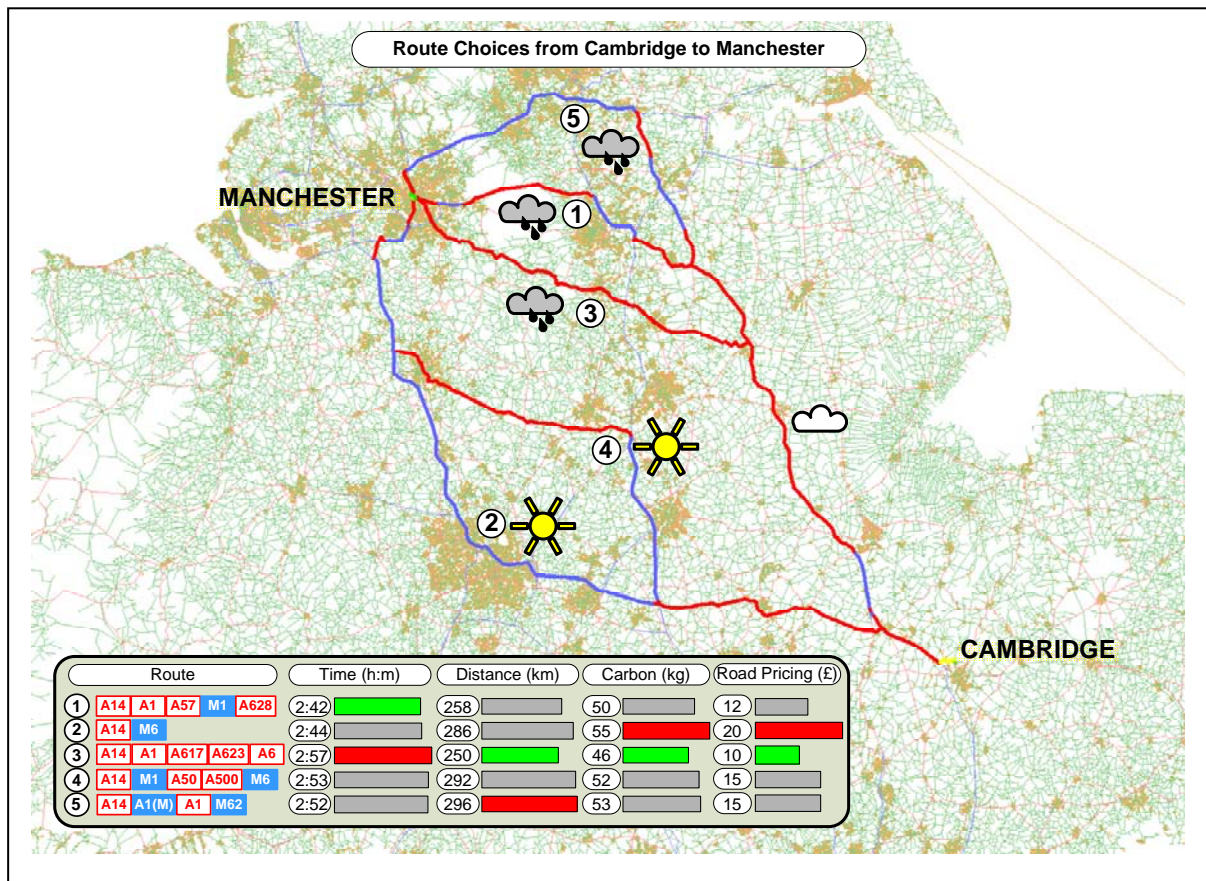


## Choice Routing

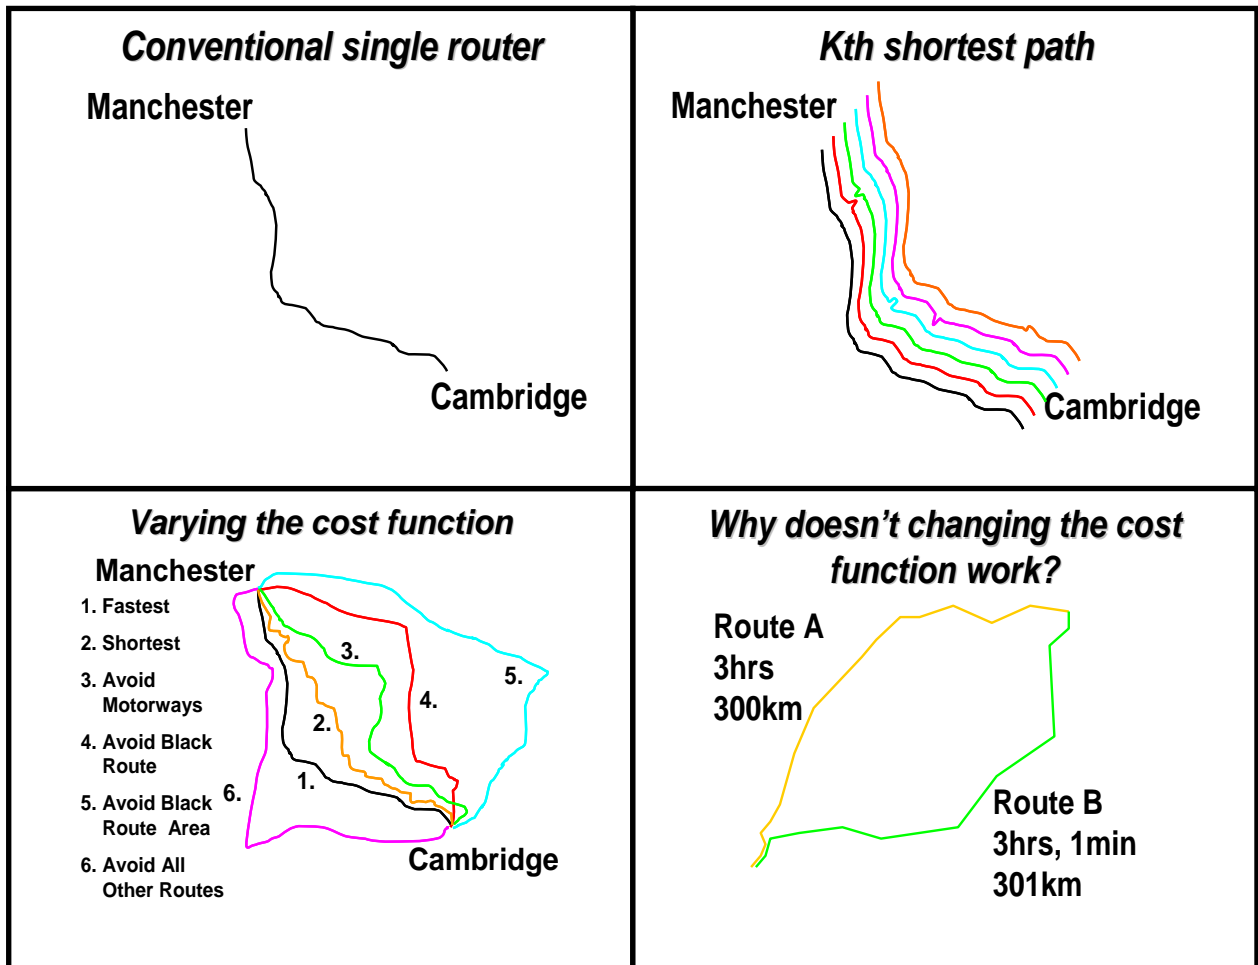
Camvit has discovered a new mathematical approach to route planning that delivers the complete set of good diverse routes. The algorithm is efficient, robust, and can handle the effects of time-dependent traffic data. A further advantage is that it layers on top of an existing “single router” and can take advantage of any optimisations built into the single router.

All previous approaches to generating sets of alternative routes have significant shortcomings and generate a poor choice of routes, this is explained below.



Here is an example of the Choice Routes between Cambridge and Manchester UK, generated in a single pass using our new algorithm. The routes are the ones that an expert driver would recommend (NB we do not have the M6 toll road in our UK map data). Notice the sensible and realistic properties of the routes generated – some routes share long stretches with others.

The set of choices can be presented to the user along with additional appropriate comparative information such as times, distances, fuel usage and, in the future, road pricing. The map of the routes is a good showcase for valuable content such as expected weather along each route, traffic incidents and POI's. The computer has done all the hard work of gathering and presenting the data on the routes and the user can now make an informed choice appropriate for the current circumstances.



