

# Chapter 7

## Modularity Maximization and Tree Clustering: Novel Ways to Determine Effective Geographic Borders

D. Grady, R. Brune, C. Thiemann, F. Theis, and D. Brockmann

**Abstract** Territorial subdivisions and geographic borders are essential for understanding phenomena in sociology, political science, history, and economics. They influence the interregional flow of information and cross-border trade and affect the diffusion of innovation and technology. However, most existing administrative borders were determined by a variety of historic and political circumstances along with some degree of arbitrariness. Societies have changed drastically, and it is doubtful that currently existing borders reflect the most logical divisions. Fortunately, at this point in history we are in a position to actually measure some aspects of the geographic structure of society through human mobility. Large-scale transportation systems such as trains and airlines provide data about the number of people traveling between geographic locations, and many promising human mobility proxies are being discovered, such as cell phones, bank notes, and various online social networks. In this chapter we apply two optimization techniques to a human mobility proxy (bank note circulation) to investigate the effective geographic borders that emerge from a direct analysis of human mobility.

---

D. Grady  
Northwestern University, Evanston, IL, USA  
e-mail: [d-grady@northwestern.edu](mailto:d-grady@northwestern.edu)

R. Brune • C. Thiemann  
Max Planck Institute for Dynamics and Self-Organization, Göttingen, Germany  
e-mail: [r-brune@northwestern.edu](mailto:r-brune@northwestern.edu); [thiemann@northwestern.edu](mailto:thiemann@northwestern.edu)

F. Theis  
Group leader CMB Professor for Mathematics in Systems Biology,  
Institute of Bioinformatics and Systems Biology, Helmholtz Zentrum München,  
Ingolstädter Landstraße 1, 85764 Neuherberg, Germany  
e-mail: [fabian.theis@helmholtz-muenchen.de](mailto:fabian.theis@helmholtz-muenchen.de)

D. Brockmann (✉)  
Northwestern University, Evanston, IL, USA  
e-mail: [brockmann@northwestern.edu](mailto:brockmann@northwestern.edu)

## 7.1 Introduction

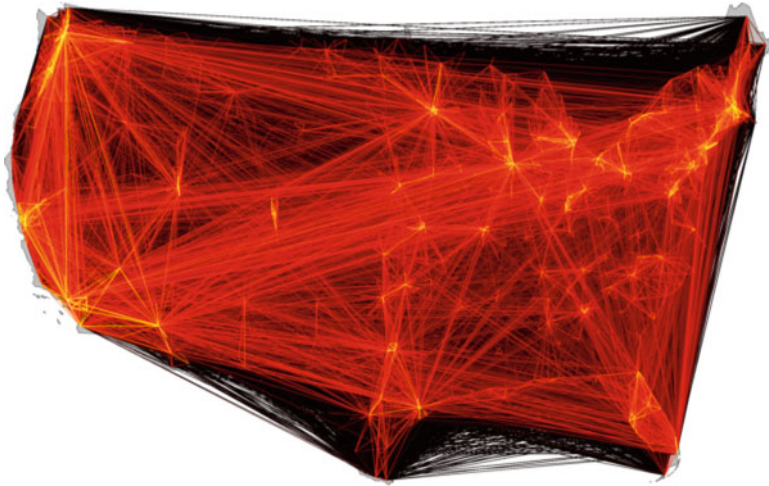
The geographic compartmentalization of maps into coherent territorial units is not only essential for the management and distribution of administrative responsibilities and the allocation of public resources. Territorial subdivisions also serve as an important frame of reference for understanding a variety of phenomena related to human activity. Existing borders frequently correlate with cultural and linguistic boundaries or topographical features [17, 34], they represent essential factors in trade and technology transfer [19, 29], and they indirectly shape the evolution of human-mediated dynamic processes such as the spread of emergent infectious diseases [11, 12, 21, 32].

The majority of existing administrative and political borders, for example in the United States and Europe, evolved over centuries and typically stabilized many decades ago, during a time when human interactions and mobility were predominantly local and the conceptual separation of spatially extended human populations into a hierarchy of geographically coherent subdivisions was meaningful and plausible.

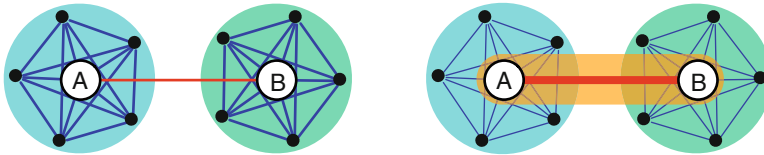
However, modern human communication and mobility has undergone massive structural changes in the past few decades [30, 34]. Efficient communication networks, large-scale and widespread social networks, and more affordable long-distance travel generated highly complex connectivity patterns among individuals in large-scale human populations [3, 9]. Although geographic proximity still dominates human activities, increasing interactions over long distances [7, 25, 38] and across cultural and political borders amplify the small-world effect [31, 41] and decrease the relative importance of local interactions.

Human mobility networks epitomize the complexity of multi-scale connectivity in human populations (see Fig. 7.1). More than 17 million passengers travel each week across long distances on the United States air transportation network alone. However, including all means of transportation, 80% of all traffic occurs across distances less than 50 km [7, 8]. The coexistence of dominant short-range and significant long-range interactions handicaps efforts to define and assess the location and structure of effective borders that are implicitly encoded in human activities across distance. The paradigm of spatially coherent communities may no longer be plausible, and it is unclear what structures emerge from the interplay of interactions and activities across spatial scales [7, 8, 25, 40]. This difficulty is schematically illustrated in Fig. 7.2. Depending on the ratio of local versus long-range traffic, one of two structurally different divisions of subpopulations is plausible. If short-range traffic outweighs long-range traffic, local, spatially coherent subdivisions are meaningful. Conversely, if long-range traffic dominates, subdividing into a single, spatially de-coherent urban community and disconnected suburban modules is appropriate and effective geographic borders are difficult to define in this case.

Although previous studies identified community structures in long-range mobility networks based on topological connectivity [28, 36], this example illustrates that the traffic intensity resulting from the interplay of mobility on all spatial scales must



**Fig. 7.1** The Where’s George? network. Multi-scale human mobility is characterized by dominant short-range and significant long-range connectivity patterns. The illustrated network represents a proxy for human mobility, the flux of bank notes between 3,109 counties in the lower 48 United States. Each link is represented by a line, the color scale encodes the strength of a connection from small (*dark red*) to large (*bright yellow*) values of  $w_{ij}$  spanning four orders of magnitude



**Fig. 7.2** A simplified illustration of generic traffic patterns between and within metropolitan mobility hubs (A and B), with two types of connections  $w_L$  and  $w_D$ , local traffic connecting individual hubs to smaller nodes in their local environment (*blue*) and long-distance links connecting the hubs (*red*). Depending on the ratio of local and long-range flux magnitude, two qualitatively different modularizations are plausible. If  $w_L \gg w_D$ , two spatially compact communities are meaningful (*left*), whereas if  $w_L \ll w_D$ , the metropolitan centers belong to one geographically delocalized module (*orange*), effectively detached from their local environment, yielding three communities altogether (*right*)

be taken into account. Obtaining comprehensive, complete, and precise datasets on human mobility covering many spatial scales is a difficult task, and recent studies have followed a promising alternative strategy based on the analysis of proxies that permit indirect measurement of human mobility patterns [7, 15, 16, 25, 30, 38].

We focus on one human mobility proxy, a dataset collected at the website [www.wheresgeorge.com](http://www.wheresgeorge.com). This website hosts a bill-tracking game called Where’s George? in which participants can tag an individual US banknote of any denomination by logging in to the website and entering the bill’s serial number along with their location. Subsequent participants who receive the bill may do the same, thereby

recording a part of the spatial trajectory the bill follows during its lifetime. We use this information to construct a network whose nodes represent counties in the continental United States and whose edges encode the number of bills exchanged between pairs of counties; details of this and a discussion of some statistics of the data are given in Sect. 7.3.

Both of our analyses rest on the idea of finding community partitions of the network, that is, dividing all of the nodes into a set of mutually disjoint groups or communities. A community of nodes can be defined in many different ways, but all definitions try to capture some aspect of the intuitive idea of a community: a set of nodes that belong together, or are more similar to one another than they are to the rest of the population.

Our first analysis in Sect. 7.2 uses a modularity maximization technique to identify community partitions. Modularity is a method of scoring any given community partition in a network. A partition with a high modularity score has many more intra-group links, and fewer inter-group links, than expected by random chance. Our optimization algorithm searches for high-modularity partitions through a stochastic, simulated annealing process.

We go on to determine community partitions in Sect. 7.6 by searching for nodes with similar topological features, namely their shortest-path tree. Each node is the root of a shortest-path tree that comprises a minimal set of the strongest links connecting that node to the rest of the network. By looking for topological similarities between shortest-path trees, we identify groups of nodes that have similar patterns of connectivity.

With both methods, once a community partition is identified, a corresponding geographic border structure is produced simply by drawing borders between counties that do not belong to the same community, and in Sect. 7.5 we discuss how a superposition of border structures alleviates some of the long-standing weaknesses of modularity maximization. The fact that communities tend to be spatially compact is one of the most surprising findings of this research, and we conclude in Sect. 7.7 by developing a method for comparing border structures and examining the degree to which effective mobility borders line up with various existing borders, such as state boundary lines, census areas, and economic areas.

## 7.2 Network Modularity

This section introduces the modularity measure and describes the simulated annealing algorithm we use for finding maximal-modularity community partitions.

We assume here that  $W$  is a square, symmetric matrix that represents a symmetric, weighted network; the elements  $w_{ij}$  are nonnegative and measure the strength of the connection between nodes  $i$  and  $j$ . Based on the idea that two nodes  $i$  and  $j$  are effectively proximal if  $w_{ij}$  is large, we search for a community partition of the nodes that has a high value of modularity [13, 24, 35]. This standard network-theoretic measure of community structure prefers partitions such that the intra-connectivity

of the modules in the partition is high and inter-connectivity between them is low as compared to a random null model. Given a partition  $P$  of the nodes into  $k$  modules  $M_n$ , the modularity  $Q(P)$  is defined as

$$Q = \sum_n \Delta F_n \quad (7.1)$$

in which  $\Delta F_n = F_n - F_n^0$  is the difference between  $F_n$ , the fraction of total mobility within the module  $M_n$ , and the expected fraction  $F_n^0$  of a random network with an identical weight distribution  $p(w)$ .  $Q$  cannot exceed unity; high values indicate that a partition successfully groups nodes into modules, whereas random partitions yield  $Q \approx 0$ . Maximizing  $Q$  in large networks is an NP-hard problem [6], but a variety of algorithms have been developed to systematically explore and sample the space of possible divisions in order to identify high-modularity partitions [13, 22].

