

Why OPERA's claim for faster-than-light neutrinos is *not* wrong

John P. Costella, Ph.D.

Melbourne, Australia

(25 September 2011)

Abstract

Two days ago I posted an analysis arguing that the claim of the OPERA collaboration for having measured faster-than-light neutrinos was based on an incorrect statistical calculation. I have since received an extremely elegant yet simple analysis by David Palmer that demonstrates that I was wrong. In this document I describe in detail the Palmer analysis, which both explicitly refutes one of the assumptions I erroneously made in my analysis, and provides an extremely simple and intuitive way of calculating the statistical significance of the OPERA neutrino experiment, putting it beyond any doubt.

Introduction

It has been an interesting 45th birthday.

On Friday—when I was still 44—the world was told that muon-neutrinos travel faster than the speed of light. Contrary to some media reports, this doesn't mean that Einstein's theory of relativity is wrong: the neutrinos weren't accelerated through the "light barrier"; but it does mean that, in some frames of reference, they were going backwards in time. Again, contrary to some media reports, that doesn't mean that *we* could go backwards in time, à la Michael J. Fox in *Back to the Future*. But it does mean that we could send *signals* backwards in time, which causes almost as much trouble to causality.

Today, I am sitting amongst my birthday presents, wishing that one of them was an iPhone 27 with muon-neutrino backwards-in-time messaging enabled. (Alas, again, relativity only allows the neutrinos to go backwards in time, not the phone.) On Friday, the results put out by the OPERA group in a preprint just hours before looked wrong: they were claiming a statistical significance of just seven nanoseconds, whereas their graphs of the "end points" didn't look anywhere near that accurate. I posted an analysis saying as much, including a back-of-the-envelope calculation appropriate for distributions with such smudgy end points. Hours later, several other particle physicists at the seminar / press conference held at CERN (at midnight my time: my first birthday present) made exactly these same points during the question and answer session at the end: the results just didn't look right. (One of them even did the same back-of-the-envelope calculation in his head as part of his question.) The answers from OPERA spokesman Dario Autiero were not satisfactory: he did not seem to grasp the need for a simple, order-of-magnitude confirmation of the uncertainty, but essentially gave answers equivalent to "we put the data into our Maximum Likelihood estimator, pushed the button, and this is what came out". Critics of my posted analysis did the same, with vague arguments that Maximum Likelihood calculations are not counting experiments (which they are).

And so I enjoyed my birthday without any revelations that would explain the discrepancy.

Today, I received an email from David Palmer that completely dissolved the fog. His analysis is elegant and simple. It pinpoints the erroneous assumption that I made. It discards the confusing implication of the OPERA paper that the statistical significance derives from the 16,111 neutrinos that they measured; in fact, almost all of them are irrelevant. And it confirms the OPERA result.

