SMART ROUTING

As a result of improvements in probe technology and smart modeling techniques, truly intelligent and real-time traffic information has arrived. Inrix's **Dr Christopher Schofield** charts the progress so far

Illustration courtesy of Nishan Sothilingam

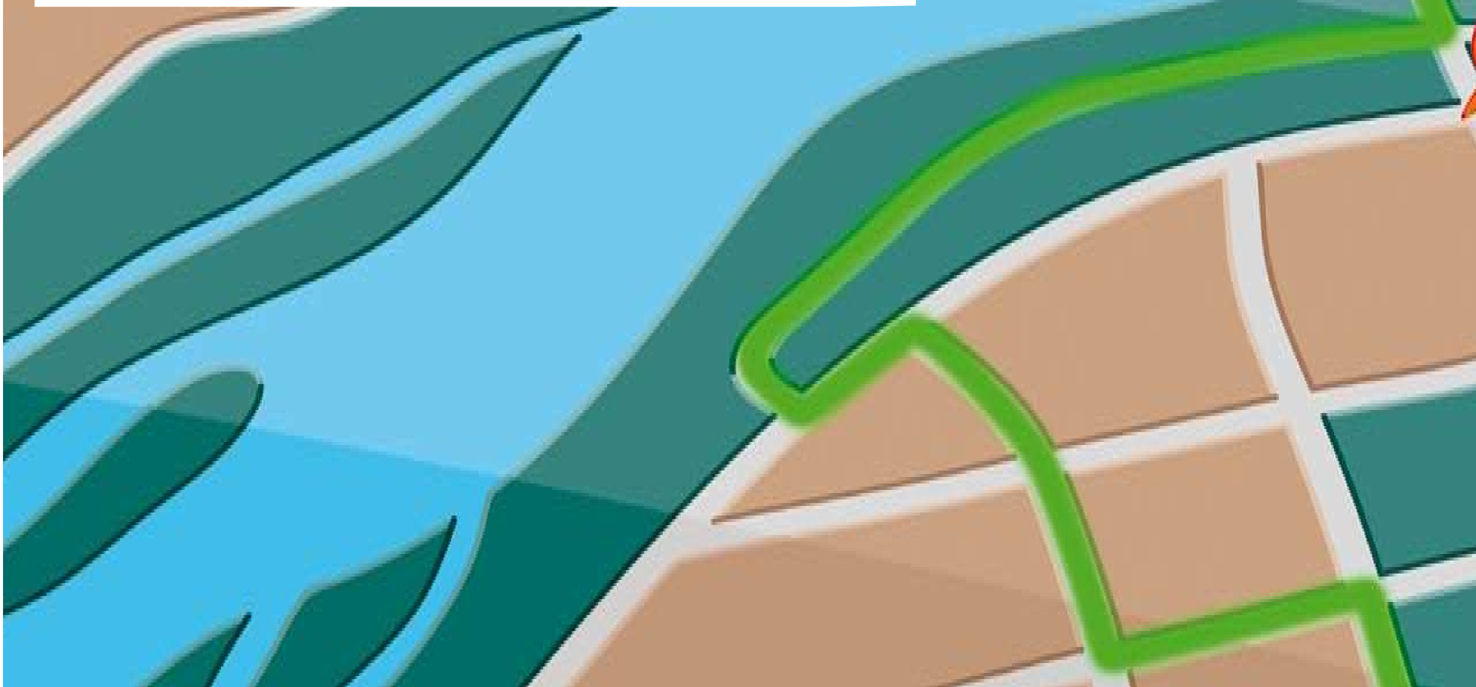

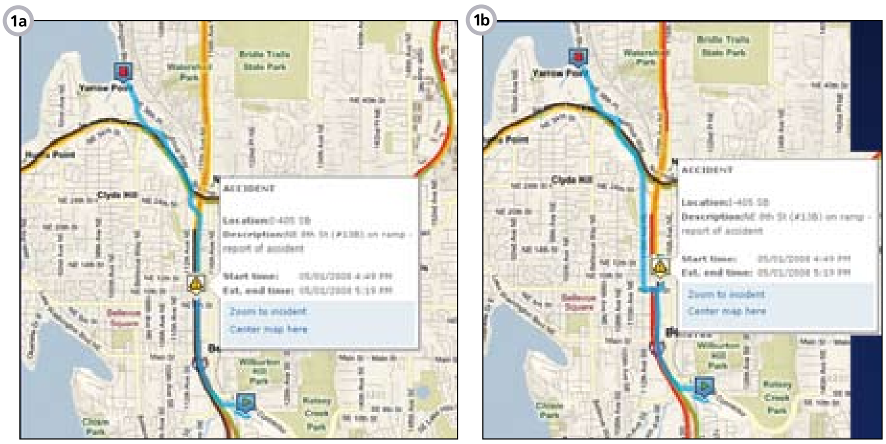

Figure 1a: Route selection identifies a 3.8-mile route, which could be traveled in six minutes at posted speed limits. However, a traffic accident and consequent slower speeds result in a 24-minute trip. **Figure 1b**, meanwhile, demonstrates that route avoidance without knowledge of traffic conditions can actually lead to a longer trip time. An accident along the original route informs the system that an alternate route should be selected. The alternate route is 4.1 miles in length and under ideal conditions would entail a 10-minute trip. However, traffic conditions on this route lead to a true travel time of 28 minutes, four minutes longer than the route with the accident