### 7.2.1 Finding Optimal Partitions

As discussed in more detail in Sect. 7.4, our method relies on finding several different high-modularity partitions, which restricts the range of applicable algorithms. For example, the deterministic divisive algorithms described by Newman and Girvan [35] cannot find several different local maxima of the modularity function. In contrast, Monte Carlo algorithms return different partitions with probabilities that monotonically increase with the corresponding modularity values, one of which is the simulated annealing algorithm described by Guimerà and Amaral [27]. Additionally, this algorithm has been found to perform the best in terms of correctly identifying modules in networks with artificial community structure in a survey by Danon et al. [13], which lead us to choose this algorithm for our work.

The partition vector  $P$  is initialized such that each of the  $N$  nodes is in its own module,  $P_i = i$ . Alternatively, one could randomly assign each node to one of a few modules to form the initial partition. We found, however, that in this case the algorithm will split these few large modules into a large number of very small modules before slowly merging them into the final result. Since splits of large modules, involving a recursive simulated annealing run, are computationally very expensive, we avoid them by starting with a partition of single-node modules.

A small modification of the partition is then made (see below) to obtain a new partition  $P'$  and its effect on the modularity value,  $\Delta Q = Q(P') - Q(P)$ . If  $\Delta Q > 0$ , the new partition is better than the old one and we replace  $P = P'$ . If  $\Delta Q < 0$ , the partition is only accepted with probability  $p_T(\Delta Q) = \exp(\Delta Q/T)$ , where  $T$  is a “temperature” that controls the typical penalty on  $Q$  we are willing to accept with the new partition  $P'$ .

This procedure is repeated a number of times, initially with a high  $T = T_0$  accepting modifications with large negative impact on modularity and therefore allowing to sample multiple local maxima. After  $O(N^2)$  modifications, the temperature is

lowered by a *cooling factor*  $c$ . When  $T$  is small enough, worse partitions are not accepted anymore and the partition  $P$  has “annealed” into a local maxima of the modularity landscape  $Q(\cdot)$ .

During each temperature step, we intersperse  $fN^2$  local with  $fN$  global modifications of the partitions, where  $f$  is a tuning parameter. A local modification is a switch of one node to another, randomly selected, module, while a global modification can be a merge of two or a split of one randomly selected module. Finding a suitable split of a module that is not immediately rejected is done by recursively running a simplified version of the simulated annealing algorithm on it: the module in question is extracted and treated as an independent network, initially randomly partitioned into two modules. Only local modifications are allowed while annealing this bipartition into a local modularity maximum. Afterwards, the split module is replaced into the full network and evaluated against the modularity value of the full partition.

We observed that the global structure of the partition is found quickly by the algorithm and mostly only local modifications are accepted at low temperatures. Since the split operations are computationally intensive, we therefore track the number of rejected split modifications in each temperature step and reduce the probability of future trials if that number is high.

To generate the large ensemble of partitions discussed in Sect. 7.5, we used  $T_0 = 2.5 \cdot 10^{-4}$  as initial temperature,  $c = 0.75$  as the cooling factor, and  $f = 0.05$ . We abort the procedure and accept the partition as “optimal” if no better partition is found in three consecutive temperature steps.

The run time of this stochastic algorithm depends in a complex way on both the size *and* the structure of the the input, and therefore the time complexity does not scale with a simple function of the input size. However, we found the algorithm to perform very well and in acceptable runtime (60–90 minutes on a 2.8 GHz processor for most runs) with the configuration given above, although these parameters are less conservative than those proposed by Guimerà et al. [27]. The large ensemble of resulting partitions (Fig. 7.15) has a tight distribution of modularity values, indicating the algorithm tends to converge onto a stable maximum.

### 7.3 A Proxy for Multiscale Human Mobility Networks

Here, we construct a proxy network for human mobility from the geographic circulation of banknotes in the United States. Movement data was collected using the online bill-tracking game at [www.wheresgeorge.com](http://www.wheresgeorge.com). Individuals participating in this game can mark individual bills and return them to circulation; other individuals who randomly receive bills can report this find online along with their current location (zip code). Our analysis is based on the intuitive notion that the coupling strength between two locations  $i$  and  $j$  increases with  $w_{ij}^H$ , the number of individuals that travel between a pair of locations per unit time, and furthermore that the flux of individuals in turn is proportional to the flux of bank notes,

**Table 7.1** Denominations in the WG dataset

Denomination	\$1	\$2	\$5	\$10	\$20	\$50	\$100
Number of bills	9,931,261	36,639	1,069,427	401,101	461,076	24,209	26,526
Fraction [%]	83.11	0.31	8.95	3.36	3.86	0.20	0.22

**Table 7.2** Absolute number and relative fraction of bills based on Federal Reserve Bank

FRB Code	Location	Count	Fraction [%]
A	Boston	799,537	6.69
B	New York City	1,325,942	11.10
C	Philadelphia	822,340	6.88
D	Cleveland	661,278	5.53
E	Richmond	948,516	7.94
F	Atlanta	1,565,732	13.10
G	Chicago	1,207,448	10.10
H	St. Louis	472,930	3.96
I	Minneapolis	360,194	3.01
J	Kansas City	713,393	5.97
K	Dallas	869,866	7.28
L	San Francisco	2,203,063	18.44

denoted by  $w_{ij}$ . Evidence for the validity of this assumption has been obtained previously [7, 8, 25] and we provide further evidence below.

As of January 15th, 2010 a total of 187,925,059 individual bills are being tracked at the website [www.wheresgeorge.com](http://www.wheresgeorge.com). Approximately 11.24% of those have had “hits”, that is they were reported a second time at the site after initial entry. The current analysis is based on a set of  $N_0 = 11,950,239$  bills that were reported at least a second time. For each bill  $n$  we have a sequence of pairs of data

$$B_n = \{Z_{n,i}, T_{n,i}\} \quad i = 0, \dots, L_n \quad n = 1, \dots, N_0$$

of zip codes  $Z_{n,i}$  and times  $T_{n,i}$  at which the bill was reported. Each  $B_n$  reflects a geographic trajectory of a bill with  $L_n$  individual legs. In total, we have 14,612,391 single legs in our database. Note that the majority (81.78%) of trajectories are single-legged reflecting a reporting probability of  $\approx 20\%$  during the lifetime of a bill.

The set of  $B_n$  represents the core dataset of our analysis. For each bill we have additional information:

1. Denomination: \$1, \$2, \$5, \$10, \$20, \$50, or \$100. The fraction of each denomination is depicted in Table 7.1.
2. The Federal Reserve Bank code, A through L, corresponding to one of 12 of the United States Federal Reserve Banks that issued the bill. The fraction of bills as a function of FRB origin is provided in Table 7.2.

We restrict the analysis to the lower 48 states and the District of Columbia (thus excluding Hawaii and Alaska) and consider only legs with origin and destination locations in these states, reducing the original dataset to 11,759,420 bills (98.40% of the original data) and 14,376,232 trajectory legs (98.38%).

The spatial resolution of the dataset is given by 41,106 zip codes, with mean linear extent of 14 km. The mean linear extent of the lower 48 states is 2,842 km defining the bounds of the system. For each zip code  $Z_i$  we use centroid information to associate with each report a longitude/latitude location  $\mathbf{x} = (\Theta, \phi)$ , such that each trajectory  $n$  corresponds to a sequence of geographic locations  $\mathbf{X}_i$  with  $i = 1, \dots, L_n$ :

$$t_n : \{ \mathbf{X}_{n,0}, \Delta T_{n,1}, \mathbf{X}_{n,1}, \dots, \Delta T_{n,L_n}, \mathbf{X}_{n,L_n} \} \quad \text{with } n = 1, \dots, N_0, \quad (7.2)$$

where  $\mathbf{X}_{n,0}$  is the initial entry location, and  $\Delta T_{n,i} = T_{n,i} - T_{n,i-1}$  are inter-report times.

### 7.3.1 Geographical Distributions

Based on these trajectories we define the density of initial entries as

$$p_{\text{IE}}(\mathbf{x}) = \frac{1}{N} \sum_{n=1}^N \delta(\mathbf{x} - \mathbf{X}_{n,0}), \quad (7.3)$$

and the density of reports as

$$p_{\text{R}}(\mathbf{x}) = \frac{1}{N} \sum_{n=1}^N \frac{1}{L_n} \sum_{i=1}^{L_n} \delta(\mathbf{x} - \mathbf{X}_{n,i}), \quad (7.4)$$

where  $\delta$  is the Dirac delta function, equal to 1 when its argument is 0 and equal to 0 otherwise.

In order to assess the spatial distribution of reports and initial entries and to quantify the correlation with the population density we compute the number of reports and initial entries for each of the  $M = 3,109$  counties in the lower 48 states. Defining for each county  $k$  a characteristic function

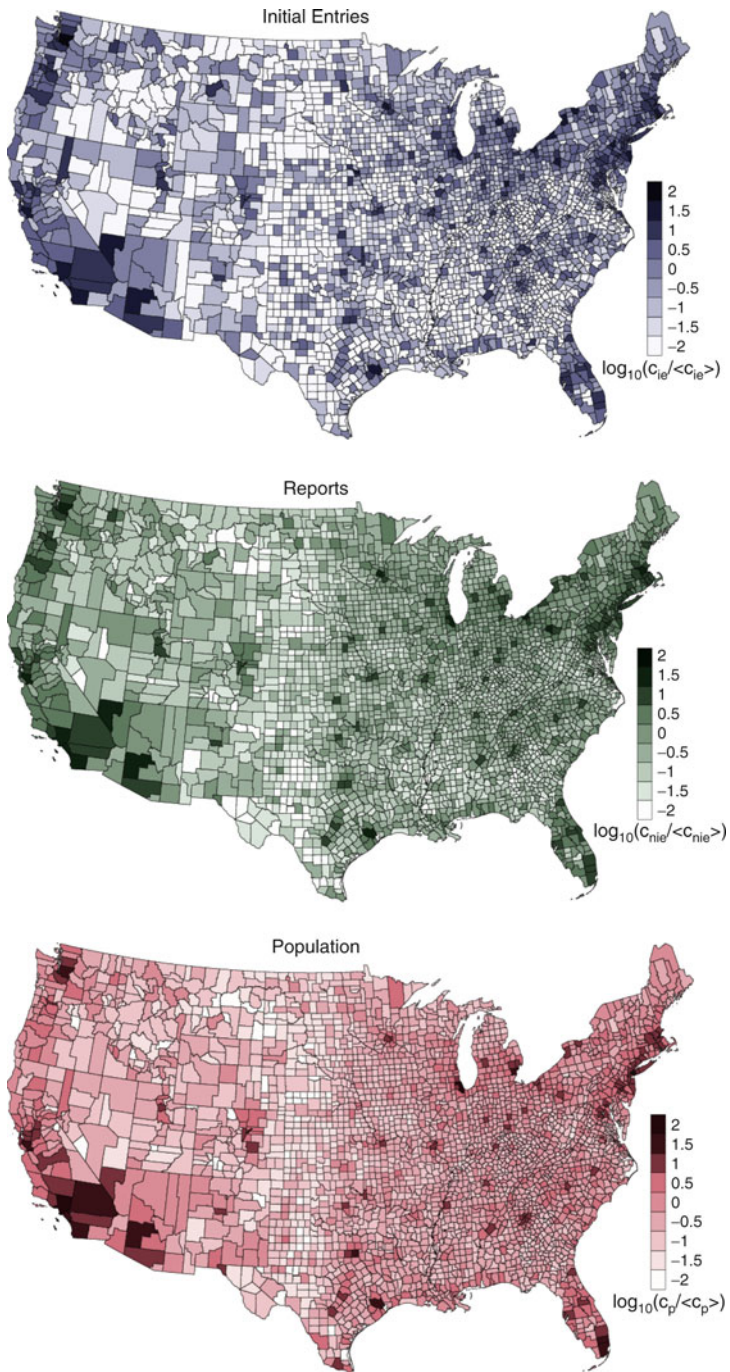
$$\chi_k(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in P_k \\ 0 & \text{otherwise} \end{cases} \quad (7.5)$$

where  $P_k$  is the polygon defining the county's interior, the number of reports and initial entries in county  $k$  are given by

$$m_{\text{R}}(k) = \langle \chi_k \rangle_{\text{R}} = \int \chi_k(\mathbf{x}) p_{\text{R}}(\mathbf{x}) \, \text{d}\mathbf{x} \quad \text{and} \quad m_{\text{IE}}(k) = \langle \chi_k \rangle_{\text{IE}} = \int \chi_k(\mathbf{x}) p_{\text{IE}}(\mathbf{x}) \, \text{d}\mathbf{x},$$

respectively. Figure 7.3 compares the distribution of reports  $m_{\text{R}}(k)$ , initial entries  $m_{\text{IE}}(k)$  and the population  $P(k)$  of the 3,109 counties. As all three quantities are





