Visualizing the Dynamics of London's Bicycle Hire Scheme

Jo Wood, Aidan Slingsby and Jason Dykes

giCentre, School of Informatics, City University London, EC1V 0HB

Abstract

Visualizing flows between origins and destinations can be straightforward when dealing with small numbers of journeys or simple geographies. Representing flows as lines embedded in geographic space has commonly been used to map transport flows, especially when geographic patterns are important as they are when characterising cities or managing transportation. However, for larger numbers of flows, this approach requires careful design to avoid problems of occlusion, salience bias and information overload. Driven by the requirements identified by users and managers of the London Bicycle Hire scheme we present three methods of representation of bicycle hire use and travel patterns. Flow maps with curved flow symbols are used to show overviews in flow structures. Gridded views of docking station location that preserve geographic relationships are used to explore docking station status over space and time in a graphically efficient manner. Origin-Destination maps that visualise the OD matrix directly while maintaining geographic context are used to provide visual details on demand. We use these approaches to identify changes in travel behaviour over space and time, to aid station rebalancing and to provide a framework for incorporating travel modelling and simulation.

Keywords: Visualization, bicycle, OD map, origin destination, treemap, flow map.

If visual means are used to help in the understanding of traffic flows, techniques are required that summarise the important character of those flows while allowing relevant detail to be identified. This paper proposes a set of techniques that provides *overview, zoom and filter, and details on demand* (Shneiderman, 1996) suitable for exploration of traffic flow data between sets of origins and destinations. The techniques are suitable for exploring transport modelling output and empirically derived flow data. It is illustrated using the London Cycle Hire Scheme (Transport for London, 2010), which provides real-time bicycle docking station status data via the web. Understanding the flows of cyclists using this scheme can have benefits for three groups of people: (a) managers of the scheme who wish to ensure it remains geographically balanced; (b) users of the scheme who wish to make best use of its facilities; (c) modellers and researchers wishing to gain some understanding of urban mobility and structure of a city (Froehlich et al., 2008, 2009). We argue that providing an interactive visual representation of the spatio-temporal flow of cycle traffic allows requirements for all three groups to be at least partially met and we contribute novel designs through which this is achieved.

The scope of this work is determined in part by data availability. We use two sources of data to generate our visualizations. In common with many similar schemes around the world, users of the London scheme are able to remove a bicycle from any docking station distributed around the city. Having acquired a bike, users then typically travel to any other docking station where cycles are docked to complete the journey record. All docking stations continually report their status to a central database as the number of bikes currently docked at any moment in time. This information is relayed via the web for mapping purposes (Transport for London, 2010) and use by the general public. We harvested this data source at five-minute intervals providing the location of all docked bikes throughout central London. Examining change over time provides some indication of the flows of cyclists around London, although it does not provide any explicit information about specific journeys or individual behaviour. Secondly, empirical origin-destination (OD) data between docking stations were provided by Serco – the scheme's operator. Each OD pair identifies the docking stations at each end of the journey and the day and time the bicycle was undocked and subsequently docked. This provides a more explicit record of flows in comparison to the docking station status dataset. We illustrate our approach to visualizing these OD data with 540,000 journeys made between mid-September and mid-October 2010. It is worth noting that the data provide no details about the specific pathway followed between docking stations, nor do they reveal directly, the characteristics of those who make the journey.

Visualizing Flows

Given the importance of flows to transport and movement mapping, it is no surprise that there is long history of embedding flows in cartographic representation (e.g., Tobler, 1987; Rae, 2009). Depending upon the distribution depicted, for small numbers of flows, joining origins and destinations with some form of line symbolisation can be suitable. However as the number of OD pairs increases, and in the case of geographically heterogeneous flows, the results become increasingly cluttered and the 'hairball' views typical of many flow maps in which little meaningful structure is revealed become more likely. In particular, long flows between distant pairs of locations can obscure shorter ones that may share part of the same path. The size of longer lines will also tend to make them more visually salient even though they may be no more significant than shorter flows in terms of process.

Andrienko et al. (2008) identify three approaches to visualizing dynamics, movement and change in large data sets – the depiction of: (a) each record directly; (b) summaries derived through aggregation, generalization or sampling; (c) proto-patterns developed through data-mining. The first of these approaches is increasingly difficult to visualise as flow lines with large data sets "*due to visual clutter and massive overplotting associated with large numbers of cases*" resulting in illegible graphics (Andrienko et al., 2008).

Graphical depictions of summaries or more abstract representations of data do offer solutions to some of the problems associated with flow mapping however and several alternatives have been proposed (see Figure 1).

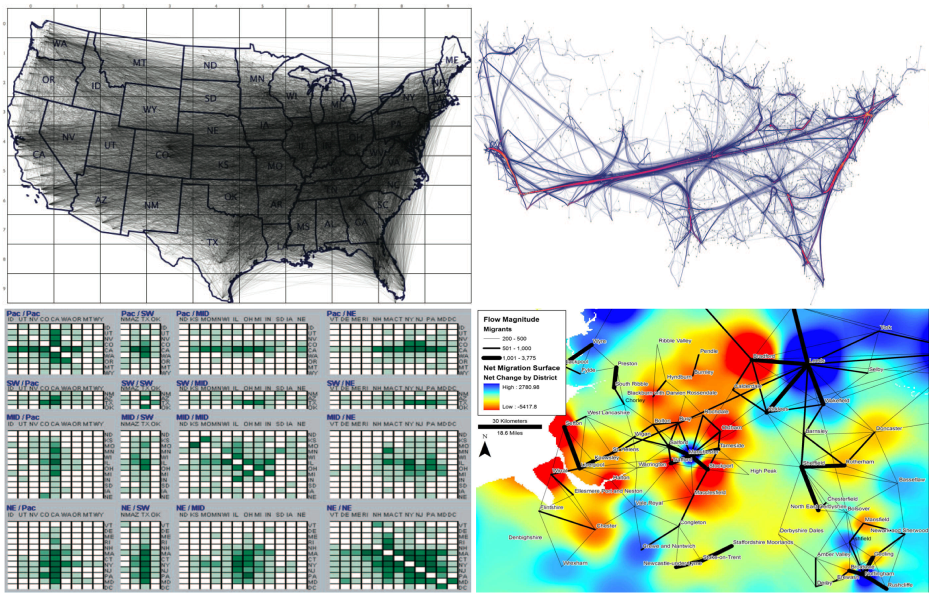

Figure 1: Four approaches to visualising flows between sets of origins and destinations. Clockwise from top-left: Semi-opaque Euclidean flow vectors; Edge bundled flow vectors (source: Holten and Wijk, 2009); Flow density mapping (source: Rae, 2009); OD matrix visualization (source: Guo et al., 2006).