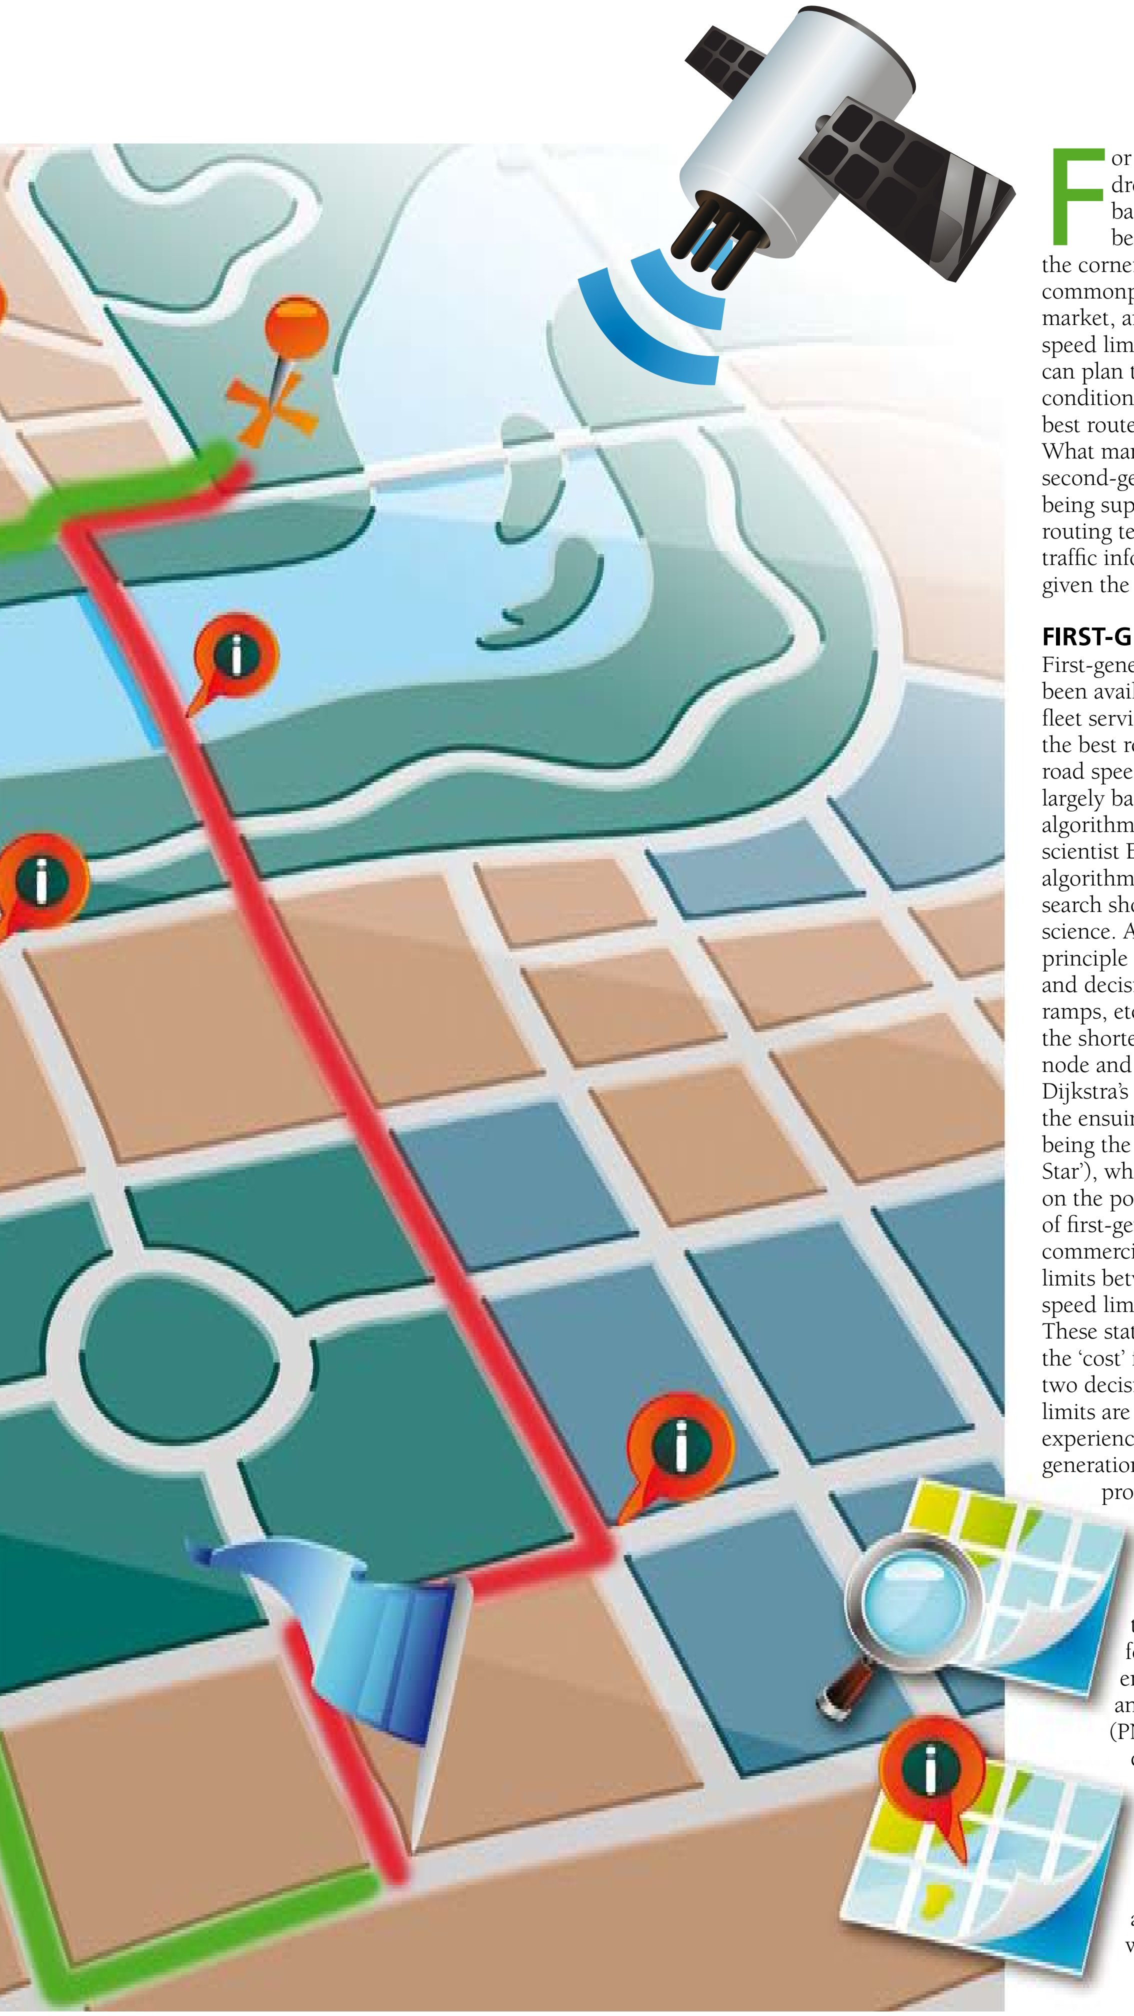

For nearly two decades, the dream of vehicle route navigation based on real-time traffic data has been anticipated as 'just around the corner'. Today such systems, if not commonplace, are at least present in the market, and squeezing out older, static, speed limit-based routing systems. Drivers can plan their travels, confident that traffic conditions are being used to select the best route and ensure the best travel times. What many do not know, however, is these second-generation systems are already being supplanted by third-generation routing technology that uses predictive traffic information to find the best route given the traffic state in the future.