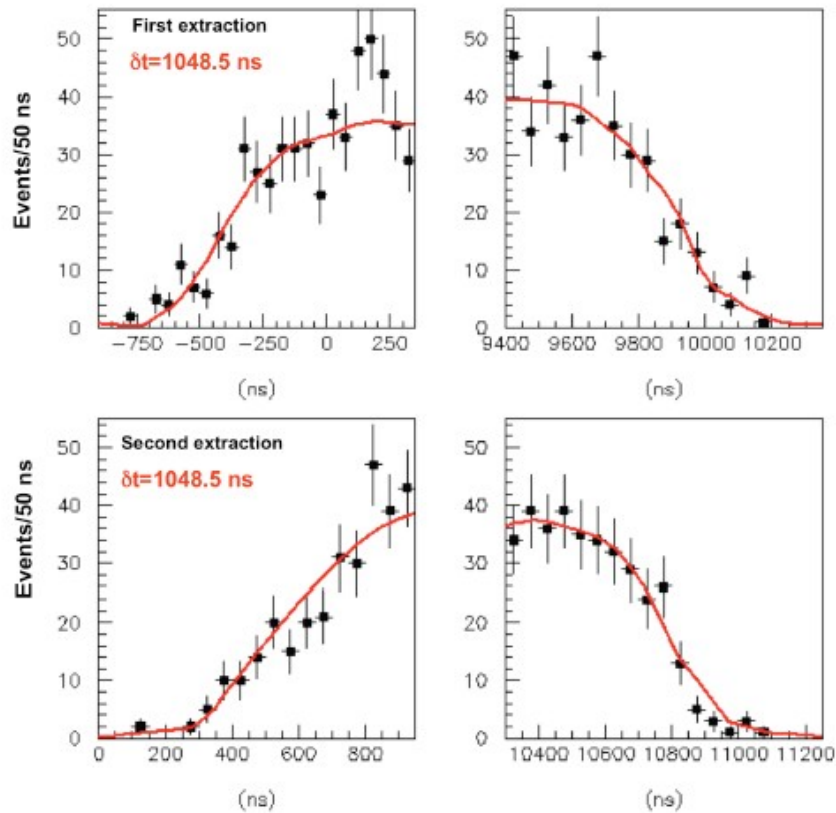


Fig. 12: Zoom of the leading (left plots) and trailing edges (right plots) of the measured neutrino interaction time distributions (data points) and the proton PDF (red line) for the two SPS extractions after correcting for δt (blind).

The key graphs

In my original posting I included Figure 12 of the OPERA paper, and I include it again, above. These graphs are the crux of the whole experiment, and in fact you can read the number of neutrinos off the graphs themselves (in each 50-nanosecond bin, at least). The reason that the parts of the graphs between these end regions (the parts of Figure 11 that are omitted in Figure 12) are almost irrelevant is that the distributions in those regions are (to a good approximation) relatively flat, and so those regions provide very little information about the shift in the travel time distribution.

Figuring out this shift simply requires figuring out how far the red curves have to be shifted left or right to match the black data points.

My observation

In my original posting, I noted that OPERA claimed that the statistical error in this shift is only seven nanoseconds, which is about one-seventh of the smallest horizontal tick-mark bins shown on each of the graphs of Figure 12. This is a tiny amount compared to the width of each “edge”. In my opinion—and, subsequently, those of other particle physicists at the seminar—it was just too small to be believable. And even now, it still just “looks wrong”. (Maybe my eyes are getting too old.)

Of course, an eyeball estimate is just that; it’s usually a good guide to where problems are, but you need to follow up with actual analysis. The next step should always be to do a rough, back-of-the-envelope, order-of-magnitude calculation to see whether the visual impression is borne out.

My assumption

My assumption, in looking at Figure 12, and re-reading the OPERA paper, was that the edges of the source distributions were simply not sharp enough—and, consequently, that there were not enough data points there—to determine the shift from the edges alone. The OPERA paper itself reinforced this belief: they state that the statistical significance derives from the 16,111 neutrinos that were measured overall, not the much smaller number of neutrinos shown in Figure 12. I therefore assumed that the best estimate of the shift would be obtained by using the whole distribution of 16,111 travel times.

My calculation

If my assumption was correct, then the simplest way to estimate the shift is to simply compute the mean of the distribution. From the Central Limit Theorem we know that the standard error of the mean will be simply the standard deviation of the source distribution divided by $\sqrt{16,111}$. Since the source distribution is (when viewed as a whole) very well approximated by a uniform distribution of width approximately 10,500 nanoseconds, the standard error of the mean is $10,500/\sqrt{12 \times 16,111} \approx 24$ nanoseconds.

The Palmer refutation

David Palmer's email to me was clear and simple: to his eyeballs, the shift looked good enough to claim a significance of seven nanoseconds. And he constructed an extremely simple and elegant back-of-the-envelope calculation to test it.

And he was right. I was wrong.

The Palmer analysis is so clear and elegant that I will now describe it in detail, and show how it provides an estimate of the statistical significance of the OPERA experiment that is so much more intuitive than the OPERA Maximum Likelihood “black box” approach that I suspect it will become the standard reference for anyone wishing to verify this part of the experiment. It also puts beyond any doubt the fact that *only the time ranges shown in Figure 12 are of critical importance to the experiment*: what happens in the 10 microseconds between the two ends is (to a good degree of approximation) of very little importance. This has important ramifications for the estimation of *systematic* errors; for example, for the proton distribution, only the parts of the red curves shown in Figure 12 are of critical importance, and it is the systematic uncertainties of the proton PDF in those regions (and, to a good approximation, those regions alone) that must be validated.