Edge bundling attempts to perform some graphical aggregation where pathways that occupy adjacent locations are combined into a smaller number of smoothed pathways representing general trends (Phan et al., 2005; Telea and Ersoy, 2010). While clutter is reduced, this process can imply a tree structure to flows where none exists and can increase the occlusion of shorter paths by longer ones.

Density mapping of pathways transforms linear flows into a continuous surface (e.g., Rae, 2009) in order to reduce visual clutter. However this technique makes the assumption that the modelled (usually Euclidean straight line) pathway between an OD pair is representative of the actual route between them. For cases where only origin and destination are known, this is potentially a misleading assumption to make. The approach also results in surfaces that make it difficult to distinguish short paths from longer ones – potentially an important distinction to make when understanding travel journey behaviour.

OD matrix (Voorhees, 1955) *visualization* offers an alternative approach that can overcome the occlusion problem and can ascribe equal graphic weight to long and short flows. Each cell in an OD matrix is symbolised with a colour representing the magnitude of flows between the origin and destination pair it represents (Wilkinson and Friendly, 2009). User studies by Ghoniem et al. (2004) have found that matrix visualization tends to be more effective than node-link flow representations for most interpretation tasks. However, they did find that for 'path finding' tasks, such as following a route through sequences of connected origin-destination pairs, the matrix view performed less effectively. One of the reasons for this disadvantage may be due to the commonly applied arbitrary or non-spatial ordering of cells in the OD matrix. Others have suggested that reordering matrix cells can be applied to reveal structure (e.g., Wilkinson, 1979; Bertin, 1983; Marble et al., 1997; Guo and Gahegan, 2006). We have shown that reordering the OD matrix into the OD Map (Wood et al., 2010) can preserve geography (discussed below) and that this provides a dense spatial overview of origin destination structure.

London Cycle Hire Scheme: Requirements and Proposed Solutions

The London Cycle Hire Scheme was launched at the end of July 2010 and by February 2011 comprised around 370 docking stations located in central London. Docking stations vary in their capacity to store bicycles ranging from 10 (Concert Hall Approach, Southbank) to 57 (Royal College Street, Camden) with a mean capacity of 24 bicycles. Total storable capacity for all stations was for about 8800 bicycles in February 2011. Each station is generally no further away than 300m from any other station and coverage in the central London zone is relatively uniform with the exception of the Royal Parks (see Figure 2 top-left). The scheme is managed by Serco who has responsibility for installing and maintaining the infrastructure (docking stations, bicycle fleet, and data transmission) and ensuring that the system is balanced – that is, that there are sufficient bicycles and spaces at each station. Balancing can be controlled by the number of bicycles made available to the scheme (around 6000 in February 2011) and by transporting bicycles between docking stations (using smaller electric vehicles able to transport 20 bicycles at a time and larger diesel vehicles able to transport greater numbers). This can be considered a specific example of fleet management (Crainic and Laporte, 1998), but unlike other systems such as courier services or taxi fleets, there is a degree of autonomous movement of the fleet by the scheme's users.

Once registered on the scheme, the first 30 minutes after releasing a bicycle are free to the user. Keeping a bicycle beyond 30 minutes incurs charges that increase in magnitude over time. The pricing structure and publicity for the scheme encourage short journeys and discourage individuals locking bicycles away from docking stations. Bicycles can be released from and docked at any station. Figure 3 shows the frequency distribution of journey times and station-station distances. The frequency distributions show general patterns of overall usage that are similar to those used to by Mooney et al. (2010) to explore the Dublin bicycle hire scheme. The log-normal distributions of both duration and distance are consistent with the observations found for the Lyon cycle scheme by Borgnat et al. (2011) with approximately similar median values. As was the case in the Lyon study, the London scheme shows a spike for journeys of duration 0 and distance

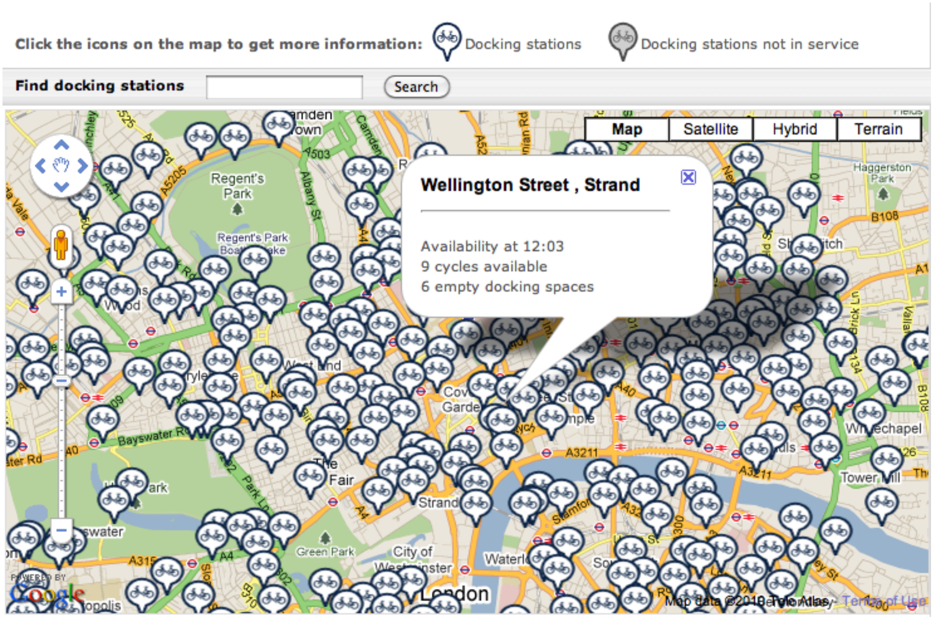

Figure 2: Transport for London's official Barclays Cycle Hire Map (https://web.barclayscyclehire.tfl.gov.uk/maps). Interactive zooming and panning allows precise location of docking stations to be found. Clicking on any station reports its status.