FIRST-GENERATION ROUTING

First-generation routing technology has been available for some years, providing fleet services with the capability to select the best routes for trucks based on known road speed limits. These systems are largely based on the well-known Dijkstra's algorithm, developed by Dutch computer scientist Edsger Dijkstra in 1959. Dijkstra's algorithm was invented to solve the graph search shortest path problem in computer science. Applied to road networks, the main principle is to treat each road as an edge, and decision points (intersections, freeway ramps, etc) as nodes. The algorithm finds the shortest path tree between a starting node and a destination node. Variants of Dijkstra's algorithm have been invented in the ensuing years, with the most popular being the A* algorithm (pronounced 'A Star'), which takes advantage of constraints on the possible routes. Most applications of first-generation routing technology use commercial map databases to derive speed limits between decision points, or estimate speed limits using the known road class. These static speeds are used to compute the 'cost' for picking a route between two decision points. Unfortunately, speed limits are poor proxies for the true speed experienced on a road, so as a result first-generation routing systems commonly provide travel times that are overly optimistic.

This technology is typically encountered in applications that do not have access to real-time traffic conditions – it is used, for instance, by unconnected embedded navigation systems and Portable Navigation Devices (PNDs). Even some connected devices use static speed limits, but may improve results by offering 'route-around' choices when real-time incident alerts are available. However, as Figure 1 illustrates, incident avoidance can actually lead to longer trip times when real-time traffic is unavailable.

SECOND-GENERATION ROUTING

In May 2006, Inrix announced the Smart Dust Network – a GPS probe network that allows drivers across North America to see real-time traffic information. This data enabled a new class of routing technology, using real-time traffic speeds on route segments rather than static, map-based speeds. You can imagine the huge difference this has made in traffic routing accuracy. Rather than selecting a route that shows the highest speeds (most often freeways), routing engines can choose alternate routes when traffic shows highway congestion.

Second-generation systems depend upon an important evolution in navigation devices: newer PNDs receive real-time updates through wireless services. Updates can include news alerts, traffic incident data and, of course, real-time traffic information. A number of applications now use real-time traffic for routing, including deCarta's Traffic Manager, which overlays traffic and incident data on maps to enable traffic-aware routing.

Although these systems have increased route selection and travel-time accuracy by

> "Third-generation routing systems have gone further, now providing lookahead estimates of traffic likely to affect a trip, and traffic-influenced alternate routes when incidents occur"

replacing speed limit data with true real-time traffic conditions, you can imagine that even real-time data does not provide the best information for making routing decisions. Consider that traffic conditions can change quickly, making a route selected at the start of a trip no longer the best choice midway to your destination. Figure 2 below illustrates just such a case.

THIRD-GENERATION ROUTING

Inrix Total Fusion combines real-time traffic information from the Smart Dust Network with traffic prediction using Bayesian networks that model how traffic conditions are affected by both local and regional variables. The models take data from numerous sources – including recent conditions on other parts of the road network, weather, road closures, even sporting events – to create a model of expected conditions in the near future. Total Fusion can provide an estimate of what traffic conditions will be like in 15-45 minutes on core roadways.

This is important because traffic conditions on a route may change, making current conditions either overly optimistic or pessimistic when estimating travel times on route segments. Third-generation routing systems estimate what traffic will be like on a road segment when the driver is expected to arrive at it. As an example, consider a driver leaving for his destination at the