## Previous Approaches

A conventional “single router” uses a variant of the Dijkstra algorithm to find the single optimum route between two points under a given cost function. How could alternative routes be generated?

The Kth Shortest Path algorithm finds the single most optimal route in the usual way. It then finds the second most optimal route by excluding each link in the most optimal route, and running the router again, perhaps 1000 times. The most optimal of these 1000 trials is the second shortest route, but this will usually be a minor variation caused by turning left then right rather than right then left at some local intersections, perhaps taking 0.1s longer. The third most optimal route requires every combination of two links to be excluded from the routes, and might require one million runs of the single router. This will find another minor variation, again probably only about 0.1s slower. As we move to further routes, the computation becomes infeasible, and in any case, even the first 1000 shortest routes will just be permutations of minor variations all within a few seconds of each other, so we would not have found any interesting alternatives that might occur at perhaps 30s longer.

Existing commercial approaches for generating alternative routes vary the cost function while a conventional “single router” is repeatedly rerun. By varying the weighting against time and distance in the cost function they generate “fastest” and



Cambridge Vehicle Information Technology Ltd.

“shortest” routes (see “Varying the Cost Function” above). By weighting against motorways, width, prettiness etc., further routes can be generated, but many of them will be similar, and some of the good ones are never generated in this way. Even more perverse routes can be found by weighting against previously generated routes, see above illustration of “Avoid Black Route”, “Avoid Black Route Area” and “Avoid All Other Routes”. These techniques do generate alternative routes, but they are not the ones an expert navigator would choose, and have many peculiarities. For example, all sensible routes from Cambridge to Manchester start off by using a section of the A14 to leave the Cambridge area. Weighting against this stretch of road would just generate unnecessary and uneconomic diversions.

Why doesn't changing the simple cost function work? The final frame above illustrates an important and subtle insight that Camvit has gained during the development of Choice Routing: rerunning the router with a different tradeoff of time and distance will never find an excellent alternative route which is very slightly longer both in time and distance. Clearly the user would be very interested in Route B in the illustration above, but for a single router it will always be masked by Route A. This is a very common occurrence. This does not just happen for complete routes, but for any sub-parts of routes, and for thousands of them at a time. The masking becomes worse if more properties such as motorway use are added to the cost function.


### **How Does Camvit Choice Routing Work?**

Conventional “single routers” begin at the source point, and spread a tree to find the earliest arrival times at any nodes encountered. Once the destination point has been reached or passed, the optimum route is found by tracing backwards along the tree from the destination to the source. After this, the tree is usually discarded. However, it compactly encodes the optimal routes to thousands of points, and there is a good clue here as to where the good alternatives might lie. We call it a Source Tree, and keep it.

If the user has specified an arrival time at the destination point, then a conventional router will be run backwards in time, beginning at the destination and spreading a tree to find the latest departure time at any nodes encountered, until the source point is reached and passed. The optimal route can then be traced backwards in the tree, from source to destination. We call this the Destination Tree.

The Camvit Choice Routing algorithm builds on top of a Company's own router. It uses this router to generate both a Source Tree and a Destination Tree for the same endpoint, and then combines the two trees in a novel way to reveal the set of best route choices.

Each tree has been attracted to chains of road segments that are locally good for getting away from the source, or towards the destination. They are thus attracted to fast radially aligned roads, and repelled from areas of slow or badly aligned roads.