0. Duration 0 hires represent failed attempts to remove a bike, or quick returns of malfunctioning bicycles. Journey distances of 0 km may represent valid journeys however, since bicycles may be returned to the same station from which they were removed (e.g. leisure cycling in circuits).

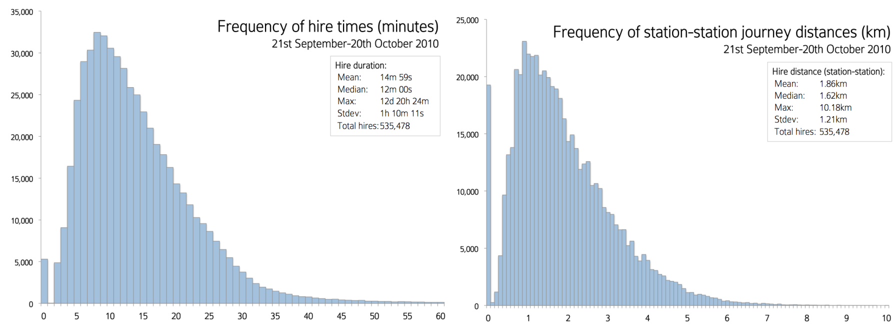

Figure 3: Frequency distributions of hire durations and journey distances. All distances measured as straight line between origin and destination stations. Travelled distances are likely to be much further. Numerical summaries exclude the 5000 failed attempts to release bikes (journey time of 0).

From a scheme user's perspective two key requirements can be established:

U1: Identify the locations of docking stations close to desired journey start and end points
U2: Identify the likelihood of there being a bike available at journey origin and a free docking space available at journey destination

From the manager's perspective two further requirements can be identified:

M1: Ensure the distribution of bicycles and spaces remains balanced.
M2: Minimise the extra rebalancing effort required to meet M1.

Requirement U1 has been met with standard 'mashup' style mapping (e.g., Wood et al., 2007) using map-tiling services such as Google Maps (see Figure 2) or OpenStreetMap (e.g., O'Brien, 2010). Mapping the city-wide docking station status and viewing change over time can give some indirect indication of the flows of journeys between stations. This can help in meeting requirements U1, U2 and M1. However use of this mashup approach means that the graphical space available for station symbolisation is relatively limited and so additional interaction is required to reveal each station's name and status – an 'elaborate' interaction to use the terminology of Yi and Stasko (2007). Understanding local and scheme-wide patterns or comparing status of more than one station can be challenging using this approach. It is therefore particularly unsuited to requirements M1 and M2. Alternative approaches for bicycle station visualization have been proposed such as perspective towers (Girardin et al., 2008) or ordered graphs (giCentre, 2010b). The problem with both approaches is that significant spatial information is either lost or difficult to interpret.

We propose three alternative graphical approaches to meet these needs. These designs summarise data about bicycles and docking stations and are combined in interactive software that provides smooth transitions between them to deliver structured overviews and details on demand.

Flow Maps

Requirements M1 and M2 require some understanding of the structure of bicycle flows over space and time. General trends in structure are useful for managing rebalancing strategies, while detailed information about selected flows allow day-to-day management and reaction to specific bicycle user behaviours. Such journey data also offer valuable information about the urban footprint and mobility around the city (Chiricota et al., 2008; Girardin et al., 2008; Froehlich et al., 2009). To provide a visual overview of general trends, we adopt a number of strategies. Following the second strategy of Andrienko et al. (2008), we combine spatial and temporal aggregation with temporal selection (e.g. by hour of the day or weekday/weekend). We use symbolisation designed to highlight the most significant flows. For any given pair of stations A,B, we symbolise flows from A to B distinctly from B to A and we attempt to remove occlusion of shorter flows by longer ones. Figure 4 shows an overview of all 540,000 origin-destination flows symbolised using our approach.

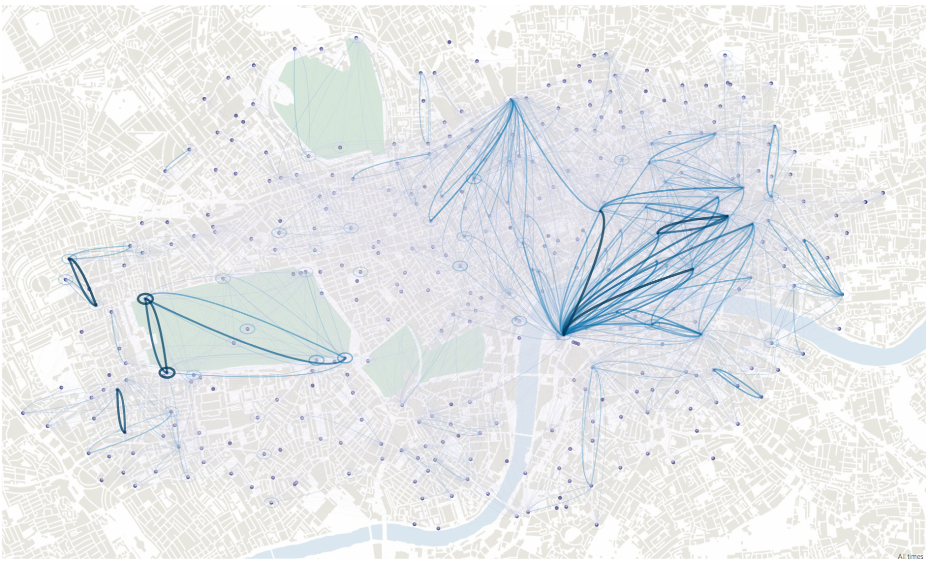

Figure 4: Origin-Destination bicycle flows 21st September - 20th October 2010. Background mapping uses Ordnance Survey data Crown Copyright and database right 2011

Before drawing, origin-destination pairs are ordered from least frequent to most frequent and then drawn in reverse frequency order. That way, small flows, largely irrelevant to M1 and M2 become obscured by larger, more significant flows. To emphasise flow magnitude, flow lines vary in transparency, thickness and colour in proportion to a weighting that depends on the relative frequency of the journey:

$$w_{od} = \left(\frac{f_{od}}{f_{max}}\right)^{1.5} \tag{1}$$

where f_{od} is the frequency of journeys from a given origin o and destination d, f_{max} is the maximum frequency between any origin and destination in the entire dataset. Line thickness is $5w_{od}$ pixels, transparency (scaled between 0=transparent, 1=opaque) is $0.2 + 0.8w_{od}$ and colour is selected from the Brewer 'Blues' sequential colour scheme Brewer and Harrower (2009) using w_{od} scaled between 0=white and 1=dark blue. Empirical experimentation showed that an exponent of 1.5 provided the right balance between emphasis of the dominant flows and the majority of less frequent flows across the city.