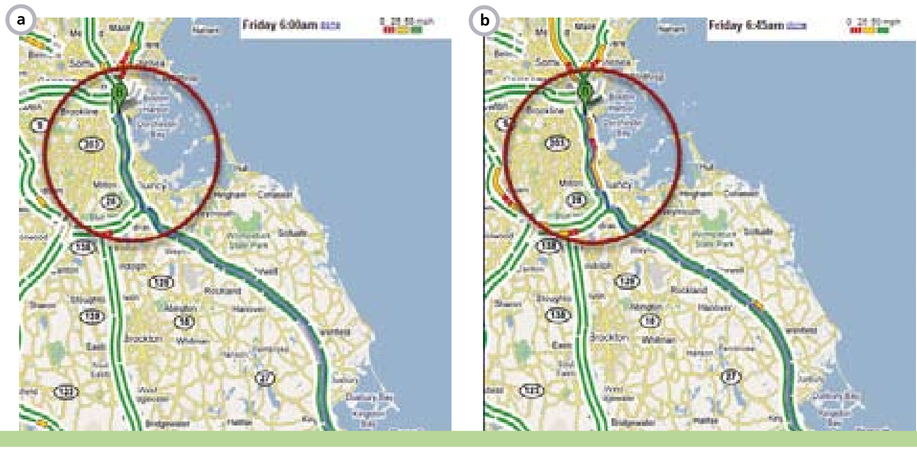

Figure 2: Second-generation routing systems produce traffic-influenced travel-time estimates. For longer routes, however, this may result in inaccurate trip times. The map in a) shows a route from the Sagamore Bridge on Cape Cod to Boston. At 06.00hrs on a Friday morning, the map shows largely free-flowing conditions on the chosen route, and the routing system predicts a 54-minute travel time for this 51-mile trip. However, as b) illustrates, 45 minutes into the trip, the traffic in the vicinity of Boston has become congested, and as a result the original route results in an 85-minute travel time, or two thirds longer than originally planned

IT JUST GOT PERSONAL

Navigation devices have become standard tools for many drivers, but key software and content components remain relatively immature. Unless the software that generates routes and estimates journey times significantly improves in accuracy and begins to deliver on its full potential, there will always be a sense of unfulfilled promise.

The next generation of traffic-speed forecasting technology – predictive traffic – is capable of making a big difference. It understands the normal daily cycle of traffic loading on roads and junctions and uses this information to forecast accurate speeds and speed ranges, for different vehicle types, for every hour of the day and day of the week. This helps navigation devices and web-based journey planners to select the most appropriate route at that time of day.

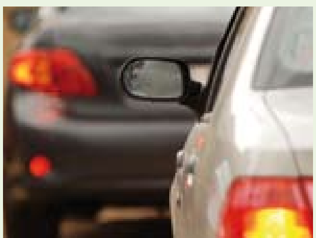

Providing drivers with timely and accurate information is key in reducing traffic congestion

However, while delivering a major improvement over current systems, predictive traffic forecasting will always be inherently inaccurate due to the very marked difference in vehicle performance and – more significantly – driver behavior.

Key to enhancing routing accuracy is understanding that every driver is different – from the individual driving fast on highways but cautiously in urban areas, to the rural expert driving fast on familiar windy minor roads yet slowly on unknown major roads and highways.

Driver personalization systems need to assess not only how a driver behaves on each road type, but also whether or not the route is regularly traveled – because familiarity has an effect on behavior. It is also essential to take into account

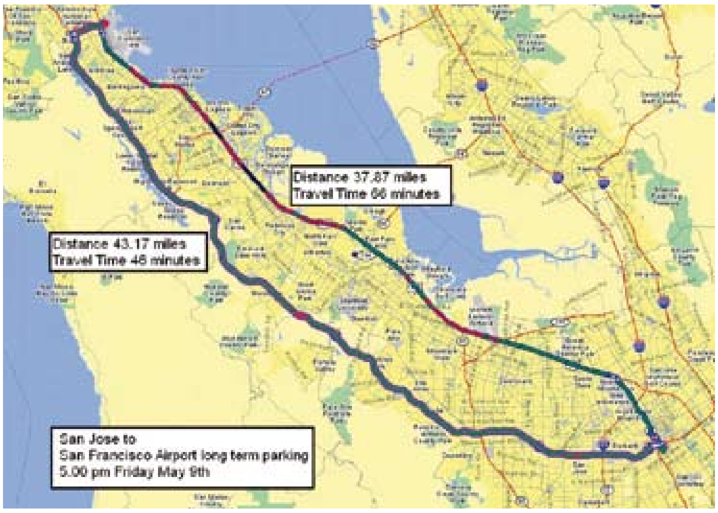

Figure 3: Inrix Total Fusion combines traffic-influenced routing with temporal lookahead. This enables route selection based not only on current traffic conditions, but on what those conditions may evolve to be during the course of the route. This figure shows a map of the San Francisco Bay area between Los Gatos and the outskirts of San Francisco. The traveler would like the best route on a Friday afternoon, from San Jose to the San Francisco airport. Seamlessly blending real-time, historic, predictive and incident data, the Inrix Fusion-based route selects a route of 43.17 miles versus the more direct path of 37.87 miles. The selected route will be 20 minutes faster given both prevailing conditions, and those likely to occur in the subsequent 45 minutes