**Fig. 7.3** The frequencies of reports (*top*) and initial entries (*middle*) correlate with the county population (*bottom*) in the lower 48 states

positive and vary over many orders of magnitude, the maps depict  $\log_{10}(m_R)$ ,  $\log_{10}(m_{IE})$  and  $\log_{10}(P)$ . Qualitatively, reports and initial entries correlate strongly with the population density. Computing the correlation coefficient of the logarithmic quantities yields  $c(R, P) = 0.933$  and  $c(IE, P) = 0.819$ . Despite the expected increase of  $m_R(k)$  and  $m_{IE}(k)$  with  $P(k)$ , only the report count increases approximately linearly with population size, whereas initial entries show a deviation for small populations. We believe that this deviation is a consequence of the social difference between the subpopulation of “Georgers” that are responsible for initiating bills and entering them into the system, “actively” playing the game, and the larger group of people that randomly receive a bill and report it, “passively” participating. This hypothesis could explain that areas with higher population densities contain a larger proportion of internet-savvy communities that are inclined to become Georgers and initiate bills. In order to exclude a potential bias caused by this effect we exclude all the legs in (7.6) that contain an initial entry as the origin, i.e. we only investigate the reduced set

$$t_{2,n} : \{ \mathbf{X}_{n,1}, \Delta T_{n,2}, \mathbf{X}_{n,2}, \dots, \Delta T_{n,L_n}, \mathbf{X}_{n,L_n} \} \quad \text{with } n = 1, \dots, N_0, \quad (7.6)$$

that excludes the first legs of all  $t_n$ . Excluding the first leg reduces the number of bills to 4,743,330. However, the key results, for example the border structures discussed in Sect. 7.5, are robust against the inclusion of initial entries. Computing mobility networks based on either set,  $t_n$  or  $t_{2,n}$  does not change the observed pattern significantly.

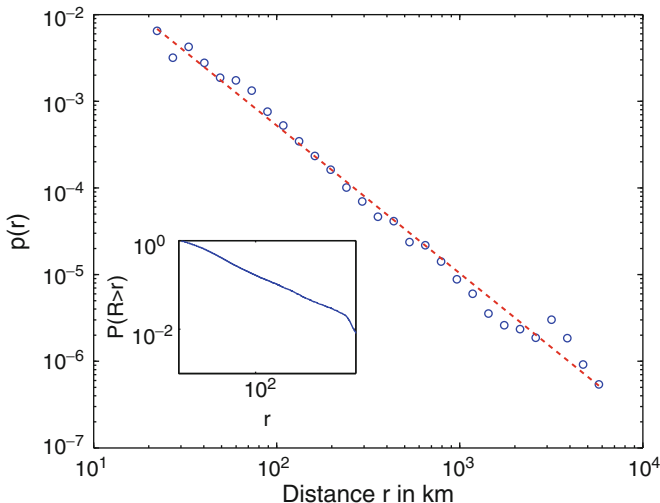
### 7.3.2 Distance and Time: Spatially Averaged Quantities

From  $t_{2,n}$  we extract pairs of spatio-temporal leg distances  $\{d_s(\mathbf{X}_{n,i}, \mathbf{X}_{n,i-1}), \Delta T_{n,i}\}$ , where  $d_s(\cdot, \cdot)$  denotes the distance on a sphere (shorter segment of the great circle that passes through both points). This type of dataset was first investigated in 2006 based on a much smaller core dataset of bill trajectories [7]. In particular, the combined probability density (pdf)

$$p(r, t) = \langle \delta(r - d_s(\mathbf{X}_{n,i}, \mathbf{X}_{n,i-1})) \delta(t - \Delta T_{n,i}) \rangle, \quad (7.7)$$

was estimated as well as marginal pdfs  $p(r)$  and  $p(t)$ . The central finding of the 2006 study was that  $p(r) \sim r^{-(1+\beta)}$  and that the time evolution of the density (7.7) can be described by a bi-fractional diffusion equation. Here, we reproduce some of the properties before we construct the mobility network used in the main text. Figure 7.4 shows the short time pdf of a bill traversing a distance  $r$  in a time  $t < \tau$  where we chose  $\tau = 4$  days. Using maximum likelihood we find this function can be described by a power-law

$$p(r) \sim \frac{1}{r^{1+\beta}} \quad \text{with } \beta = 0.7056 \pm 0.0659.$$



**Fig. 7.4** The estimated probability  $p(r|t < \tau)$  of a bill traversing a distance  $r$  in time  $t < \tau$  where  $\tau = 4$  days. In red a maximum likelihood fit of the the function  $p(r) \sim r^{-(1+\beta)}$  with  $\beta = 0.7056$

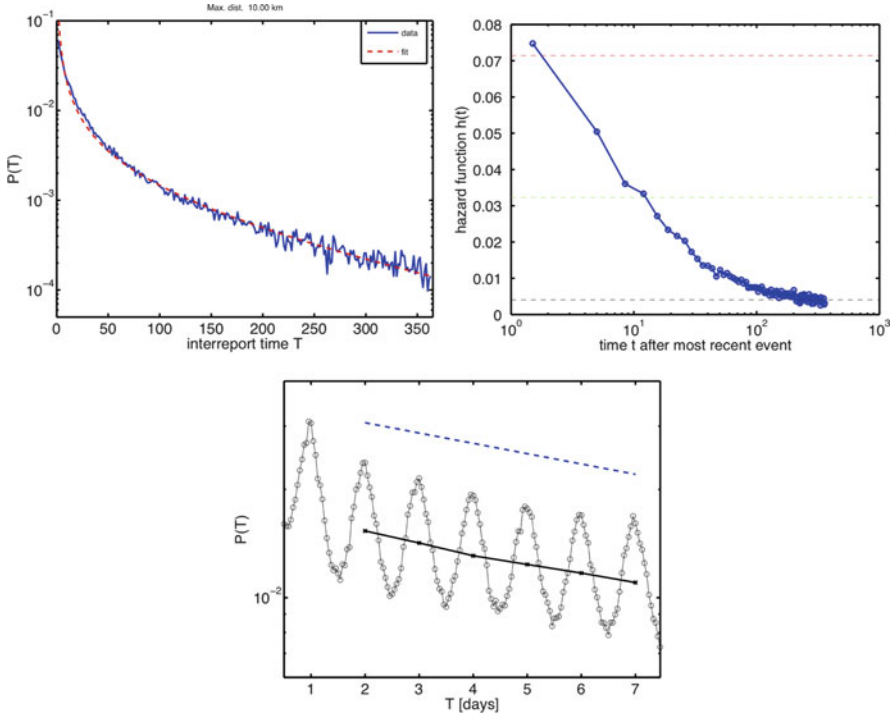
This power law describes the dispersal characteristics on a population-averaged level. The short-time distance pdf represents a dispersal kernel and for small times  $t$  approximates the instantaneous rate of traversing a distance  $r$ .

Complementary to this, temporal aspects of the process can be revealed by computing the pdf for the time  $t$  between reporting events given that these occur within a small radius  $r > r_0$ . Figure 7.5 depicts  $p(t)$  for all legs with  $r < 10$  km and a minimal inter-report time of  $t_{\min} = 1$  day. The inter-report times are described well by a power law moderated by an exponential factor

$$p(t) \sim t^{-\alpha} e^{-t/T_0} \quad \text{with} \quad T_0 = 248 \pm 27, \quad \alpha = 0.99 \pm 0.05. \quad (7.8)$$

The observed power-law decay  $\sim t^{-1}$  for times  $t \ll T_0$  is intriguing. These type of decays have been observed in a multitude of contexts involving human activity, for instance the time between consecutive phone calls [25], emails, [5] and the number of words between two identical words in texts [1]. A consequence of this law is bursting behavior, i.e. given an event occurred at time  $t_0$  the probability rate that an event occurs immediately after the first is higher than expected from ordinary Poisson statistics. This behavior is best illustrated by the so-called hazard function  $h(t)$  that quantifies the instantaneous probability rate of an event happening at  $t$  given that the last event occurred at  $t = 0$ . If we let

$$P(\tau > t) = \int_t^\infty p(s) ds,$$



**Fig. 7.5** Inter-report time statistics. (Left) The function  $p(t|r < r_0)$  for  $r_0 = 10$  km. The observed function can be accounted for by an initial algebraic decay  $t^{-1}$  moderated by an exponential function for large arguments. The red dashed curved is a fit obtained from maximum-likelihood estimation. (Right) The hazard function  $h(t)$  that represents the instantaneous rate of an event at time  $t$  provided that an event occurred at  $t = 0$ . The dashed lines represent reporting rates of once per 2 weeks (top), once per month (middle) and once per  $T_0 = 248$  days (bottom). (Bottom)  $p(t)$  for very short times. A zoom-in resolves daily oscillations modulated by the decay observed on the left. These oscillations indicate that users tend to report to the website at the same time of the day with the highest probability

be the cumulative probability that the second event occurs at a time  $\tau$  later than  $t$ , the hazard function is defined by

$$P(\tau > t) = e^{-\int_0^t h(s) ds}.$$

For a Poisson process with rate  $\gamma$  we have

$$h(t) = \gamma \Rightarrow P(\tau > t) = e^{-\gamma t}.$$

The hazard function can be computed according to

$$h(t) = -\frac{d}{dt} \log [P(\tau > t)] = \frac{p(t)}{P(\tau > t)}.$$

Figure 7.5 depicts the function  $h(t)$  for inter-report times in the WG data. For small times ( $t < 1$  week) the probability rate for a report is of the order of one report per two weeks, which is also the expected time between two reports in this time window. For larger times ( $t > 100$  days) the constant value of  $1/T_0$  is approached, equivalent to one report in  $3/4$  of a year. Possible explanations of the bursting behavior and the initial algebraic decay in  $p(t)$  are a strong behavioral heterogeneity of players that participate in the game or an effective queueing in the system, i.e. bills may enter shops and initially have a comparatively high likelihood of leaving, being “on top of the stack.” As time passes these bills may “get stuck” and equilibrate to the long time scale present in the system.