The geometry of the flows was symbolised using Bezier curves following the 'link curves' approached used by Fekete et al. (2003) to represent tree structures. Unlike Fekete et al. (2003), we used a tighter radius of curvature (90 degrees rather than 60 degrees from the Euclidean OD axis) and a control point closer to an end point (one sixth the Euclidean distance from destination to origin rather than one half) to ensure a more geographically embedded representation. By making the curves asymmetric, we can distinguish origin (and unlike Fekete et al. (2003) we use the straighter end of the curve) from destination (the sharper end of the curve). This is important in being able to reveal asymmetries in origin-destination patterns between two stations. For example in Hyde Park (large green area left of centre in Figure 4), there are a larger number of flows from the NW to the SE corner of the park than their are in the opposite direction. Because the degree of curvature is a function of Euclidean distance between origin and destination, this form of symbolisation is less likely to result in longer flows occluding shorter ones. Where origin and destinations are the same station, the flow is symbolised by an ellipse orbiting the docking station.

The most significant flows are apparent revealing at lest two distinct patterns of behaviour. Cycling within Hyde Park (left-central part of map) is evident. Temporal selection shows that this behaviour is most common at weekends and peaks around midday and early afternoon. Two main locations of commuter behaviour are evident, associated with the mainline railway stations Kings Cross/St Pancras (top centre of map) and Waterloo (centre of map, just south of the River Thames). Additionally, several elongated loops are evident showing popular journeys between pairs of stations in both directions.

This approach is amenable to temporal selection, for example by aggregating all flows that occur at a particular hour of the day. Figure 5 shows the morning and evening rush hour flow patterns. In addition to the obvious tidal flow in and out of the centre of London during the day, the overviews show that the main morning destinations do not match the evening origin stations. This has been explained by Serco, who examined the maps, as a result of their intervention strategies in providing bicycles at selected popular origin stations in the morning and removing them from popular destination stations in the evening.

There are still some limitations in this form of flow mapping. Using colour and thickness to represent flow frequency gives an overview of relative flow frequency, but is difficult to judge absolute numbers of flows. Interactive zooming and panning around the map helps to some extent along with a labelled frequency histogram (bottom right of Figure 5) , but is limited in conveying absolute flow magnitudes. This approach also does not overcome the *salience bias* produced by longer lines – while representing longer journeys, they may not be indicative of the most important journeys when addressing requirements M1 and M2. Indeed, in terms of load rebalancing, clusters of shorter directed flows may be most significant in terms of meeting managers' needs. We therefore consider an additional approach for showing details on demand that attempts to remove salience bias and all flow occlusion.

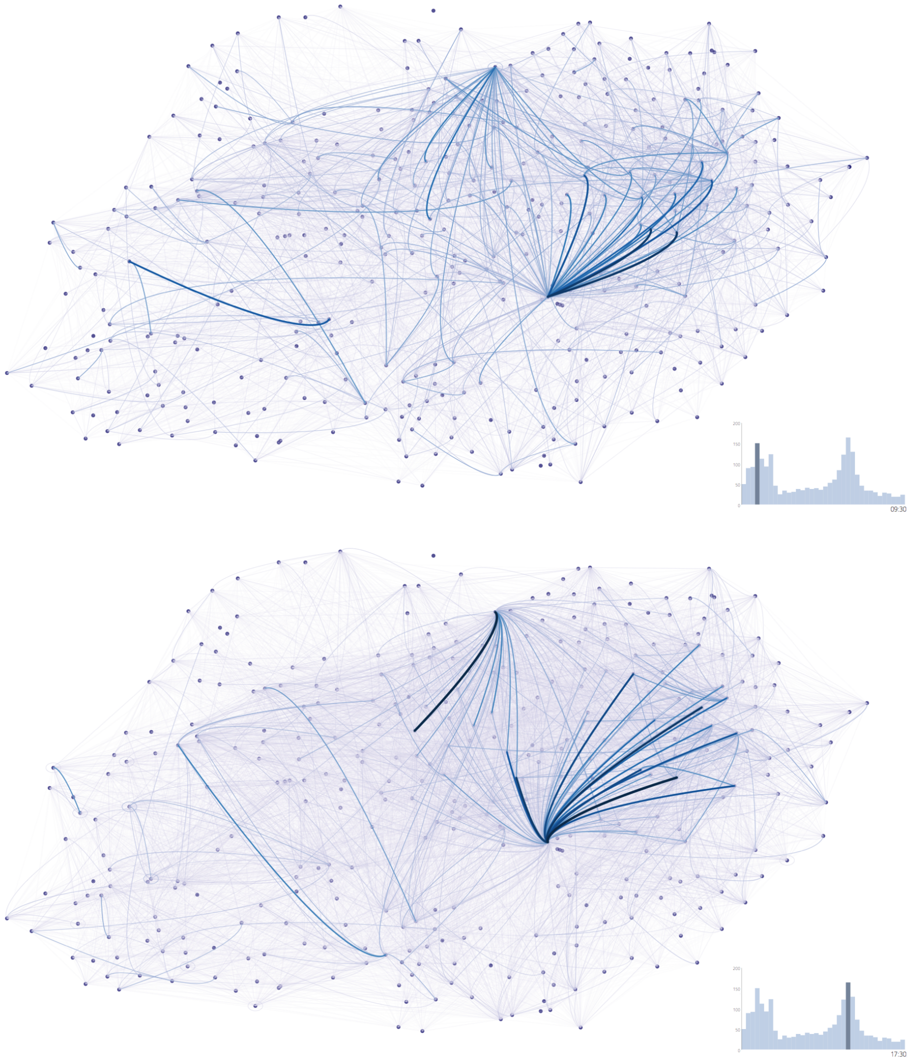

Figure 5: Origin-Destination bicycle flows 21st September 20 - 20th October 2010 showing morning (upper) and evening (lower) bicycle flows.