The Palmer analysis

David Palmer's analysis makes use of the observation that the edges shown in Figure 12 are, to a first approximation, close enough to Normal in shape for a simple estimate of their uncertainty. If we *remove* the roughly 10 microseconds of data between each left and right edge, and “smash” the two end regions together (with the middle deleted section simply thrown away), the resulting distribution will be close enough to Normal to make a good estimate of the statistical uncertainty.

Firstly, for the “plateau” level, we know that there were 16,111 neutrinos, and measuring the widths of the red curves at half-height (which is a good enough approximation because the ends are roughly symmetrical) from Figure 12, the width of the first extraction is about 10,300 ns, and that of the second extraction about 10,230 ns, so the “plateau” rate of neutrinos is around 0.785 neutrinos per ns, or around 39.2 neutrinos per 50 ns bin. (This is actually an iterative calculation: you need to know


the answer to estimate the half-height.) Now, my only contribution to the Palmer analysis is to suggest a repeatable way of estimating where to “cut” each of the four graphs (vertically, as if with scissors) before “pasting” them back together. The magnitude of the maximum slope of a Normal distribution (i.e. the magnitude of the slope at the points of inflection) is given by

$$\frac{1}{\sigma\sqrt{2\pi e}} \approx \frac{0.242}{\sigma}$$

This allows us to determine an approximate value of σ by considering the magnitudes of the maximum slopes of the proton PDFs in Figure 12, the reciprocals of which (normalised to the plateau count) are around 510 and 530 ns for the left edges, and 330 and 360 ns for the right edges, giving us an approximation for σ of 310, 320, 200 and 220 ns respectively.

We now need to determine where to “cut” the left and right distributions before “stitching” them together. For this, we can use the fact that the half-height of a Normal distribution occurs at a distance $\sqrt{2\ln 2} \approx 1.18$ standard deviations away from the mean. The half-height times are at approximately -380 and 530 ns for the left edges, and at approximately $9,920$ and $10,780$ ns for the right edges, which means we should “cut” the left edge distributions at approximately -20 and 910 ns, and the right edge distributions at approximately $9,680$ and $10,520$ ns.

Rounding these numbers off to the nearest bin, we obtain the following “cut and paste” distributions:



To good approximation, almost all of the information in the OPERA experiment is in these data points. If my eyeballs are working properly again, there are 919 neutrinos represented here. In other words, of the 16,111 neutrinos measured by OPERA, the other 15,192 (in the “plateau”) are almost irrelevant.

The Palmer analysis now observes that the measurement of the edges really just reduces down (in this approximation) to measuring the means of these approximately-Normal distributions. The average of the four “standard deviation” approximations measured above is about 260 ns. In total we have 919 neutrinos, so the standard error of the estimate of the mean will be approximately $260/\sqrt{919} \approx 8.7$ ns.

This is only a rough back-of-the-envelope calculation, but it contains most of the information in the OPERA data set. It provides an intuitive estimate of the statistical uncertainty that is completely compatible with the 6.9 ns quoted by OPERA.

Intuitively, it makes good sense that including the information in the “Bart hair” of each extraction’s distribution in an overall Maximum Likelihood analysis will improve the statistical fit slightly. The 919 neutrinos above in the “end regions” determine most of the statistical power of the experiment; the other 15,192 neutrinos only reduce the statistical error by a further 1.8 ns.

Apart from refuting my previous erroneous analysis, this analysis should hopefully assist those analysing the estimated *systematic* uncertainties in the OPERA experiment. Roughly speaking, what occurs in the first and last microsecond of each burst is overwhelmingly of greatest importance; what happens in the ten microseconds in between is less critical.

Conclusions

The Palmer analysis puts beyond all doubt the statistical significance of the OPERA result. It is extremely elegant, and yields a statistically sound approximation that is simple enough that even high school students can check the calculations. It also allows a rough quantitative estimate of the relative importance of the “end regions” and the “plateau region” in analysing the systematic uncertainties in the OPERA experiment.

The “blunder”, the “embarrassing gaffe”, is mine and mine alone. I am happy to wear that ignominy: the OPERA result—if its estimates for systematic errors withstand scrutiny, and if it is subsequently confirmed in future experiments—would arguably be the most important discovery in physics in almost a century. Looking stupid is a small price to pay if it brings us closer to determining whether such a monumental outcome is real or imaginary.

Follow-ups:

[Do OPERA's tachyonic neutrinos make sense?](#)
[Could the OPERA tachyon be the unbroken Higgs?](#)