### 7.3.3 Definition of the Mobility Network

From the trajectories defined by (7.6) and the characteristic functions of the counties (7.5) we construct a matrix  $\tilde{w}_{ij}$  that counts the number of legs which originate at county  $i$  and terminate at  $j$ ,

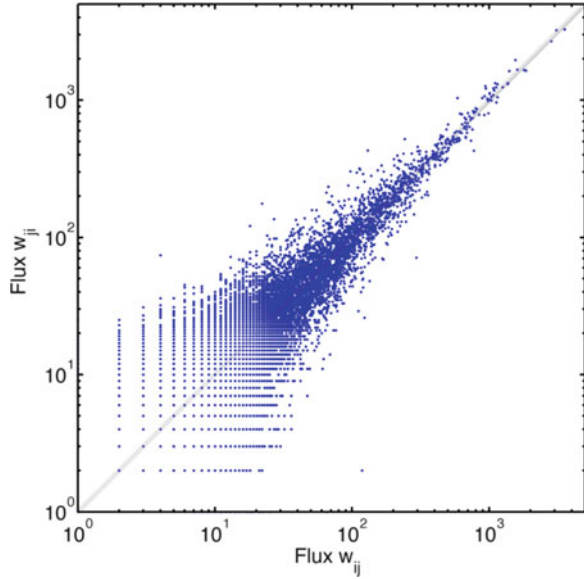
$$\tilde{w}_{ij} = \sum_{n=1}^N \sum_{k=2}^{L_n} \chi_i(\mathbf{X}_{n,k-1}) \chi_j(\mathbf{X}_{n,k}) \Theta(T - \Delta T_{n,k}),$$

where  $\Theta(\cdot)$  is the Heaviside step-function. In order to exclude potential biases induced by initial entries we ignore the first leg of all trajectories ( $k = 2$  in the above sum). This choice is motivated by the fact that the community of individuals that initiate bills might be less representative than those that find bills and report them. Indications that this might have an effect are supported by the different scaling behavior of initial entry frequencies with population as compared to report frequencies with population. The factor  $\Theta(T - \Delta T_{n,k})$  excludes legs that have an inter-event time larger than time  $T$ . The matrix  $\tilde{w}_{ij}$  need not to be symmetric, as the flux of bills from  $i \rightarrow j$  need not equal those that travel  $j \rightarrow i$ . However, as Fig. 7.6 indicates the flux matrix is statistically symmetric. Plotting  $\tilde{w}_{ij}$  against  $\tilde{w}_{ji}$  indicates a clear mean linear relationship. Since we base our analysis on the flux of money between two given counties we symmetrize the network and use  $w_{ij}$  in our analysis defined by

$$w_{ij} = \frac{1}{2} (\tilde{w}_{ij} + \tilde{w}_{ji}),$$

which of course also depends on the time threshold parameter  $T$ . Choosing the optimal value for  $T$  is a trade-off between trying to estimate instantaneous flux, i.e. choosing  $T$  as small as possible, and using as many legs as possible to decrease fluctuations, i.e. choosing large values for  $T$ . Choosing a value for  $T < 30$  days for instance rules out bills that visit a Federal Reserve Bank in between reports in counties  $i$  and  $j$ , as bills that enter FRBs do not return to circulation until approximately 3–4 weeks after entering the FRB. To make sure that our results do

**Fig. 7.6** Symmetry of flux network  $\tilde{w}_{ij}$



not significantly change as the parameter  $T$  is varied we performed the analysis for various values of  $T$  ranging from a few days to  $T = 1$  year. The computed border structure does not significantly depend on the value of  $T$ . Decreasing  $T$  thins out the network and reduces the overall connectivity, yet the effects are similar to bootstrapping the network randomly, a process that also does not change our results and is discussed in Sect. 7.7.1.

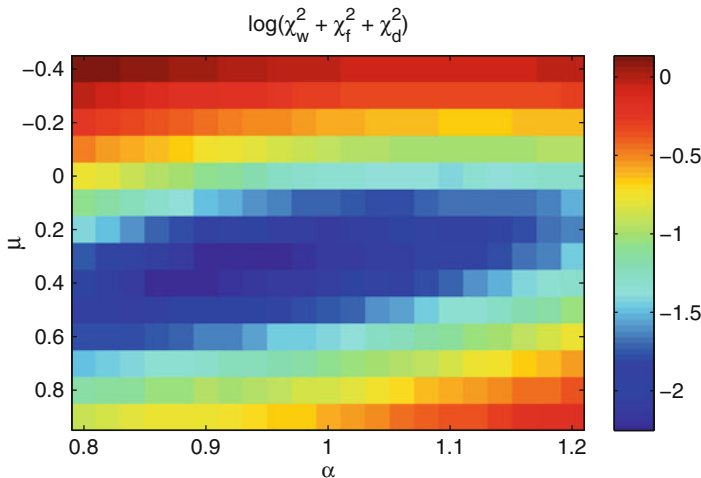
### 7.3.4 Gravity as a Null Model

In addition to the empirical data described above, we also construct a synthetic mobility network based on the well-known gravity law hypothesis [2, 10, 42] to serve as a null model. In gravity models, the interaction strength between a collection of sub-populations with geographic positions  $x_i$ , sizes  $N_i$  (obtained from census data,<sup>1</sup>) and distances  $d_{ij} = |x_i - x_j|$  is given by

$$p_{ij} \propto \frac{N_i^\alpha N_j^\beta}{d_{ij}^{1+\mu}} \quad (7.9)$$

in which  $\alpha$ ,  $\beta$ , and  $\mu$  are non-negative parameters.

<sup>1</sup><http://www.census.gov>.



**Fig. 7.7**  $\chi^2$  goodness-of-fit for different parameters of the gravity law. The minimum is at  $(\alpha, \mu) = (0.96, 0.3)$

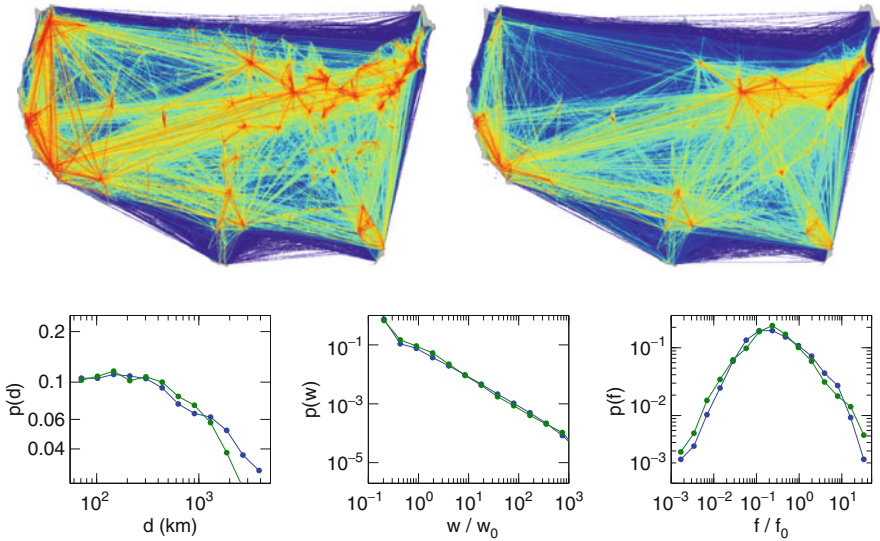
To create a model network comparable to our data, we first compute  $p_{ij}$  for all counties  $i$  and  $j$  in the continental U.S. and normalize them such that  $\sum_{i,j} p_{ij} = 1$ . We then interpret these values as probabilities for a travel event to happen between the two counties (or, speaking in terms of the original data source, a dollar bill report). Thus, starting with all-zero link weights  $w_{ij}$ , we repeatedly draw a pair of nodes according to  $p_{ij}$  and increase the corresponding  $w_{ij}$  by one, until approximately the same connectivity (number of non-zero  $w_{ij}$ ) as in the real-data network is reached.

We generated gravity networks for different parameter values and gauged them against our real data by comparing the distributions of first-order network statistics to find the best fit to our data. Distributions have been compared by log-binning the values and computing the  $\chi^2$  statistic

$$\chi^2 = \sum_i^n \frac{(N_i^G - N_i^R)^2}{N_i^R}$$

where  $n$  is the number of bins and  $N_i^G$  ( $N_i^R$ ) is the number of values from the gravity (real-data) network in bin  $i$ .

Our real data is symmetric and node fluxes are proportional to population sizes, therefore we assume  $\alpha = \beta \approx 1$  to narrow down the search volume in parameter space. We computed  $\chi^2$  for the distribution of link weights, node fluxes and geographical distances and used the sum of them,  $\chi_w^2 + \chi_f^2 + \chi_d^2$ , as the goodness-of-fit measure. Figure 7.7 shows this quantity for  $(\alpha, \mu) \in [0.8, 1.2] \times [-0.4, 0.9]$ , from which we concluded that  $\alpha = \beta = 0.96$  and  $\mu = 0.3$  are the best parameter choices. The resulting network and first-order statistics are shown in Fig. 7.8.



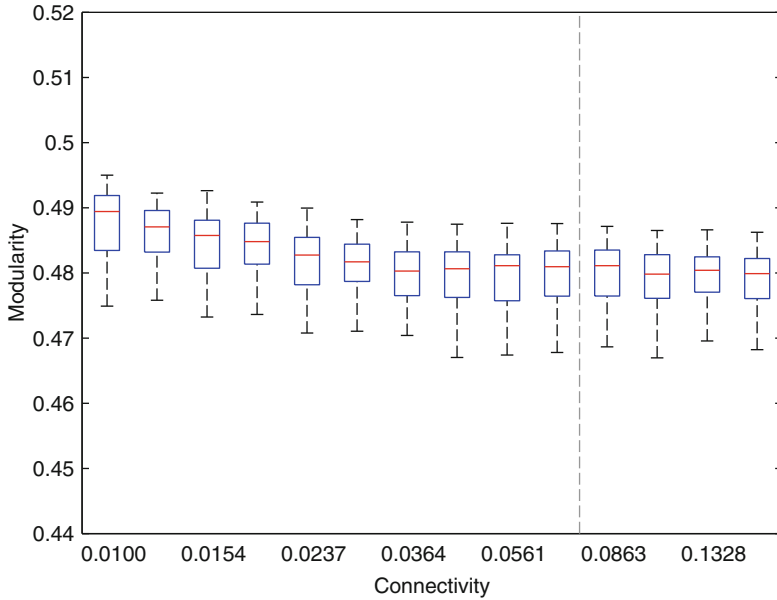
**Fig. 7.8** Comparison of the real-data network (*top left*) and the gravity model network with  $\alpha = \beta = 0.96$  and  $\mu = 0.3$  (*top right*). The bottom plot shows the distributions of geographical distances  $d$ , link weights  $w$ , and node fluxes  $f$  in the real-data network (*blue lines*) and the gravity model (*green lines*)

Similar to the bootstrapping procedure described in Sect. 7.7.1, we tested the robustness of the community structure of the model network by generating snapshots of the network at different connectivities and computing an ensemble of 80 high-modularity partitions for each snapshot. We found that the modularity statistics are stable around the target connectivity of 0.0765 (Fig. 7.9).