Spatial Matrices: Geographic Small Multiples

Two additional approaches complement the geographic overview provided by flow mapping. In each we maximise the graphical space available for data depiction by organising each station symbol in a grid while minimising the geographical distortion involved using the gridded spatial layout algorithm of Wood and Dykes (2008). We can use the extra graphical space and the structure provided by the grid to symbolise the status of each station in a manner that is consistent and comparable – through the 'small multiples' of Tufte (1983). We could do so by depicting each data item or by summarising according to the options described by Andrienko et al. (2008). Here, we summarise, showing current and recent status: the height of a blue bar at each station indicates the proportion of bikes currently docked (the higher the blue bar, the greater the proportion of docked bikes). The station's status is further summarised as being in one of three possible states: Full/nearly full indicated by a dark brown border; empty/nearly empty as a light border; normal usable status a mid brown border. This form of symbolisation provides an overview of the scheme's status at any point in time (see Figure 6). It is able to demonstrate the clear tidal flow of usage with stations towards the periphery of the zone being emptied during the morning rush hour along with a filling of the stations towards the centre during the working day. Live results can be seen through our *BikeGrid* application (giCentre, 2010a). This approach, while not showing flows directly, does address requirements U1 and U2 as well as providing details necessary for managing M1 and M2.

We use animated transitions to help relate the geographies of the map and spatial treemap following the findings of Heer and Robertson (2007). Figure 7 shows stills from an animated sequence between conventional geographical layout and gridded layout. The River Thames and major parks are retained in the treemap, and throughout the transition, to help maintain geographic context and continuity through the transformation.

To allow trends over time to be indicated and predicted (necessary for requirements U2, M1 and M2), usage over the previous 24 hours is also depicted for each station (see Figure 8). Without intervention, recent trends (towards the right of each graph) and patterns over the previous day (towards the left of each graph) can be viewed concurrently and spatially. Froehlich et al. (2009) show the importance of recent trends in predicting future station status. Rebalancing of stations (M1 and M2) can be directed towards stations that show strong tidal commuter flows being either emptied or filled rapidly during commuter rush hours (Figure 6).

Addressing U2, users of the scheme can identify stations that are more likely to receive rebalancing (e.g. Belgrove St in Figure 8) or ones that are less likely to fill or empty completely (e.g. Warren St compared with Euston Road in Figure 8).

Hierarchical Spatial Matrices: The OD Map

The gridded spatial treemap can be used to show changes in bicycle volume at each station over time, but it says nothing directly about the journeys made between stations. We adapt the gridded projection of docking stations described above to provide a view of flows that complement the overview provided by flow mapping. To do this we implement an interactive OD Map (Wood et al., 2010). This is an extension of the one-level gridded spatial treemap to two levels. As before, each station is represented as a cell in a grid at its approximate geographic location. Within each of the 400 cells we create a further grid of the same 400 docking station cells using the same layout. These small grids represent all the destination stations at which journeys that started in each large cell (origin station) ended. By colouring each of the destination cells using either a local or global sequential colouring scheme (Harrower and Brewer, 2003), a visual impression of the quantity and spread of journeys can be given (see Figure 9). Importantly, this provides an equal graphic space for all potential origin-destination pairs, so removes the salience bias present in flow maps.

The overview of the OD map reveals that some stations serve as sources of bicycles that are distributed widely over the city. These are often located close to key commuter railway stations (Paddington, Euston, Kings Cross and Waterloo), an inference made by Mooney et al. (2010) in their preliminary analysis of the Dublin scheme. However our spatial overview draws attention to some contrasting local patterns in London – with some stations serving a much more local neighbourhood. Figure 10 shows a zoomed detail demonstrating this contrast. Crosshairs are drawn in each large origin station showing its location 'destination space'. The cells close to the crosshairs always represent short journeys while the further from the crosshairs, the greater the distance between origin and destination station. As we might expect, shorter journeys tend to be more common than longer ones (darker blue cells tend to be closer to the crosshairs). However there are important variations. We can distinguish commuter behaviour from leisure behaviour (see Figure 10) as well as identify the number of journeys made through interaction with the OD map. This provides geographic details on demand that complement the structural overview given by the flow mapping. It allow us to interrogate the flow patterns of geographically focussed areas without occlusion by more distant journeys outside our area of interest. This process supports requirements M1 and M2 where rebalancing activity is planned for specific geographic regions.

Conclusions

We have presented three techniques that offer potential to meet both user and management requirements for the London Cycle Hire Scheme. They incorporate some recent advances in information visualization that make efficient use of graphical space allowing the encoding of additional docking station information and ensure smooth transitions between alternative geographies. We argue that the complexity of the movements of cycle users over space and time benefits from a visual approach to analysis. Unlike more conventional mapping the techniques presented here require some familiarisation before they can be interpreted quickly. We incorporated animated transitions between conventional geographic and gridded views to assist in this familiarisation along with zooming and panning for detailed exploration. We found the ability to add the River Thames and larger parks as significant orientation landmarks helped in this process especially as they both had effects on travel behaviour of cyclists.

Following established good practice in the design of flow mapping allowed structural overviews to be presented that reveal patterns not directly observable without careful spatial and temporal aggregation. OD mapping complements the overview and offers one of the few techniques for visualizing origin-destination data that is able to show both short and longer journeys simultaneously and geographically in a manner that is scalable and unbiased. This has allowed us to distinguish docking stations that appear to serve a more local neighbourhood from ones that act as city-wide hubs and consider ways in which these are geographically distributed – with implications for cycle scheme users and managers.

