EDAX results from a homogenised sample.

Furthermore, there is a significant fraction of Al in the β phase. Even in a homogenised sample, where the Al tends to have diffused into the α phase, this has been measured as 5.95 ± 0.29 at-%. This, of course, demonstrates the weakness of the assumption that the β phase is pure Ti-V. Aluminium, having a much lower atomic radius, tends to decrease the unit cell parameters of the metal on alloying substitutionally. In the above calculation, it would greatly increase the calculated fraction of vanadium, as observed. Iron is also present in the β phase. With its larger atomic radius, it should have the opposite effect. Overall, the unbalanced effects of the Al and Fe make it almost impossible to accurately calculate the composition of each phase directly from the XRD data.

⑦ Finite element analysis of the homogenisation process

The homogenisation of the oxygen profiles in the samples is modelled by a very simple one-dimensional finite element technique. The model assumes that the oxygen moves by diffusion through a sealed body of uniform composition. It is assumed that there is no limit on the solubility of oxygen in the matrix, which is far more accurate in α-Ti than β-Ti, and that the diffusion at any point is entirely governed by the local concentration gradient as described by Fick's 1st Law. Any effects of the oxygen concentration on the stability of the phase are neglected.

Two computer programs have been written in the Fortran language (version '77) to run the model. Although their user interfaces and the non-critical parts of the programs are almost identical, they operate by completely different algorithms.

Each program searches for a file of input data called "Homo.inp". This file consists of a list of concentration levels estimated at regular intervals through the sample. The program counts the entries in this list, stopping as soon as it encounters a negative value. It asks the operator for the parameters of the sample: its length (from which it calculates the distance between the points whose concentration levels are given in "Homo.inp") and its diffusion coefficient for the mobile element (in this case the element is oxygen). The operator must also input the duration of the homogenisation process and the number of time-steps the model should calculate. The accuracy of the model increases if the time-steps are small and if the data points are finely spaced. However, the rounding errors increase with every calculation, so the number of data points and time-steps should not be increased excessively. Before commencing the homogenisation, the programs calculate the mean concentration. This is re-calculated from the final values after the homogenisation has taken place. A comparison of these two mean values provides a way of finding how much matter has been "lost" during the calculations.

The first program, HomoSimp.f, uses a very simple algorithm: running through the sample from left to right, it examines the elements in pairs and transfers a small quantity of matter between them, calculated from Fick's 1st Law. This is repeated for the desired number of time-steps. The algorithm is not perfectly symmetrical; if run from right to left, the result would be slightly different. However, this effect becomes insignificant if each time-step and *x*-step is very small. After every time-step, the concentration profile is output to a file called "Homo.out".

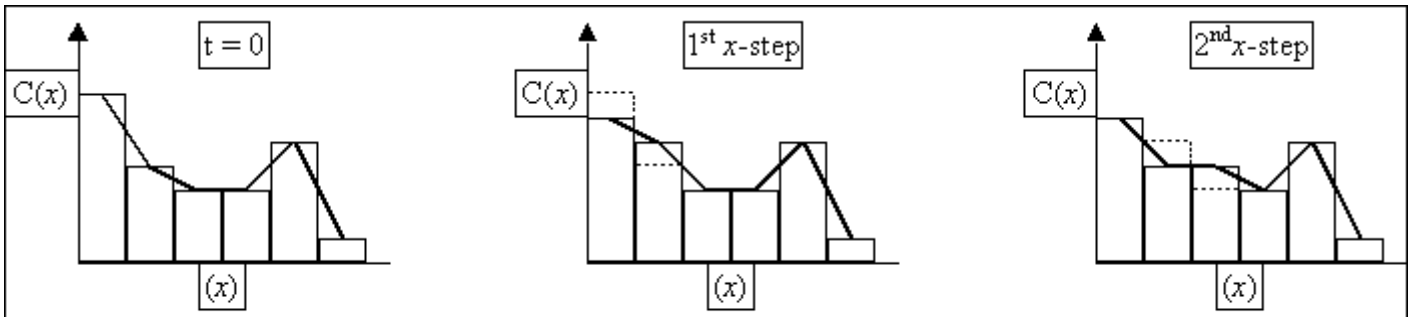

Fig. 29 – Schematic representation of the action of HomoSimp.f on the concentration data

The key equations of HomoSimp.f are: $\frac{dC}{dx} = \frac{\mathrm{Conc}(x+dx,t)-\mathrm{Conc}(x,t)}{dx}$

$$\mathrm{Conc}(x,t+dt) = \mathrm{Conc}(x,t) + D\frac{dC}{dx}dt$$

$$\mathrm{Conc}(x+dx,t+dt) = \mathrm{Conc}(x+dx,t) - D\frac{dC}{dx}dt$$

(although in the next *x*-step, the element currently referenced by $x + dx$ will be altered further)