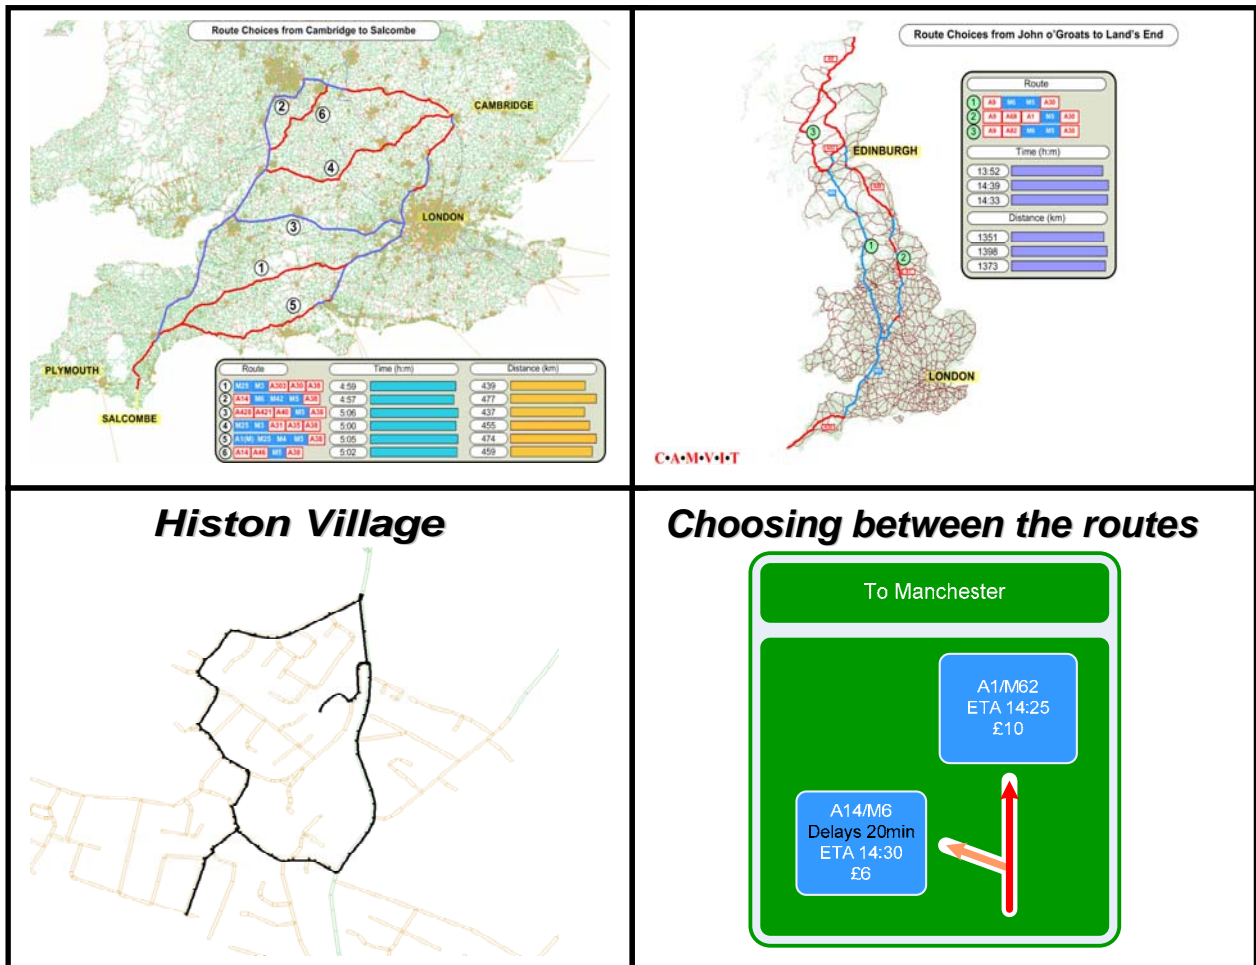


We now scan through the two trees, marking chains of road segments that are traversed in the same direction in both trees. They can easily be found by following the back pointers in the two trees, or alternatively, by looking at the sum of the costs that is stored at each node in each tree. This “via cost” is constant along each chain, but declines towards the source or destination, so we have coined the term “plateaux” for them.

What emerges is a small number of long chains, and millions of small fragments. These long plateaux are the first clue to the good alternative routes, and we select the best using a “goodness” criteria.

Each selected plateau is extended into a complete route by tracing back in the trees from each end. One of these plateaux (actually the one with the highest goodness), is guaranteed to be the single most optimal route that a single router would have found. The goodness is used to decide how many of the routes may be interesting to the user, which usually varies between 1 and 7, although we limit to 5 by default to prevent confusion.

The routes found are the very ones that an expert navigator would propose, no perverse alternatives emerge. The routes are often very close in both time and distance, and where, for geographical reasons, there is only one good route, only one will be found.