height of rush-hour. If his route extends beyond the next 15-30 minutes, it may be that current conditions overstate the congestion he will actually encounter in the second half of his trip. The Total Fusion engine can estimate that congestion will subside in the latter half of the trip – offering a route that may currently appear to be heavily congested – as the best overall choice. This non-intuitive result derives from the extensive history used to predict conditions. The system also understands that congestion on a part of the network not included in a route may affect the route in the future. In this way, the impact of accidents peripheral to a trip route can be used to select an alternate route, and improve travel time. Figure 3 shows how the third-generation technology can provide the best route.

In a recent study, ABI Research estimated that 12% of cars will have embedded telematics by 2010. An important distinction among telematics-enabled vehicles are those that possess one-way communications (i.e., receive data) versus those with two-way capabilities. Two-way systems enable the vehicle to both receive up-to-date traffic data and route selections, and deliver current speeds to the traffic data network. An example of this is the new Ford SYNC with Traffic, Directions & Information. Employing the third-generation Routing Engine, it provides the forefront in routing technology: Ford SYNC-enabled cars communicate with Total Fusion to receive current traffic information and routes, and provide updates to the Inrix system.

NAVIGATING THE FUTURE

With the wide availability of accurate real-time traffic flow data in North America, Japan and a few European countries, in-vehicle routing systems have finally come of age. Third-generation routing systems have gone further though, now providing lookahead estimates of traffic likely to affect a road-user's trip, and traffic-influenced alternate routes when incidents occur. Based upon the success of GPS probe network and data fusion technologies, the coverage and quality of real-time and predictive traffic flow data is now rapidly expanding across Europe and some countries in Asia. With the upcoming introductions of two-way connected vehicles and navigation devices, and the growing sophistication of telematics-enabled vehicles, drivers using these new technologies can finally expect to receive accurate turn-by-turn directions, in-city drive-time estimates, and routes optimized for time, distance and fuel efficiency.

Dr Christopher Schofield joined Inrix in October 2008 as principal scientist. He is responsible for the technical vision, new research directions, and fundamental technology of Inrix's products.

time of day: if a driver is on the orbital road of a large city in rush-hour, there is no chance to demonstrate personal driving habits, as the speed is dictated by the surrounding congestion.

However, this accuracy is also dependent upon the validity of the underlying data used for the system. If the speed-forecasting technology is not taking into account all of the complex influences that affect journeys on today's busy roads, then the underlying forecasts cannot achieve the required accuracy and flexibility in the route-planning process, which undermines the relevance of driver personalization.

Taking all these factors into account, driver personalization becomes more precise as it learns each driver's unique habits, increasing estimated journey-time accuracy. For example, if a suggested route is primarily on a highway, the predicted journey time will be proportionately faster for the confident highway driver than for the individual more comfortable on rural lanes.

Giving a driver good quality information enables them to make the best possible route choices

Leveraging this depth of information and personalization also enables drivers to take more control by accepting a route that reflects their driving style.

The significantly improved route accuracy delivered by driver personalization can also transform the financial value of navigation devices, whether portable or line-fit. Car-makers could offer more differentiation by automatically incorporating driver personalization into line-fit devices at the same time as seats and mirrors are adjusted to each driver's settings.

With the accuracy delivered by driver personalization, route acceptance levels increase. Now advertisers can be offered the chance to deliver real-time offers based on a high likelihood that an individual will be taking a specific route at a known time.

Predictive traffic systems are an important step in the maturity of speed forecasting. But only by taking into account the personal factor and incorporating individual driving behavior can navigation devices deliver the accuracy required to meet user expectations and, maximize the potential for targeted location-based advertising.

John Holland is the CEO of UK-based Journey Dynamics