The second program, HomoPro.f, uses a completely different algorithm, based on a standard method[3] used in finite element analysis: running through the sample from left to right, it examines the elements in groups of three, and adjusts the level of the middle one by some quantity calculated from the concentration gradients to the left and the right. At the edges of the sample, it is assumed that no transfer can occur to or

from the environment, so the end-most elements are treated as if reflections of the penultimate elements existed on the outer side of the end-most elements. A critical difference between this algorithm and that of HomoSimp.f is that HomoSimp.f uses only one array to store the profile. It changes two values in the array with each *x*-step, and these new values are used in the next *x*-step. HomoPro.f, on the contrary, uses two arrays. The values in the second array refer to the next time-step; they are calculated from the first array without any changes being made to the first array. At the end of each time-step, the values in the first array are replaced with those from the second. These are output to Homo.out. This process is repeated for the desired number of time-steps.

The key equations of HomoPro.f are:

$$\frac{\mathrm{d}C}{\mathrm{d}x} = \frac{\mathrm{Conc}(x+\mathrm{d}x,t)+\mathrm{Conc}(x-\mathrm{d}x,t)-2\times\mathrm{Conc}(x,t)}{\mathrm{d}x}$$

$$\mathrm{Conc}(x,t+\mathrm{d}t) = \mathrm{Conc}(x,t) + D\frac{\mathrm{d}C}{\mathrm{d}x}\mathrm{d}t$$

… and at the edges:

$$\frac{\mathrm{d}C}{\mathrm{d}x} = \frac{\mathrm{Conc}(x=2\times\mathrm{d}x,t)-\mathrm{Conc}(x=\mathrm{d}x,t)}{\mathrm{d}x} \text{ and } \frac{\mathrm{d}C}{\mathrm{d}x} = \frac{\mathrm{Conc}(x=x_{\mathrm{Total}}-\mathrm{d}x,t)-\mathrm{Conc}(x=x_{\mathrm{Total}},t)}{\mathrm{d}x}$$

(x_{Total} = sample width)

Since HomoPro.f calculates the entire $\mathrm{Conc}(t+\mathrm{d}t)$ profile from $\mathrm{Conc}(t)$, rather than from a mixture of $\mathrm{Conc}(t)$ and $\mathrm{Conc}(t+\mathrm{d}t)$, it is inherently less dependent on the order of the *x*-steps. This removes one source of error present in HomoSimp.f. However, since HomoPro.f alters each element separately, rather than shifting mass from one to another, it is more prone to loss or gain of mass throughout the homogenisation procedure.

Fig. 30 **Comparison of input and output from the model**

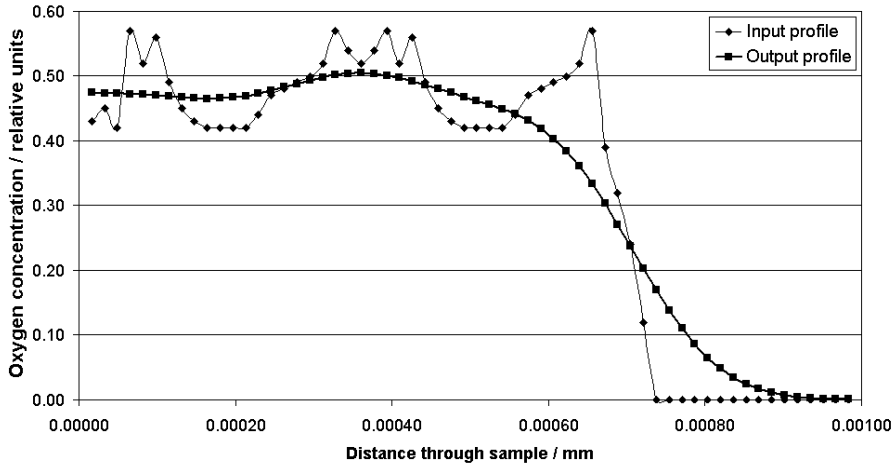

The outputs of the two programs were compared. After processing 60 elements for 50 000 time steps (3×10^6 concentration values), HomoSimp.f calculated a 6×10^{-4} % change in the mean concentration. HomoPro.f processed the same input with exactly the same parameters and calculated a 1.12×10^{-2} % change. These errors are negligible. Indeed, in Fig. 30, the output profile from only one program is shown, since the two are almost identical.