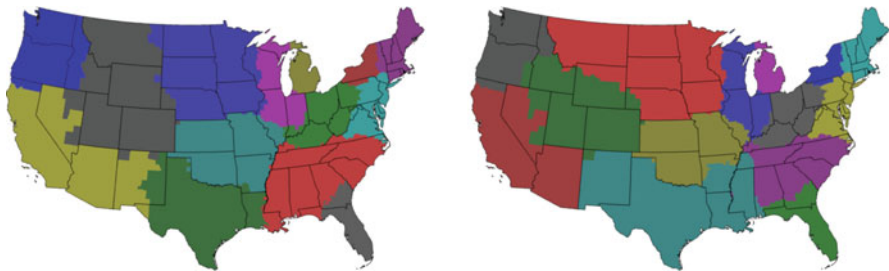
## 7.4 Degeneracy and Superposition

Given a mobility network constructed from the Where's George data, we then apply the optimization algorithm described in Sect. 7.2 to generate community partitions. Since the optimization process is stochastic, the resulting partition varies between realizations of the process. Two representative examples of high-modularity partitions are displayed in Fig. 7.10. Note that, although modularity only takes into account the structure of the weight matrix  $W$  and is explicitly blind to the geographic locations of nodes, the effective large-scale modules are spatially compact in every map. Consequently, although long-distance mobility plays an important role, the massive traffic along short distances generates spatial coherence of community patches of mean linear extension  $l = 633 \pm 250$  km. Note however that although each maps exhibits qualitative similarities between detected large-scale subdivisions and although each of the maps possess a high modularity score,





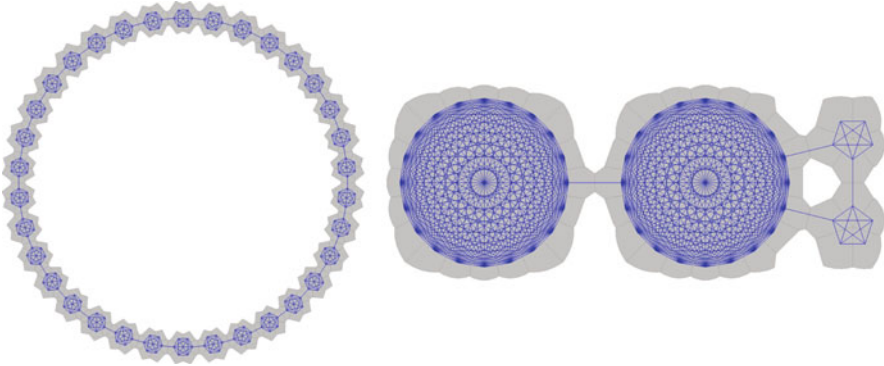
**Fig. 7.9** Distributions of modularity values for an ensemble of 80 partitions each computed for snapshots of the model network at different connectivities. The dashed line corresponds to 0.0765, the connectivity of the real-data mobility network



**Fig. 7.10** High-modularity community partitions of the WG mobility network. The stochastic algorithm produces different partitions when run many times; these are two representative examples. Modularity values are 0.6808 (*left*) and 0.6807 (*right*)

obvious structural differences exist; in fact, even if it were possible to determine a partition with maximal modularity, any such partition is not in principle unique. It is thus questionable whether any single effective map can be considered the most plausible partition.

Theoretical concerns aside, recent work [23] has identified practical issues with modularity maximization, in particular the so-called *resolution limit*. We demonstrate that a superposition of community partitions can alleviate these issues



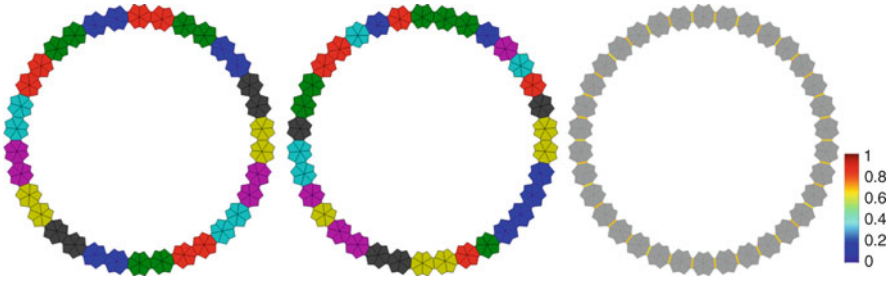
**Fig. 7.11** Two networks that expose the resolution limit problem with modularity. The shaded areas indicate an artificial geography for nicer visualization of the boundaries in the next figures. (Left) A ring of 34 cliques, each of 6 nodes and connected to their neighbors by single links. (Right) A network of two 20-node cliques and two five-node cliques

with the modularity score, and to this end discuss its known shortcomings in more detail.

In fact, it is straightforward to construct networks of which several distinct partitions with equal and maximum modularity value exist. This *degeneracy of modularity* was independently found by Good et al. [26] and marked as a drawback of the modularity measure.

Fortunato and Barthélemy [23] also report on the resolution limit of modularity. The authors present two artificial, unweighted networks that exhibit an intuitively very clear community structure, yet partitions exist that do not reflect this structure but have a higher modularity value than the partition that does. In particular, these networks are constructed by connecting multiple fully connected graphs (“cliques”) with single links (Fig. 7.11). It is clear that every clique should be grouped into one module, but the best partition according to modularity will group multiple cliques together. This only occurs if the cliques are small (in terms of number of links) compared to the full network, thus the modularity measure cannot detect communities below a certain resolution limit.

Our proposed method combines an ensemble of partitions by focusing on the boundaries of a partition (“Which adjacent nodes are separated into different modules?”) rather than its volumes (“Which nodes are grouped together?”) and then computing for each boundary the fraction of partitions in which it exists. Because we are interested in geographically embedded networks and modules are virtually always spatially compact in our case, we can restrict ourselves to boundaries that are also real geographical borders between nodes. However, the idea can be easily generalized to non-geographical networks, at the expense of convenient straightforward visualization. Since all partitions in the ensemble have a high modularity value, this method highlights similarities and differences in degenerated partitions,



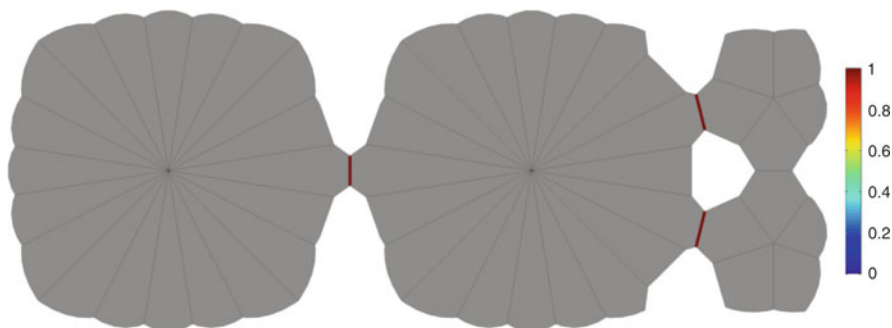
**Fig. 7.12** (Left) The optimal partition in the clique ring groups pairs of cliques together (the same color is used for multiple modules). (Center) Example of a partition found by the modularity optimization algorithm. (Right) Superposition reveals boundaries in the clique ring between every clique. Color codes the fraction of partitions in which the boundary was found. We use  $T = 2.5 \cdot 10^{-4}$ ,  $c = 0.75$ , and  $f = 0.5$  for this example and the next

yielding a unique “partition” (or to be more precise, a map) of the network and thus overcoming the degeneracy problem.

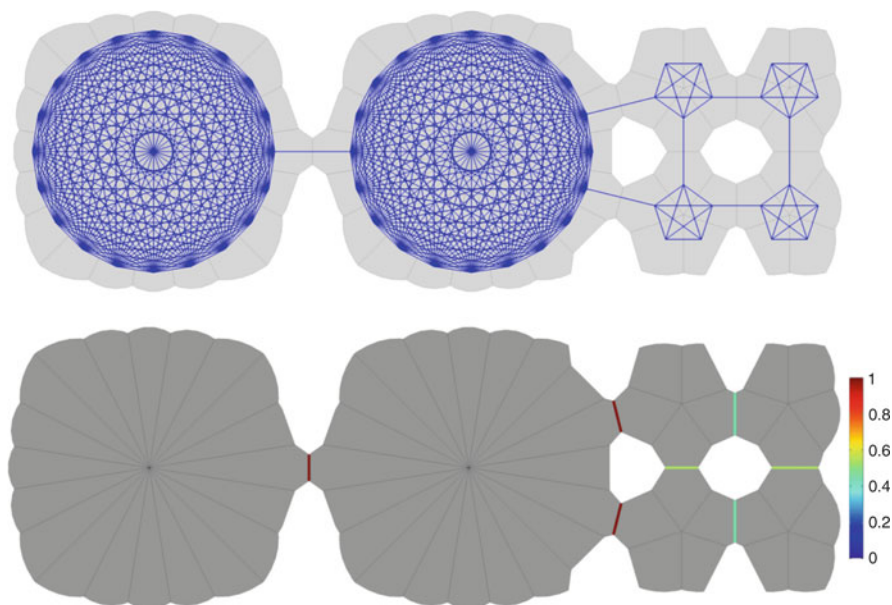
In our method, any single partition obviously suffers from this limitation as well. However, the resolution limit can be alleviated by looking at an ensemble, if enough small modules exist to create degeneracies. To illustrate this, we applied our method to the two example networks from Fortunato and Barthélemy [23]. Figure 7.11 (left) shows a ring of 34 6-cliques, all connected to their neighbors by a single link. The intuitive partition in which each clique is in its own module has modularity  $Q_{\text{real}} = 0.9081$  while a partition that groups pairs of cliques together has  $Q_{\text{opt}} = 0.9099$ . However, two distinct partitions exist that group pairs of cliques. Thus, an ensemble of optimal partitions will be composed out of those two partitions, yielding a boundary map in which every boundary between two cliques appears in 50% of the ensemble partitions. For nicer visualization, we created an artificial geography for this network and computed partitions and boundaries, shown in Fig. 7.12. Due to the nature of our algorithm, the resulting partitions contain a few  $n$ -tuples of cliques and single-clique modules that have not been split or merged into perfect clique-pairs before the termination criterion, and thus the observed boundaries are stronger than expected.