Several possibilities for further development are offered by these techniques. Modelling longer term daily, weekly and seasonal trends will become possible as the scheme matures (Borgnat et al., 2011). This would allow anomalous behaviour at any given place or time to be identified assisting in meeting both user and management requirements. Equally a range of models, such as those utilised by Froehlich et al. (2009), and their errors can be visualized geographically, on a station-by-station basis and also to consider journeys between all pairs of stations using these approaches. For efficient rebalancing, the gridded projection of space could be integrated with the existing road network and traffic monitoring and modelling. The impact of changes to the scheme such as the introduction of new stations, changing the number of active bicycles and the addition of casual use (credit card hire without scheme membership) could be visualised and compared with observed and modelled patterns. We are interested in understanding in greater depth the barriers and incentives that influence cycling behaviour. These include influences such as weather, perception of safety in response to daylight hours, policy and publicity activities. The framework and techniques suggested here allow the impact of such candidate influences to be assessed visually.

Acknowledgements

We are grateful to Graham Bright of Serco for providing the sample origin-destination data and discussions on management requirements for balancing the scheme. Live data provided with the support of Adrian Short's API for monitoring station status. Base geographical mapping of building outlines and River Thames provided Ordnance Survey's VectorMap District dataset.

References

Andrienko, G., N. Andrienko, J. Dykes, S. Fabrikant, and M. Wachowicz (2008). Geovisualization of dynamics, movement and change: key issues and developing approaches in visualization research. *Information Visualization 7*(3), 173–180.

Bertin, J. (1983). *Semiology of graphics*. University of Wisconsin Press.

Borgnat, P., C. Robardet, J. Rouquier, P. Abry, E. Fleury, and P. Flandrin (2011). Shared bicycles in a city: A signal processing and data analysis perspective. *Advances in Complex Systems 14*(3), 415–438.

Brewer, C. and M. Harrower (2009). Colorbrewer: Color advice for maps. http://colorbrewer2.org.

Chiricota, Y., G. Melançon, T. T. P. Quang, and P. Tissandier (2008). Visual exploration of (French) commuter networks. In *Geovisualization of Dynamics, Movement, and Change, Girona, Spain, 5-8 May 2008*, Spain.

Crainic, T. and G. Laporte (1998). *Fleet Management and Logistics*. Springer.

Fekete, J., D. Wang, N. Dang, A. Aris, and C. Plaisant (2003). Overlaying graph links on treemaps. In *Proceedings of the IEEE Symposium on Information Visualization, Poster Session*, pp. 82–83.

Froehlich, J., J. Neumann, and N. Oliver (2008). Measuring the pulse of the city through shared bicycle programs. In *International Workshop on Urban, Community, and Social Applications of Networked Sensing Systems - UrbanSense08, Raleigh, North Carolina*, pp. 16–20.

Froehlich, J., J. Neumann, and N. Oliver (2009). Sensing and predicting the pulse of the city through shared bicycling. In *International Joint Conference on Artificial Intelligence*, Volume 21, Pasadena, California, pp. 1420–1426.

Ghoniem, M., J. Fekete, and P. Castagliola (2004). A comparison of the readability of graphs using node-link and matrix-based representations. *Proceedings of the IEEEE Symposium on Information Visualization Infovis 2004*, 17–24.

giCentre (2010a). *BikeGrid: Cycle hire docking station viewer, http://gicentre.org/bikeGrid*.

giCentre (2010b). TFL bike availability over the last 24 hours, http://gicentre.org/tfl_bikes.

Girardin, F., F. Calabrese, F. D. Fiore, C. Ratti, and J. Blat (2008). Digital footprinting: Uncovering tourists with User-Generated content. *IEEE Pervasive Computing 7*(4), 36–43.

Guo, D., J. Chen, A. MacEachren, and K. Liao (2006). A visualization system for space-time and multivariate patterns (VIS-STAMP). *IEEE Transactions on Visualization and Computer Graphics 12*(6), 1461–1474.

Guo, D. and M. Gahegan (2006). Spatial ordering and encoding for geographic data mining and visualization. *Journal of Intelligent Information Systems 27*, 243–266.

Harrower, M. and C. Brewer (2003). Colorbrewer.org: an online tool for selecting colour schemes for maps. *The Cartographic Journal 40*(1), 27–37.

Heer, J. and G. Robertson (2007). Animated transitions in statistical data graphics. *IEEE Transactions on Visualization and Computer Graphics 13*(6), 1240–1247.

Holten, D. and J. V. Wijk (2009). Force-Directed edge bundling for graph visualization. *Computer Graphics Forum 28*(3), 983–990.

Marble, D., Z. Gou, L. Liu, and J. Saunders (1997). Recent advances in the exploratory analysis of interregional flows in space and time. In Z. Kemp (Ed.), *Innovations in GIS 4*, pp. 75–88. London: Taylor & Francis.

Mooney, P., P. Corcoran, and A. C. Winstanley (2010). Preliminary results of a spatial analysis of dublin city's bike rental scheme. In M. Haklay, J. Morely, and H. Rahemtulla (Eds.), *Proceedings of the GIS Research UK 18th Annual Conference (London, England, 2010)*, pp. 149–154. University College London.

O'Brien, O. (2010). *London Boris Bikes Map*. http://bikes.oobrien.com.

Phan, D., X. Ling, R. Yeh, and P. Hanrahan (2005). Flow map layout. *IEEE Symposium on Information Visualization Infovis 2005*, 219–224.

Rae, A. (2009). From spatial interaction data to spatial interaction information? geovisualisation and spatial structures of migration from the 2001 UK census. *Computers, Environment and Urban Systems 33*, 161–178.

Shneiderman, B. (1996). The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the IEEEE Symposium on Visual Languages*, pp. 336–343.

Telea, A. and O. Ersoy (2010). Image-based edge bundles: Simplified visualization of large graphs. *Computer Graphics Forum 29*(3), 843–852.

Tobler, W. (1987). Experiments in migration mapping by computer. *The American Cartographer 14*(2), 155–163.

Transport for London (2010). *Barclays Cycle Hire Map*. https://web.barclayscyclehire.tfl.gov.uk/maps.

Tufte, E. (1983). *The visual display of quantitative information.* Graphic Press. Cheshire, CT, USA.