The algorithm employed by HomoSimp.f is extremely crude, so it might be expected that it would produce different results to those from HomoPro.f. However, on comparison, the outputs of the two programs are identical, to within 0.03 %, decreasing for larger concentration values. This suggests that with small time-steps and a large number of finite elements, the cruder algorithm is adequate.

⑧ Electro-deoxygenation of oxygenated Ti-6Al-4V samples in molten calcium chloride

As discussed in section ①, the standard reduction of titanium ore is an expensive process. The Fray-Farthing-Chen process[1], developed in Cambridge, uses an electrochemical method to deoxygenate TiO_2 in a simpler and cheaper process. In this section, the optical phase analysis and micro-hardness techniques described in sections ③ & ④ are applied to deoxygenated samples, and the resulting profiles are compared with a solution of Fick's 2nd Law.

The samples used in the electrolysis had been oxidised at 700 ºC for 115 hours and homogenised at 900 ºC for 287½ hours. It is assumed that they were identical. Other such samples were tested in the Eltra OHN2000 oxygen analyser by the elemental fusion technique described in section ④, giving their oxygen content as 1.5664 ± 0.02 wt-%. One further sample was tested to find the average hardness and β fraction of the samples before electrolysis; these are plotted on the graphs in Figs. 32-37.

The electrolysis was performed in a sealed furnace (see Fig. 31) containing an inert atmosphere of argon. This was essential, since the samples would oxidise rapidly in an oxygen atmosphere at the required temperatures. Since $CaCl_2$ becomes molten at about 800 ºC, and its mobility increases with temperature above that point, the procedure was carried out at 950 ºC.