The second network proposed in Fortunato and Barthélemy is constructed from two 20-cliques and two 5-cliques (Fig. 7.11 (right)). Here, the two smaller cliques are merged into one module by the optimal partition ( $Q_{\text{opt}} = 0.5426$ ), although one would again expect each of them to be in its own module ( $Q_{\text{real}} = 0.5416$ ). Our method is not able to capture the intuitive community structure in this case (Fig. 7.13), because no degeneracy exists (the partitions in which only one of the small cliques are grouped with the large one, but not the other, are too far from the optimum to be produced by the algorithm,  $Q_{\text{deg}} = 0.4959$ ).

But if we extend the network such that four small cliques exist, the partition which groups all cliques into their own modules is still suboptimal to any partition that groups together more than one of the small cliques, but degeneracies are created



**Fig. 7.13** Boundaries found in the clique network shown in Fig. 7.11 (*right*). Our algorithm is not able to find a boundary between the two small cliques



**Fig. 7.14** Modification of the clique network in Fig. 7.11 (*right*). Because there are multiple high-modularity partitions that group the smaller cliques into pairs, our method can detect the correct community structure in this case

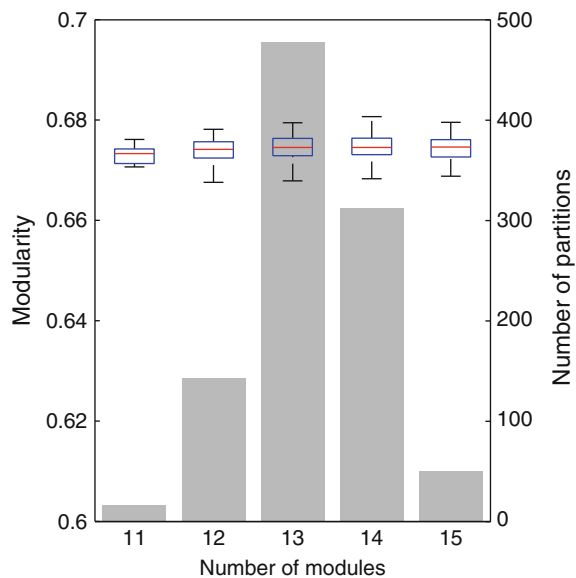
and the ensemble of partitions reveals the true community structure in this network (Fig. 7.14).

In conclusion, our method is able to dissolve both the degeneracy and resolution limit problems if enough small modules exist to create degeneracies. In fact, we will observe small “building blocks” in the WG data that are not seen in single partitions but emerge from the superposition of a partition ensemble.

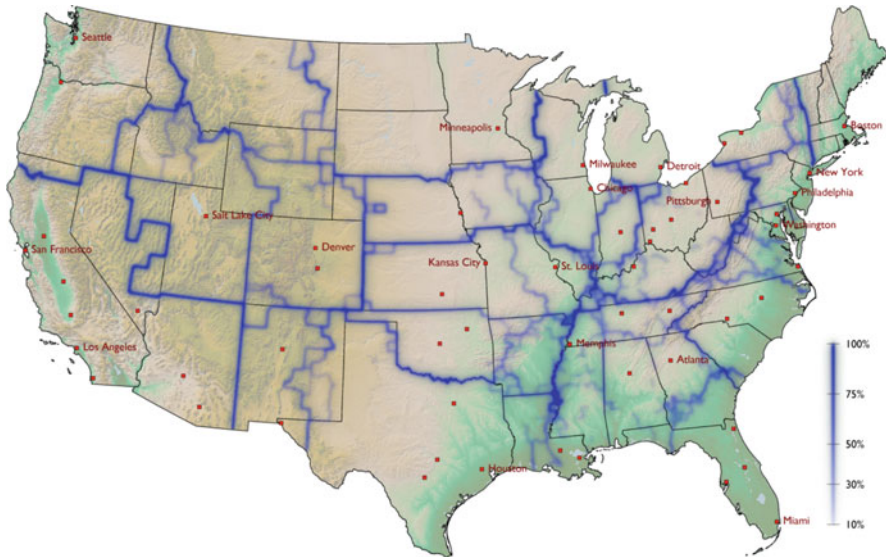
## 7.5 Assessment of Border Structures

Using the algorithm described in Sect. 7.2, we compute an ensemble of 1,000 partitions of the WG mobility network, all exhibiting a high modularity ( $Q = 0.6744 \pm 0.0026$ , see also Fig. 7.15 for the distribution of modularity values) and spatially compact modules, and perform a linear superposition of the set of maps. This method extracts features that are structural properties of the entire ensemble. The most prominent emergent feature is a complex network of spatially continuous geographic borders (Fig. 7.16). These borders are statistically significant topological features of the underlying multi-scale mobility network. An important aspect of this method is the ability to not only identify the location of these borders but also to quantify the frequency with which individual borders appear in the set of partitions, a measure for the strength of a border.

Investigating this system of effective mobility borders more closely, we see that although they correlate significantly with territorial state borders ( $p < 0.001$ , see Sect. 7.7) they frequently occur in unexpected locations. For example, they effectively split some states into independent patches, as with Pennsylvania, where the strongest border of the map separates the state into regions centered around Pittsburgh and Philadelphia. Other examples are Missouri, which is split into two halves, the eastern part dominated by St. Louis (also taking a piece of Illinois) and the western by Kansas City, and the southern part of Georgia, which is effectively allocated to Florida. Also of note are the Appalachian mountains. Representing a real topographical barrier to most means of transportation, this mountain range only partially coincides with state borders, but the effective mobility border is clearly



**Fig. 7.15** Ensemble statistics of geographic subdivisions for a set of  $N = 1,000$  partitions. The number of modules  $k$  in each subdivision is narrowly distributed around 13 (grey bars), and so are the conditional distributions of modularity (superimposed whisker plots). The ensemble mean is  $\bar{Q} = 0.674 \pm 0.0026$

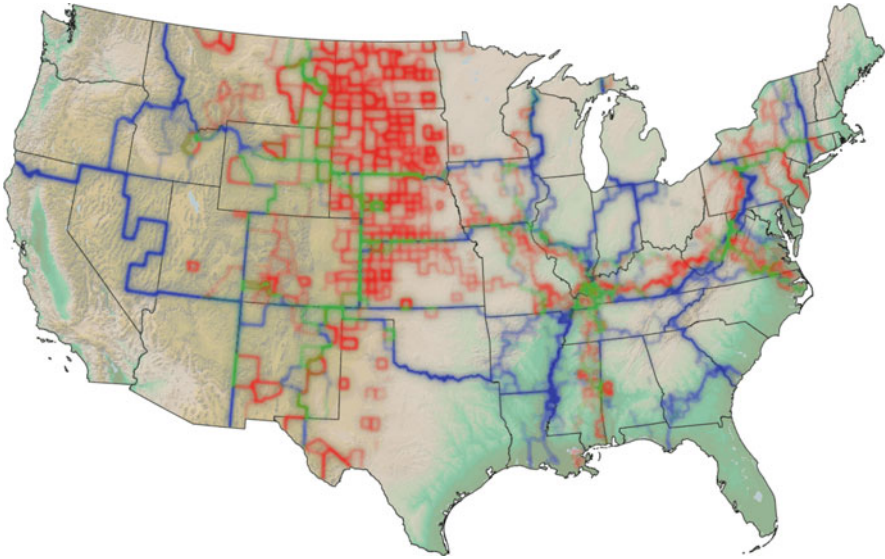


**Fig. 7.16** Effective borders emerge from linear superposition of all maps in the ensemble (*blue lines*). Intensity encodes border significance (i.e. the fraction of maps that exhibit the border). Black lines indicate state borders. Although 44% of state borders coincide with effective borders (*left pie chart*), approximately 64% of effective borders do not coincide with state borders. These borders are statistically significant features of the ensemble of high modularity maps, they partially correlate with administrative borders, topographical features, and frequently split states

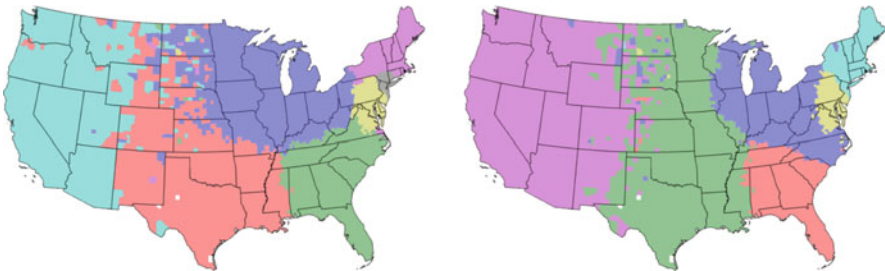
correlated with it. Finally, note that effective patches are often centered around large metropolitan areas that represent hubs in the transportation network, for instance Atlanta, Minneapolis and Salt Lake City. We find that 44% of the administrative state borders are also effective boundaries, while 64% of all effective boundaries do not coincide with state borders.

### 7.5.1 Comparison to Gravity Models

We also investigate whether the observed pattern of borders can be accounted for by the prominent class of gravity models [2,10,42], frequently encountered in modeling spatial disease dynamics [42]. In these phenomenological models, it is assumed that the interaction strength  $w_{ij}$  between a collection of sub-populations is given by (7.9), and we construct such a model according to the procedure described in Sect. 7.3.4. Although their validity is still a matter of debate, gravity models are commonly used if no direct data on mobility is available. The key feature of a gravity model is that  $w_{ij}$  is entirely determined by the spatial distribution of sub-populations. We, therefore, test whether the observed patterns of borders (Fig. 7.16) are indeed determined by the existing multi-scale mobility network or rather indirectly by the



**Fig. 7.17** The border structure of the gravity network (*red*) partially coincides with the borders in the original data (*blue*), but not significantly. The overlap is shown in green, for significance tests see Sect. 7.7



**Fig. 7.18** Sample partitions of the gravity network. Although they share qualitative features with those from the original network (Fig. 7.10), generic partitions of the gravity model network are structurally different, typically exhibiting fewer modules per partition, in different locations and with less spatial compactness

underlying spatial distribution of the population in combination with gravity law coupling. Figure 7.17 illustrates the borders we find in a network that obeys (7.9).

Comparing this model network to the original multi-scale network we see that their qualitative properties are similar, with strong short-range connections as well as prominent long-range links. However, maximal modularity maps typically contain only five subdivisions with a mean modularity of only  $\bar{Q} = 0.4791$ . Because borders determined for the model system are strongly fluctuating (Fig. 7.18), they yield much less coherent large-scale patches. Some specific borders, e.g. the Appalachian rim, are correctly reproduced in the model. The difference between the borders of the model system and the empirical data is statistically significant

(see Sect. 7.7), and we conclude that the sharp definition of borders in the original multi-scale mobility network and the pronounced spatial coherence of the building blocks are an intrinsic feature of the real multi-scale mobility network and cannot be generated by a gravity model that has a maximum first-order statistical overlap with the original mobility network.