Voorhees, A. (1955). A general theory of traffic movement. *Proceedings of the Institute of Traffic Engineers*, 46–56.

Wilkinson, L. (1979). Permuting a matrix to a simple pattern. *Proceedings of the Statistical Computing Section of the American Statistical Association*, 409–412.

Wilkinson, L. and M. Friendly (2009). The history of the cluster heat map. *The American Statistician 63*(2), 179–184.

Wood, J. and J. Dykes (2008). Spatially ordered treemaps. *IEEE Transactions on Visualization and Computer Graphics 14*(6), 1348–1355.

Wood, J., J. Dykes, and A. Slingsby (2010). Visualization of origins, destinations and flows with OD maps. *Cartographic Journal 47*(2), 117–129.

Wood, J., J. Dykes, A. Slingsby, and K. Clarke (2007). Interactive visual exploration of a large spatio-temporal dataset: Reflections on a geovisualization mashup. *IEEE Transactions on Visualization and Computer Graphics 13*(6), 1176–1183.

Yi, J. S. and J. Stasko (2007). Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics 13*(6), 1224–1231.

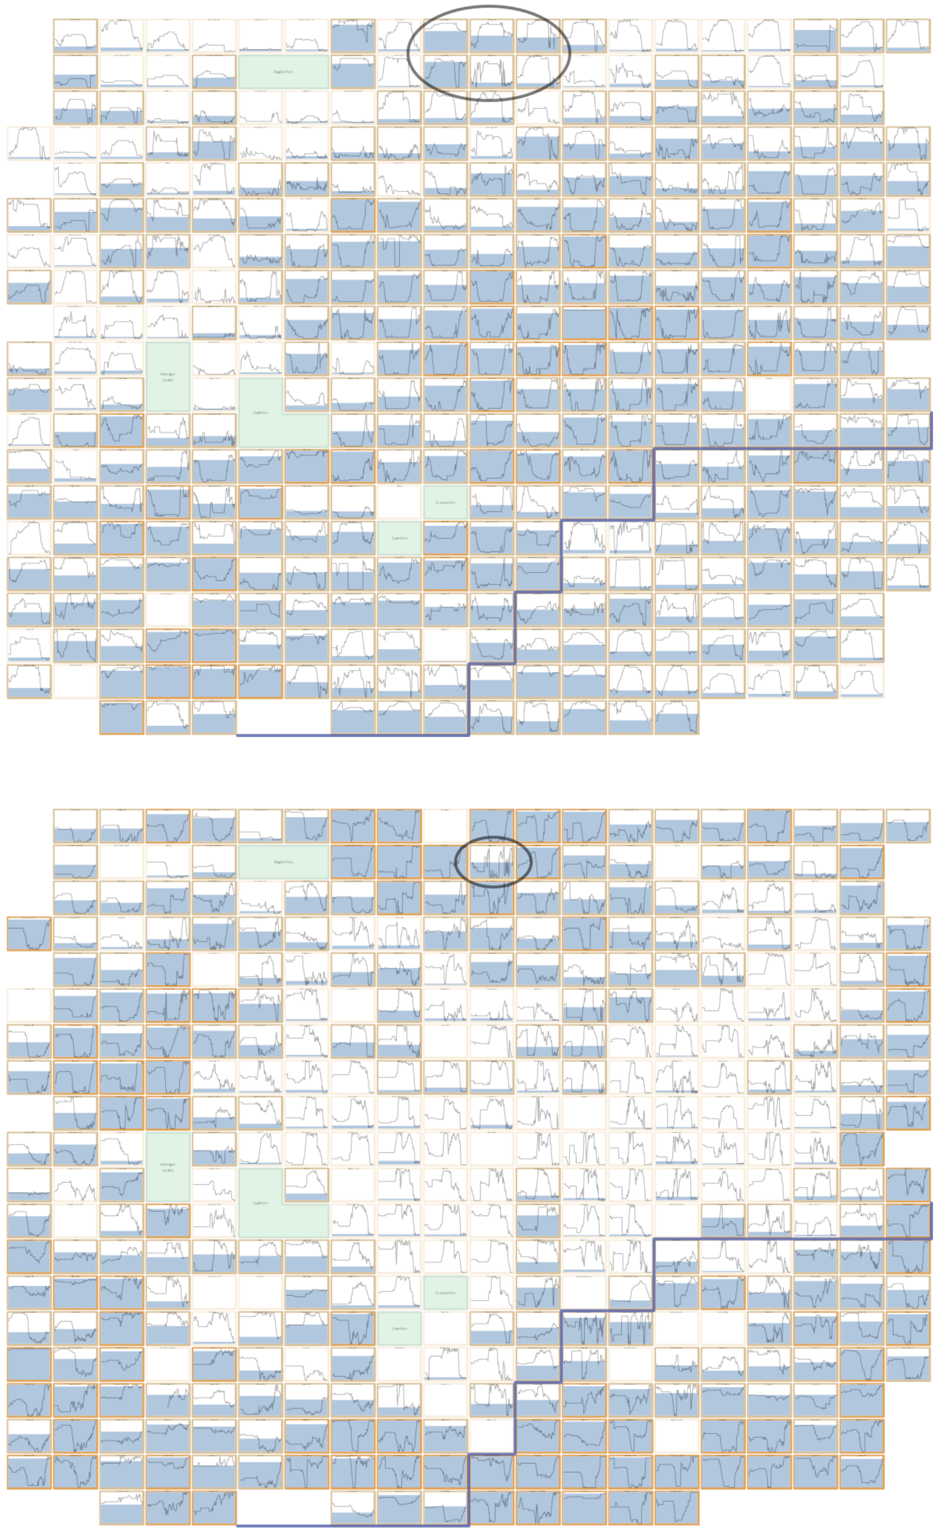

Figure 6: Docking stations' status at 10:30am weekday (top) and 10:30pm weekday (bottom). Evidence of commuter tidal flows evident with some outliers such as Belgrove St (Kings Cross [row 2, col 11]) due to interventions to restock and maintain spaces.