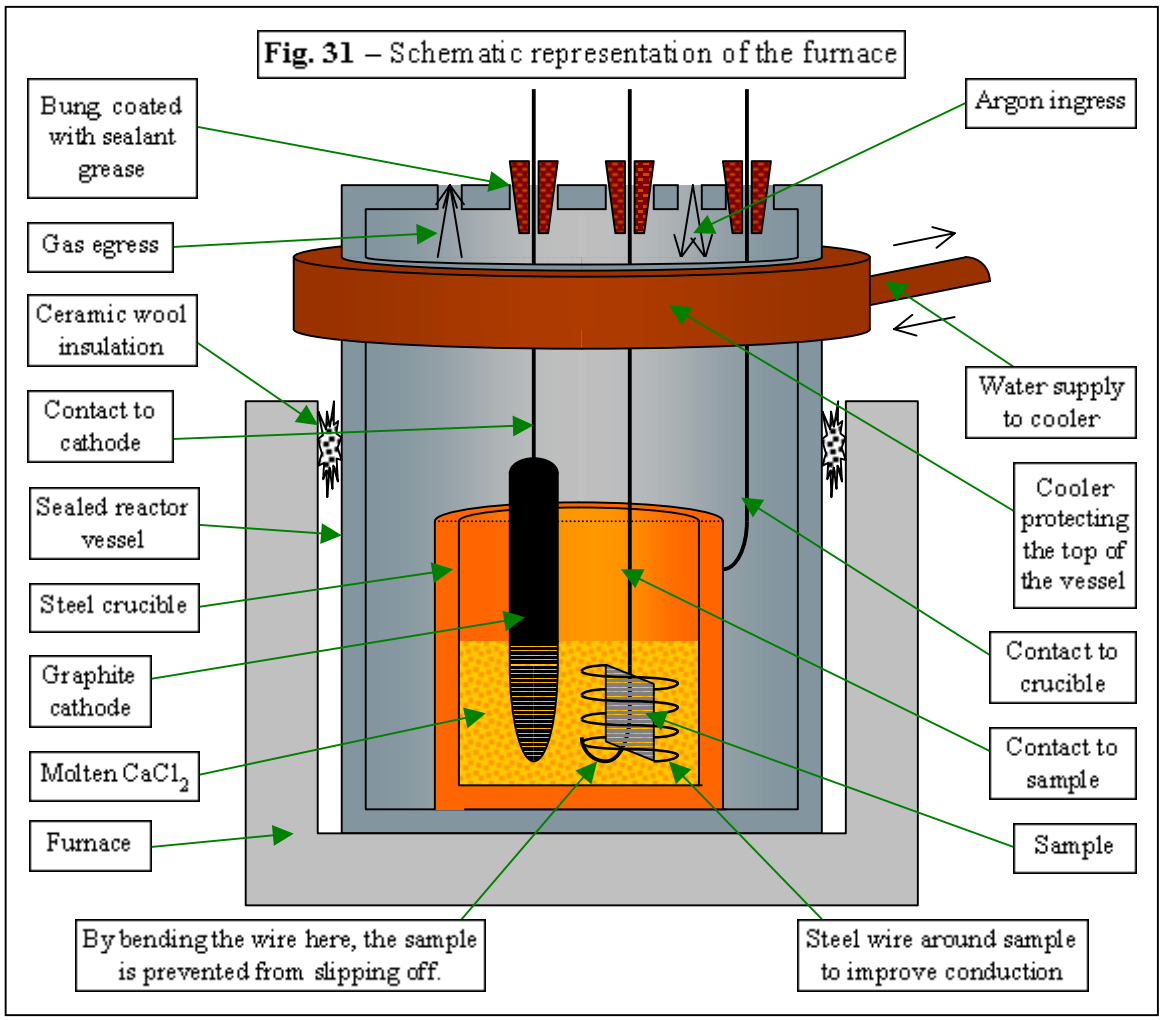

The crucible was filled with dehydrated calcium chloride powder and lowered into the reactor vessel on the end of a thick (2 mm diameter) molybdenum wire, attached firmly by a screw, thus ensuring a good electrical contact. The graphite cathode was lowered into the crucible, also on a thick wire. At this point the lid was placed onto the reactor vessel, sealing it completely. The wires passed through small holes in the lid, which were plugged with rubber bungs coated in sealant grease. As the furnace heated up, the gas tubes were connected to the lid to permit the ingress and egress of argon, thus ensuring an inert atmosphere. The cooler was connected to the water supply, protecting the top of the furnace from the internal temperature; the operator often has to manipulate apparatus adjacent to the furnace lid, and this would be dangerous if the lid were permitted to become hot. Once the salt had melted, a pre-electrolysis current of 2.5 ± 1.0 A was passed between the cathode and the crucible, removing any water present in the salt.

After at least an hour of pre-electrolysis, the voltage and current were noted. The sample was inserted quickly, and its wire connected to the positive terminal of the power source. For the electrolysis, the sample was therefore the anode. Oxygen was taken from the sample as O^{2-} ions, which were attracted to the graphite cathode, where they formed CO_2. This waste gas was allowed to leave the furnace through the appropriate hole in the lid, as shown in Fig. 31.

The waste gases were not released immediately into the atmosphere; they were first passed through ceramic wool and desiccant crystals, then bubbled through two water traps. These have been omitted from Fig. 31.

The voltage and current across the system were noted at the beginning and end of the electrolysis. From these it was hoped to calculate the quantity of oxygen removed from the sample, making the assumption that the flow of oxygen was the only current-carrying activity. To make this assumption more accurate, the surface area of the anode was minimised; rather than making a basket of thin steel wire, the sample was tied closely to the molybdenum main contact with a minimum quantity of steel wire. This reduced the current carried by other mechanisms, such as possible deoxygenation of the steel wire.

Various difficulties were encountered during this part of the work, such as a leak in the crucible that necessitated a day's delay while the reactor was cleaned, with the result that only three samples were successfully electrolysed. The relevant measurements are shown in Table 6.

Sample number	Duration / hrs	Pre-electrolysis current, I_0 / A	Initial current, I_1 / A	Final current, I_2 / A
22	½	4.3 ± 0.2	0.99 ± 0.01	0.73 ± 0.01
23	1½	3.2 ± 0.2	0.64 ± 0.01	0.31 ± 0.01
26	4½	1.24 ± 0.05	1.10 ± 0.01	0.01 ± 0.01