## 7.6 Shortest-Path Tree Clustering

The methods already discussed successfully extract the structure of geographic borders inherent in multi-scale mobility networks. Bootstrapping the network indicates that these structures are surprisingly stable in response to perturbations of the network, but neither the modularity measure nor the stochastic algorithm we use to discover partitions provide specific information about the substructures in the network that make these borders so robust. What feature of the network, more specifically which subset of links if any, generates the observed borders? In order to address this question and further investigate the structural stability of the observed patterns, we developed a new and efficient computational technique based on the concept of shortest-path trees (SPT). Like stochastic modularity maximization, this technique identifies a structure of borders that encompass spatially coherent regions (Fig. 7.19), but unlike modularity this structure is unique. More importantly, it identifies a unique set of connections in the network, a network backbone, that correlates strongly with the observed borders.

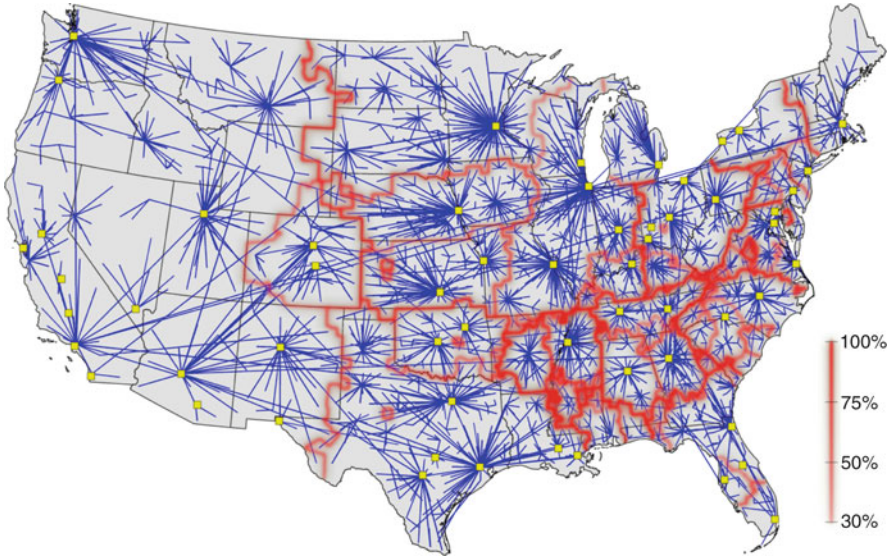
This second method for identifying community partitions, based on topological features of the analyzed network, has three parts. Given a network with  $N$  nodes containing a single connected component, we first compute a shortest-path tree for each node in the network. At least three widely-known algorithms are applicable (Dijkstra, Floyd–Warshall, and Bellman–Ford) and various optimizations are possible; in addition, if the input is sparse some of these algorithms improve in time complexity. In the worst case, however, this can be computed in  $O(N^3)$  time.

Second, we compute a dissimilarity score for each pair of shortest-path trees, and using the dissimilarity functions described below, this can also be accomplished in  $O(N^3)$  time.

Third and last, we apply hierarchical clustering to the table of dissimilarity scores, which also takes  $O(N^3)$  time for a naive implementation (because we compute the smallest element of an at-largest- $N$ -by- $N$  table  $N$  times). Therefore the entire suggested procedure takes  $O(N^3)$  time.

As mentioned, various optimizations are possible for computing shortest-path trees and hierarchical clustering, and these algorithms are so widely used that high-quality, efficient implementations are easily available. In fact, we find that the second step, computing dissimilarity scores, actually dominates the running time although it is by far the simplest computation; this is due to the fact that we use interfaces to pre-compiled, canned routines for steps 1 and 3, while step 2 is a naive MATLAB script. In practice the entire analysis can be run start to finish in under a half hour for our network of  $N = 3,109$  nodes on a circa-2008 laptop.





**Fig. 7.19** By comparing the border structure from SPT clustering with the ensemble of significant links (those that appear in at least half of the shortest-path trees) we identify topological structures which reveal the core of the network that explains the majority of border locations. This core is represented by the network in blue consisting of star-shaped modules centered around large cities (yellow squares)

### 7.6.1 Computing Shortest-Path Trees

The shortest path from vertex  $i$  to vertex  $j$  is the series of edges that minimizes the total effective distance  $d = \sum 1/w_{ij}$  along the legs of the path [14]. The distance along an edge for us is the inverse of the edge weight, as a highly-weighted edge indicates that two vertices are effectively proximal. (There are no edges with an infinite distance, because we do not define an edge between vertices if there is zero weight.)

The shortest-path tree  $T_i$  rooted at node  $i$  is the union of all shortest paths originating at  $i$  and ending at other nodes. We use the MATLAB interface<sup>2</sup> to the Boost Graph Library<sup>3</sup> to compute shortest-path trees. To prevent random fluctuations in our data from overwhelming the signal, we add a weak link between neighboring counties.

<sup>2</sup>[http://www.stanford.edu/~dgleich/programs/matlab\\_bgl/](http://www.stanford.edu/~dgleich/programs/matlab_bgl/).

<sup>3</sup>[http://www.boost.org/doc/libs/1\\_41\\_0/libs/graph/doc/index.html](http://www.boost.org/doc/libs/1_41_0/libs/graph/doc/index.html).

### 7.6.2 *Measuring Tree Distance*

A shortest-path tree can be easily represented as a vector of vertex labels  $T = [t_k], k = 1 \dots N$ , such that  $t_k$  is the label of the parent of vertex  $k$ , with a special symbol (perhaps 0) used to indicate the root. There are no disconnected nodes in the mobility network, thus each tree vector represents a single tree and not a forest. This representation lends itself to straightforward and meaningful comparisons between two trees.

We define two related measures of the dissimilarity between two trees. The first, called *parent dissimilarity*, asks the question, how many of the vertices in  $T_A$  do not have the same parent in  $T_B$ ? We denote this by  $z_p(T_A, T_B)$ , and it is exactly the general Hamming distance of two symbol sequences, that is, the number of places where corresponding labels in  $T_A$  and  $T_B$  do not match. The second, called *overlap dissimilarity*, asks the question, how many edges do the two trees *not* share? It is defined as  $z_o(T_A, T_B) = s_{\max} - s(T_A, T_B)$ . Here,  $s_{\max}$  is the largest number of edges two trees could share, which is the number of vertices less one (since the root does not contribute an edge).  $s(T_A, T_B)$  is the number of edges that  $T_A$  and  $T_B$  *do* share, and where  $z_p$  asks essentially the same question considering edges to be directed,  $z_o$  considers edges to be undirected. Also note that although we consider only the topology of trees when measuring their dissimilarity, the topology is determined by the weight of edges in the original graph and thus the mobility dynamics. For both measures, possible  $z$  values range from 0 (completely identical trees) to  $N$ , the number of nodes in the network.

We compute both measures for each distinct pair of trees in our network and find that they are highly correlated (the Pearson correlation coefficient of the two sets is 0.9980). For this reason, and because of the more straightforward interpretation, we focus exclusively on  $z_p$ . The parent dissimilarity values in our data range from 2 to 240.

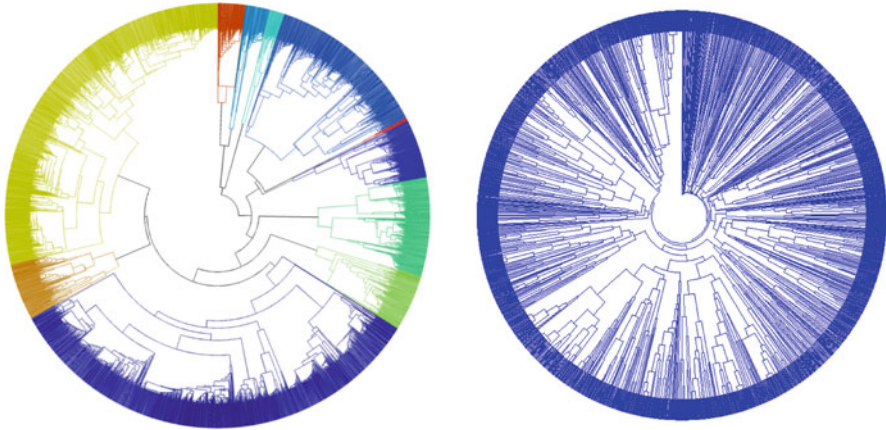
To test the stability of this measure we also added various amounts of noise to the original weight matrix; for example, adding 1% noise means that we adjusted each entry by a random number such that its perturbed value is within 1% of its original value. We then compute the set of shortest-path trees for the perturbed weight matrix, calculate the tree dissimilarities, and then compute the Pearson correlation of the original dissimilarities and the perturbed. The results (0.9995 for 0.1% noise, 0.9984 for 1% noise, 0.9937 for 5% noise) indicate the method is robust against small perturbations, and in addition we do not observe significant changes in the structure of borders determined by the perturbed matrices.

### 7.6.3 *Hierarchical Clustering and Borders*

The measures described above produce a dissimilarity matrix well-suited for use with hierarchical clustering [20]. This technique iteratively groups data points together into clusters that are less and less similar; it begins by identifying the

**Table 7.3** Cophenetic correlation coefficients [37] for various linkage functions using parent dissimilarity of trees and inverse weight of links

Linkage	$z_p(T_i, T_j)$	$1/w_{ij}$
Single	0.6584	0.1883
Average	0.8048	0.3757
Complete	0.7197	0.1400

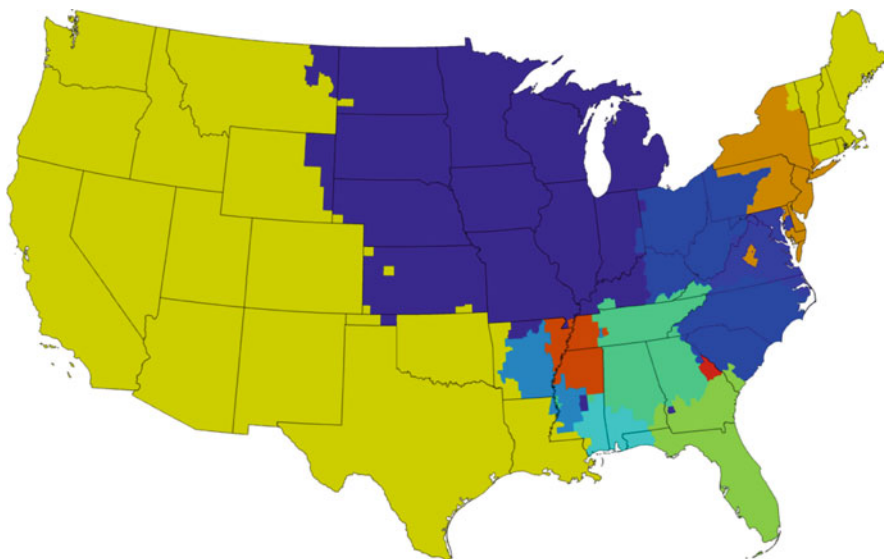


**Fig. 7.20** Dendrograms from hierarchical clustering. (Left) Using the parent dissimilarity matrix and average linkage. Colors correspond to a particular community partition depicted in Fig. 7.21. (Right) Using the inverse weight matrix with noise and average linkages. Even inspection by eye reveals immediately that clustering of the inverse weight matrix produces a poor fit to the data, pointing to the need for some type of “pre-conditioning,” here provided by SPT dissimilarity

two points with the lowest dissimilarity and grouping them together, then finding the next-most-similar data point or group, and so on. When it is necessary to compare the dissimilarity of one point (or group of points) with another group of points, a linkage function is used. There are several commonly-used linkage functions; we compute single linkages (comparing the shortest distance between two groups), average linkages (the average distance between two groups), and complete linkages (the greatest distance between two groups) and find that the average linkage produces the best fit to our data (Table 7.3).