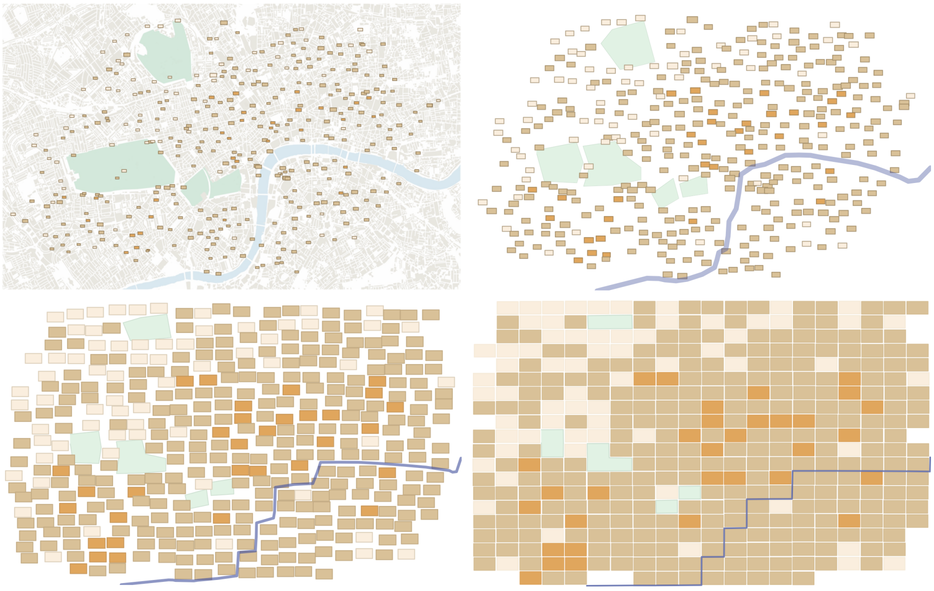

Figure 7: Stills from an animated transition between geographic and grid views of 370 docking station sites (see http://gicentre.org/bikeGrid). Three shades of brown indicate empty, normal and full docking stations; green symbols the Royal Parks and blue the River Thames. Interactive morphing between the two views helps to establish the new gridded geography.

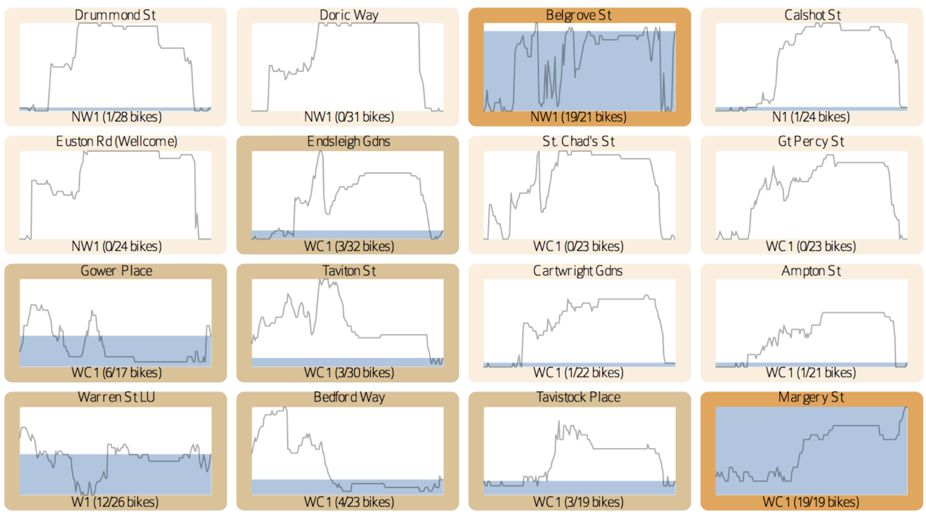

Figure 8: Zoomed in portion of northern-central stations for a 10:30am weekday showing the 'last 24 hours' timelines. Note the rapid fluctuations in Belgrove St indicating intervention in contrast to adjacent Calshot St that shows a 'natural' tidal flow.

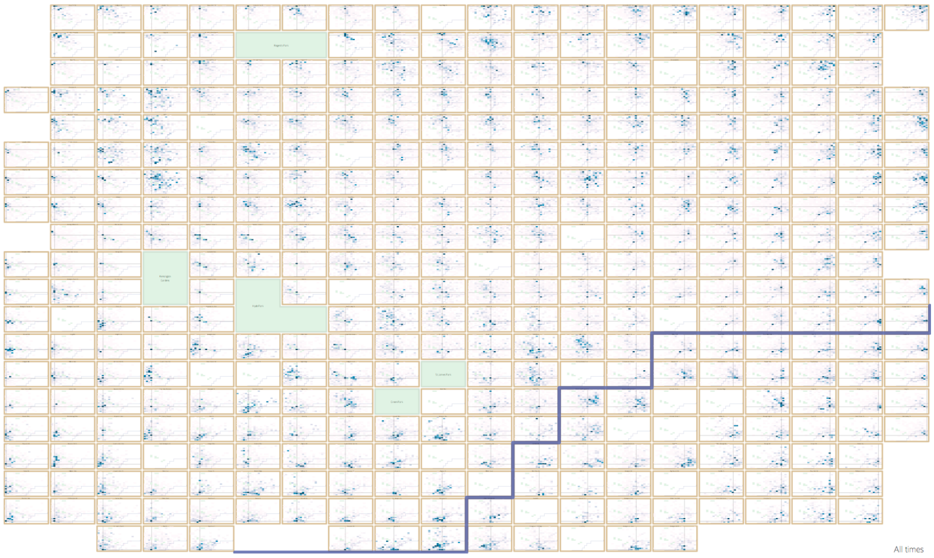

Figure 9: OD Map showing 21st September-20th October journeys between all docking stations. The overview allows stations with a large geographic spread of destinations (e.g. Winsland St [row 7, col 4]) to be distinguished from stations with more local travel (e.g. Porchester Place [row 7, col 5]).

Figure 10: Enlarged section of the OD Map showing docking stations in the Hyde Park area. Cumberland Gate (top-left) shows journeys in and around Hyde Park. In contrast, Green St (top centre) shows a stronger commuting function with journeys to the more distant Kings Cross and Waterloo areas. Grosvenor Square (top-right) shows a high proportion of circular journeys that originate and end at the same station