Table 6

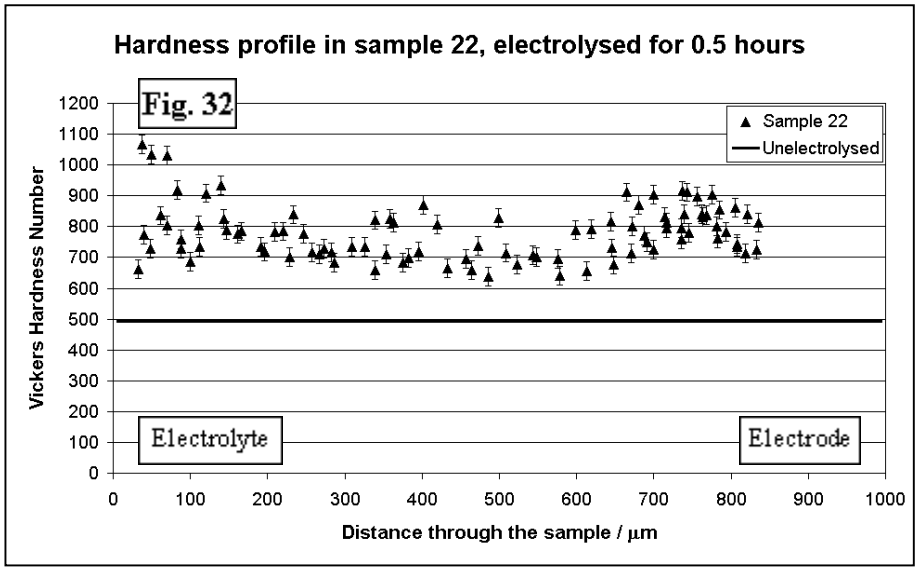

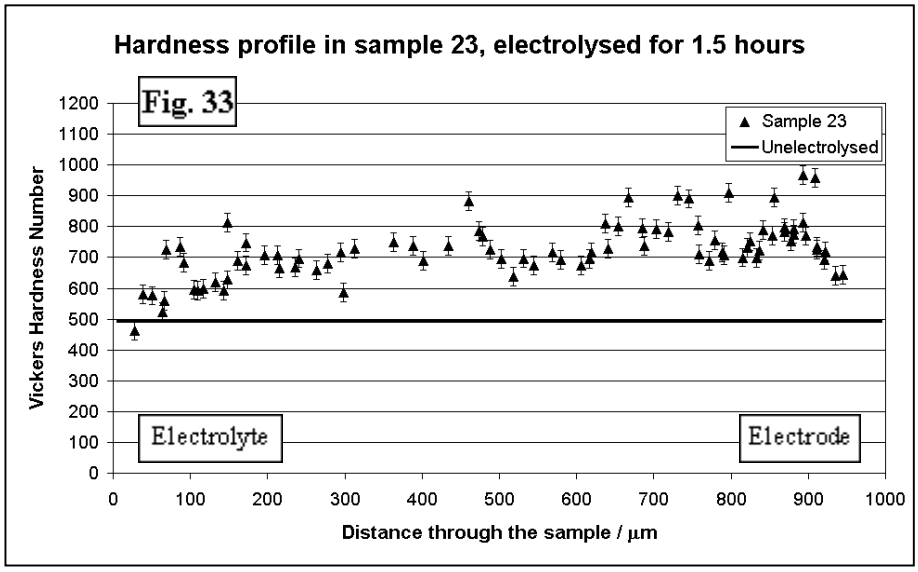

In all cases, the electrolyte was molten $CaCl_2$, the voltage both before and during electrolysis was 2.5 ± 0.05V and the temperature was 950 ºC.

The measured currents are double what might be predicted. This is not a random error; the fluctuation of the readings is reflected in the errors quoted above. It is most likely that some other mechanism is allowing a current to flow between the anode and cathode. Clearly it is impossible to calculate the quantity of oxygen removed from the samples accurately. Measurement of the mass lost by a sample during electrolysis would be similarly inaccurate, since the salt tends to enter any porous region of the sample, affecting its mass.

30 minutes of electrolysis has significantly increased the hardness of sample 22 (Fig. 32). The hardness is very uniform across most of the sample, but the 150 μm closest to the electrical contact could not be measured; it was so brittle that any attempt to indent it caused the development of microscopic cracks,

as described in section ④. It seems very likely that the temperature and chemical effects of the treatment have caused the sudden increase in hardness rather than any effect associated with the oxygen content – indeed, an immediate decrease in oxygen content (and hence hardness) is predicted below.