The algorithm works excellently at all length scales. In the Cambridge to Salcombe example above the six different choices are very close in time. Between the extreme points in Britain only three significant choices are found, while on a scale of about 1 km. in Histon village near Cambridge only two alternatives are generated because of private roads.

## Handling Historical and Real-Time Traffic Data

None of these examples take into account the effects of traffic congestion, but Choice Routing can handle time-dependent costs for each road link. To give the best information to the driver we would recommend combining the routes found with and without traffic data so that it becomes very clear why a normally good route should be avoided at a particular time. Traffic incidents on any of the Choice Routes can be indicated. This allows the driver to either use the route they normally take after making allowance for the congestion, to make a diversion around the trouble, or to use one of the other Choice Routes that may be faster on that occasion.

## Choice Points en route

When driving along a Choice Route a small number (usually only two or three) Choice Points are encountered on the way, where Choice Routes diverge. Since



Cambridge Vehicle Information Technology Ltd.

these are very special points they could be flagged on an in-vehicle display well before they occur. Information to help the choice could be given, for example, in the form of a road sign which could show information which characterizes each route (see road sign in “Choosing between the routes” above). In this example the left turn onto the A14/M6 route is normally better, but due to traffic conditions today the A1/M62 route may be an option. A conventional “single router” would flip between these two routes without explanation and would have no idea that we are approaching a significant Choice Point. With Choice Routing we can see the reason why the normally good route has a problem. We may choose to go that way anyway since a 20 minute delay is not too bad, or we could choose to go on the faster (and more expensive) route, which might, though, be unfamiliar to us.

### **Choosing By Driving**

On the approach to one of these Choice Points, we can inform the driver about the most up-to-date conditions on each route, according to their preferences. We make this information available perhaps five or ten minutes before the junction, to allow plenty of time for a decision to be reached.

Once the driver has decided which route to choose, they do not have to tell the system. We will observe their behaviour at the Choice Point, and switch guidance to whichever path they choose. In case they stray from the path, we actually switch guidance to the first point on the plateau for the selected route, so that they will be guided there rather than onto some other route. Once that point is reached, guidance is switched to the destination, as that will take us along the plateau anyway.

### **Camvit's Business Model**

Camvit would like to transfer this new technology into your company by providing consultancy in the form of workshops and expert advice. Our experience is that a typical router could be upgraded to one that performs Choice Routing in less than a couple of weeks. We recommend that your engineers implement the new code, as it will then conform to your standards, and is extensible and maintainable independently of us. We would of course be available to support you whenever it is useful, either with changes to the code, or for demonstrations and explanations to your potential customers.

We would bundle a licensing agreement for the technology to be evaluated internally and demonstrated to your customers for 3 years. A license to use the technique commercially can be negotiated at a later stage based on a sharing of the extra revenues that are generated by the use of the technique.



Cambridge Vehicle Information Technology Ltd.

**Company contact details:**

Name            Alan Jones  
Address        Camvit, 67 Narrow Lane, Histon, Cambridge, CB24 9YP, UK  
Tel:             +44 1223 710357  
Email:         ahjones@camvit.com  
Web:            http://www.camvit.com

**Company profile:**

Cambridge Vehicle Information Technology is a private limited company formed by several research engineers from AT&T Laboratories Cambridge Limited (previously Olivetti Research Limited). Our past projects encompassed most areas of that Laboratory's work since its formation 18 years ago including networking, multimedia and tracking systems.

From this unique standpoint we have developed vehicle navigation technologies which will lead to a new generation of commercial systems. We are keen to pursue some of these ideas in partnership with established companies. Camvit has a long-standing close relationship with Cambridge University Computer Laboratory where both Directors are Visiting Fellows. Recent joint research work includes a Wireless Broadband WAN and consultancy for research into Traffic Flows.

Cambridge Vehicle Information Technology Ltd. was formed from [Cotares Ltd.](#), a Cambridge-based Consulting and Research Company.

---

A summary of our ideas is found on our website <http://www.camvit.com>, where the pdf brochure can also be downloaded in English, German, French and Japanese.



Camvit acknowledges the support of the East of England Development Agency and UK T&I.

---