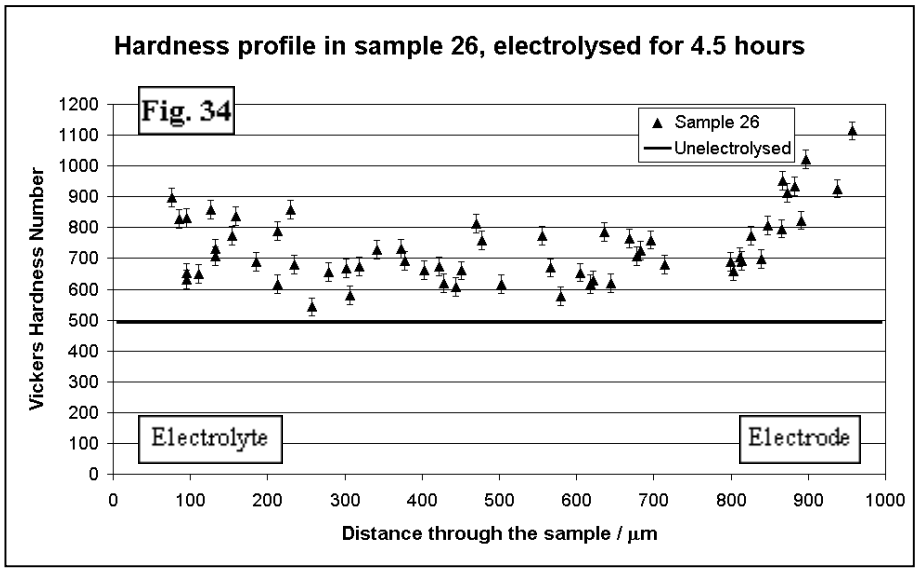

Sample 23, electrolysed for 90 minutes, has a hardness comparable to the majority of sample 22 in the region closest to the electrical contact (Fig. 33). However, the hardness decreases significantly towards the surface exposed to the electrolyte. This presumably indicates a lower oxygen content, allowing freer movement of dislocations, and hence easier plastic deformation.

Sample 26 has the greatest scatter in its hardness (Fig. 34). No clear trend can reliably be observed, although the hardness is lower than that of sample 22. This indicates, as above, that the oxygen content has decreased.

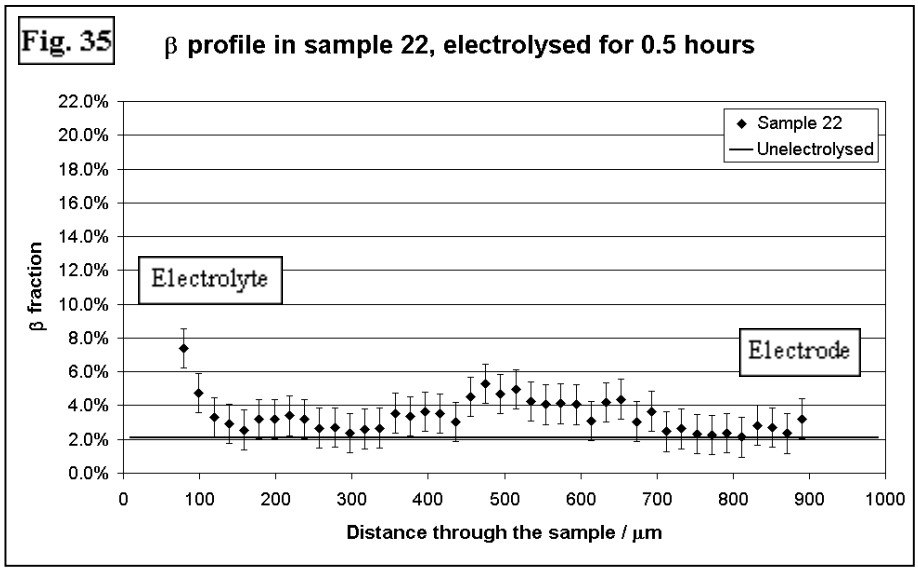

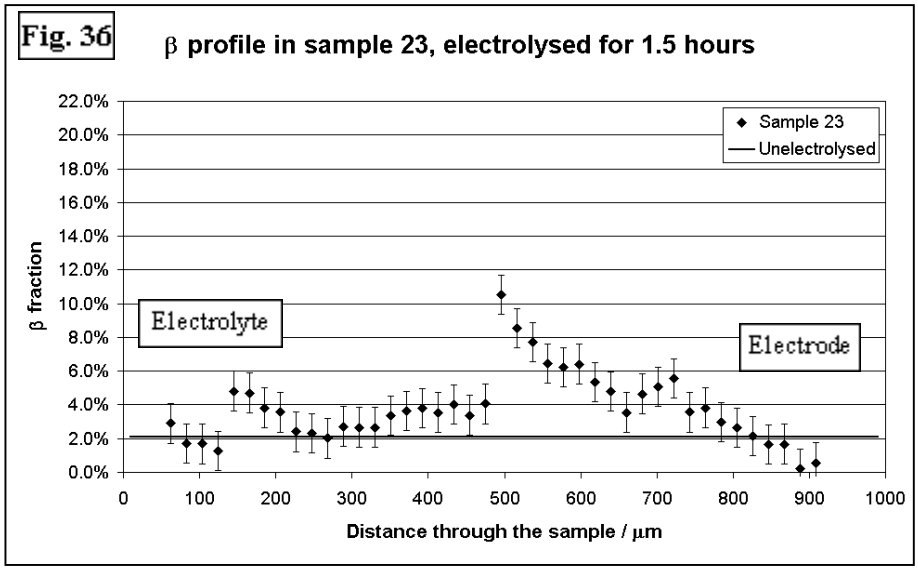

The β profiles show similar effects. Sample 22 has the same β-fraction throughout as an un-electrolysed sample (Fig. 35). The β profile of sample 23 (Fig. 36) is inexplicable; this sample would be re-analysed if time permitted. It does, however, show that there is marginally more β phase in sample 23 than sample 22, indicating a decrease in the oxygen fraction. Finally sample 26, which was electrolysed for 4½ hours, shows a great increase in the fraction of β phase (Fig. 37), especially nearer the surface that was exposed to the salt. The oxygen content has clearly been reduced to a fraction of its pre-electrolysis level. From the micrographs, it may be observed that the microstructure is comparable to that of an as-received sample, indicating that the oxygenation has been reversed almost perfectly.

The decrease of the oxygen content during the electrolysis may be modelled approximately. It is assumed that the motion of the oxygen within the sample occurs by a diffusion mechanism, and that this process limits the electrolysis since diffusion mechanisms are slow in solids – the oxygen anions may be assumed to move much faster in the electrolyte. Fick's 2nd Law, $\frac{\partial C}{\partial t} = D\frac{\partial^2 C}{\partial x^2}$, is solved to predict a concentration profile in the shape of an error function, starting from a uniform concentration.

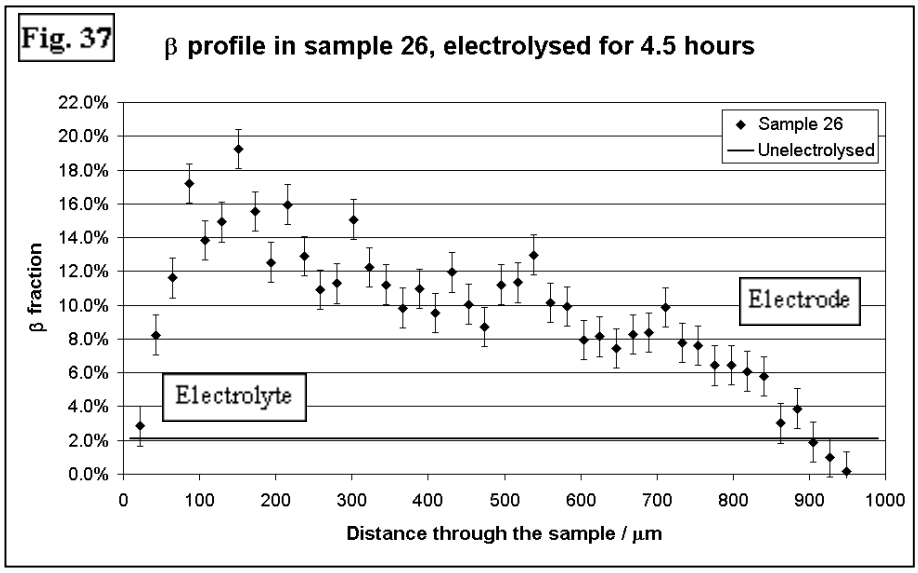

The appropriate solution is:

$$C(x,t) = \left(C^{\alpha} - C^{\alpha E}\right)\mathrm{erf}\left(\frac{x_* - x}{2\sqrt{Dt}}\right) + C^{\alpha E}$$

in which…

C^{α} = concentration of oxygen in the α phase of the sample.

$C^{\alpha E}$ = concentration at the interface of the α phase and the electrolyte.

x = distance from contact wire,

x_* = distance of surface exposed to electrolyte.

For these samples, $x_* = 1$ mm.

Theoretically, this form (Fig. 38) would be observed in all the experimentally obtained profiles. In fact, the hardness profiles all suggest that other factors have influenced the hardness of the samples. Such factors could include increased porosity, which would affect the plastic deformation mechanisms of the metal, increasing the measured hardness. Alternatively, the $CaCl_2$ could have had some chemical effect on the sample which might increase the hardness. The effect is seen to be very fast-acting, but not continuous, since all three electrolysed samples show a similar increase in hardness despite their very different electrolysis durations.

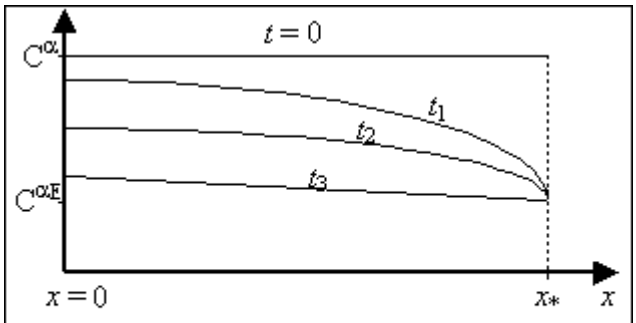

Fig. 38 – The predicted forms of the concentration profile

The β-fraction profiles more closely resemble the model; sample 22 (Fig. 35) has not apparently been greatly affected by the short electrolysis while sample 26 (Fig. 37) shows a far lower oxygen content comparable to an as-received sample. It has clearly been fully deoxygenated across most of its thickness. The model predicts a very low oxygen concentration at the sample / electrolyte interface in every electrolysed sample. This is not observed by either the hardness measurements or the optical analysis.

⑨ Conclusion

The work describe in this report successfully analyses the characteristics and properties of the Ti-6Al-4V samples in various stages of oxygenation and deoxygenation. The relationships between the oxygen content of the alloy and its phase fractions and hardness have been described. Although the comparison of the empirical results and theoretical predictions was not always conclusive, qualitative trends are observed to be in accordance with the theory, except for the deoxygenated samples. Further work, refining the relationships between the oxygen content and the hardness and the phase fractions, would be desirable. This relationship could be applied to electro-deoxygenated samples, revealing whether or not the oxygenation is perfectly reversed. If it is found that the deoxygenation induces properties different from those observed in an un-oxygenated sample, these properties could be applicable.

⑩ Acknowledgements

This work was supported by funding from Dr G. Z. Chen[*], who also provided the Ti-6Al-4V samples, and by the loan of resources from Prof. H. K. D. H. Bhadeshia[†].

Thanks are due to T. Mach, who assisted with some of the practical work, in particular the hardness testing of the oxidised and as-received samples, and who developed the method used in the optical phase analysis. The author would also like to thank F. Clarke, D. Vowles, M. Vickers, T. Mosdall and K. Roberts for invaluable technical assistance.

● References

1. G. Z. Chen, D. J. Fray & T. W. Farthing: *Metallurgical and Materials Transactions B*, 2001, vol. 32B, pp. 1041-1052.
2. G. Z. Chen, D. J. Fray & T. W. Farthing: *Nature*, 2000, vol. 407, pp. 361-364.
3. H. K. D. H. Bhadeshia: *Materials Science & Technology*, 1989, vol. 5 (2), pp. 131-137.
4. R. I. Jaffee: *Progress in Metal Physics*, *The Physical Metallurgy of Titanium Alloys*.
5. P. Kofstad: *Journal of the Less-Common Metals*, 1967, vol. 12, pp. 449-464.
6. *Cost Effective Titanium Component Technology for Leading-Edge Performance - Seminar Papers*, 1999.
7. H. E. Swanson, R. K. Fuyat & G. M. Ugrinic: *National Bureau of Standards (U.S.), Circular*, 1954, vol. 359, pp. 1-73.
8. E. S. Makarov & L. M. Kuznetsov: *Zhurnal Strukturnoi Khimii*, 1960, vol. 1, pp. 170-177.
9. R. J. Wasilewski: *Transactions of the Metallurgical Society of Aime*, 1961, vol. 221, pp. 1231-1235.
10. R. M. Wood: *Proceedings of the Physical Society, London*, 1962, vol. 80, pp. 783-786.
11. W. Rostoker: *Journal of Metals*, 1952, vol. 4, pp. 981-982.
12. B. W. Levinger: *Journal of Metals*, 1953, vol. 5, p. 195.
13. J. Spreadborough & J. W. Christian: *Proceedings of the Physical Society, London*, 1959, vol. 74, pp. 609-615.
14. O. N. Carlson & C. V. Owen: *Journal of the Electrochemical Society*, 1961, vol. 108, pp. 88-93.
15. W. J. James, M. E. Straumanis: *Journal of the Electrochemical Society*, 1960, vol. 107, p. 69.
16. M. E. Straumanis: *Journal of Applied Physics*, 1959, vol. 30, pp. 1965-1969.
17. E. Sandor & W. A. Wooster: *Nature*, 1958, vol. 182, pp. 1435-1436.
18. A. Wexler & W. S. Corak: *Physical Review*, 1952, vol. 85, pp. 85-90.
19. A. U. Seybolt & H. T. Sumsion: *Journal of Metals*, 1953, vol. 5, pp. 292-299.
20. S. Beatty: *Journal of Metals*, 1952, vol. 4, pp. 987-988.
21. S. A. Bradford & O. N. Carlson: *Transactions of the American Society for Metals*, 1962, vol. 15, pp. 421-422.
22. R. G. Ward & T. P. Hoar: *Journal of the Institute of Metallurgy*, 1961-62, vol. 90, pp. 6-12.

[*] School of Chemical, Environmental & Mining Engineering, University of Nottingham
[†] Department of Materials Science & Metallurgy, University of Cambridge