The result of the hierarchical clustering algorithm is a linkage structure that can be represented graphically with a dendrogram (Fig. 7.20). The radial lines in the dendrogram represent vertices in our network or groups of vertices, and the arcs represent a link that joins groups together in the hierarchy. The nearer an arc is to the center of the circle, the greater the dissimilarity between the groups joined by the arc.

Each arc corresponds to a geographic border between a set of counties, and the closer the arc is to the center of the circle, the more significant the border. At the outermost level, the dendrogram necessarily puts a border around each individual

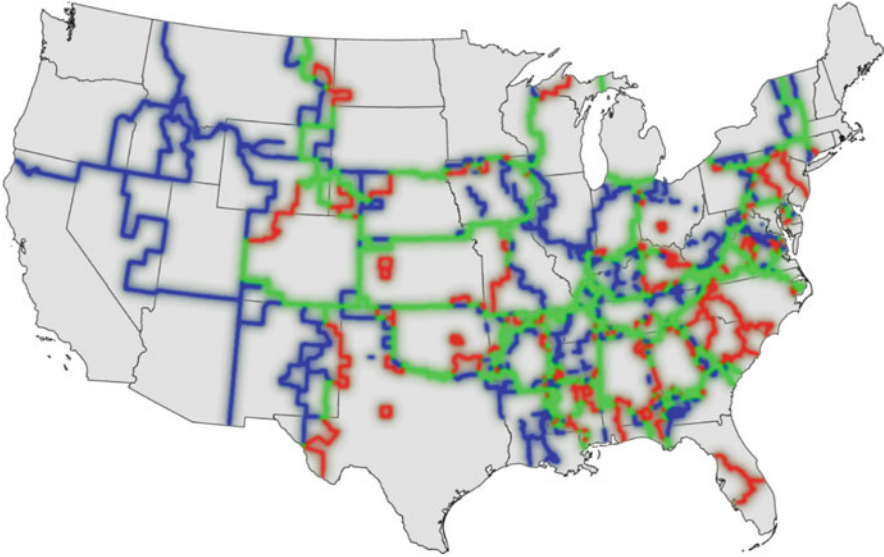


**Fig. 7.21** The geographic partition determined by cutting the dendrogram of Fig. 7.20 at a height of 95

county, and we threshold at 30% of the height of the tree (corresponding to a dissimilarity  $z_p = 41.6019$ ) for the analysis of the WG network.

As you can see in Fig. 7.21, the high-level groups identified by this procedure are spatially coherent, but may be divided into spatially disjoint regions at certain heights in the dendrogram.

Hierarchical clustering is also sometimes applied directly to the inverse of the weights,  $1/w_{ij}$ . We have investigated this method as well and find that it has several shortcomings. First, to apply a hierarchical clustering algorithm requires computing a dissimilarity for every pair of data points; since many pairs of counties are not directly connected by a link in our network ( $w_{ij}$  is zero), the inverse does not exist and it is consequently necessary to add some noise to the weight matrix at the very first step, representing pairs of vertices that are “extremely distant” but not disconnected. Second, the linkage structures produced from this approach fit the data poorly (Table 7.3). Last, one can see by visual inspection of the dendrograms in Fig. 7.20 that this approach does not yield significant information. Comparing the dendrograms for the  $z_p$  and  $1/w$  matrix, we see that in the shortest-path tree approach most of the links that appear higher in the tree (closer to the center) are linking together two groups that are strongly dissimilar from one another (seen by comparing the height of the parent link to the heights of the children links). In the inverse weight method, this is not true: links high in the tree are linking groups that are quite similar; that is, inverse weight clustering does not identify groups of strongly dissimilar vertices.

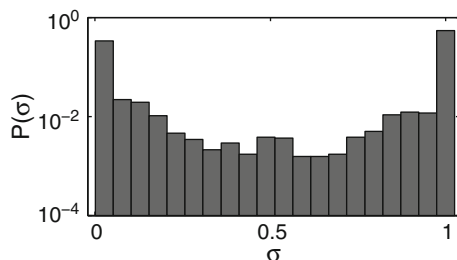


**Fig. 7.22** Comparing borders from modularity maximization (*blue*) with SPT clustering (*red*) reveals a significant overlap (*green*). The cumulative topological overlap (see Sect. 7.7) is 0.5282 indicating that the SPTD method represents an alternative computational approach to border extraction

Although the method yields a unique sequence of topological segmentations, the observed geographic borders exhibit a strong correlation with those determined by modularity maximization (Fig. 7.22).

#### 7.6.4 Link Significance

The key advantage of this method is that it can systematically extract properties of the network that match the observed borders. A way to demonstrate this is to measure the frequency  $\sigma$  at which individual links appear in the ensemble of all SPTs, which is conceptually related to their link betweenness [35]. Computing this *link significance*  $\sigma$  for each connection, we find that the distribution  $P(\sigma)$  of the network is bimodally peaked (Fig. 7.23). This is a promising feature of  $P(\sigma)$  as it allows labeling links as either significant or redundant without introducing an arbitrary threshold which is necessary for more continuously distributed link centrality measures. Extracting the group of significant links and constructing a subnetwork from these links only we observe that this subnetwork matches the computed border structure. By virtue of the fact that the most frequently shared links between SPTs are local, short-range connections we see that the SPT boundaries enclose local neighborhoods and that the boundaries fall along lines where SPTs



**Fig. 7.23** The distribution of link significance  $\sigma$ , defined for each link as the number of shortest-path trees the link appears in, exhibits a strong bimodal distribution. This implies that SPTD can sort links into important or not, and that  $\sigma$  is approximately a binary variable

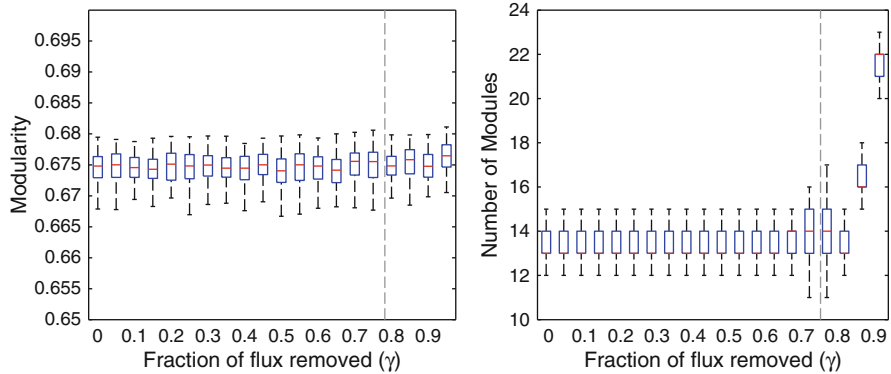
do not share common features. Note that effective metropolitan areas around cities can be detected with greater precision than modularity, although the western US is detected as effectively a single community.

Finally, we performed statistical analyses that quantify the overlap of the effective, mobility-induced borders with those provided by census-related systems. We choose the set of borders separating the states, the borders defined by the districts of the 12 Federal Reserve Banks, and the borders of Economic Areas [39]. We discuss this analysis in more detail in Sect. 7.7, but briefly, we find a significant correlation with economic boundaries ( $p < 0.001$ ,  $z$ -score 8.024 for the modularity borders and  $p < 0.001$ ,  $z$ -score 13.29 for the SPT borders).

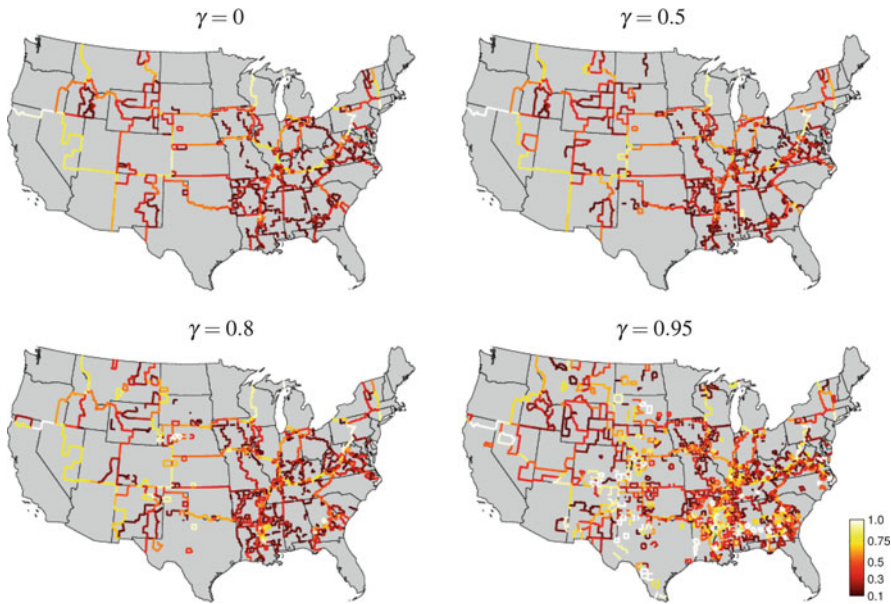
## 7.7 Significance and Comparison of Border Structures

### 7.7.1 Bootstrapping the Where's George Data

In order to test the robustness of our method against random data removal, we performed the following bootstrapping analysis. Starting with the full dollar bill dataset, and the resulting network weight matrix  $W$  with elements  $w_{ij}$ , we randomly remove single dollar bill reports until the total flux  $f = \sum_{i,j} w_{ij}$  is reduced by a factor  $\gamma$ . Using this method we constructed several networks for  $0 \leq \gamma \leq 0.95$  and computed an ensemble of 100 partitions for every value of  $\gamma$ , using the simulated annealing algorithm described in Sect. 7.2. We find that the modularity value is unaffected by bootstrapping even if 95% of the total flux is removed, although the number of modules in each partition rises as the network is thinned out more than 85% (Fig. 7.24). Also, the boundary structure emerging from superposition of all partitions is very robust under this procedure (Fig. 7.25). At 20% of the original flux ( $\gamma = 0.8$ ) virtually all of the boundaries found in the complete network are still identified, although the sparsity of the data evokes some singular counties. Even with only 5% of the flux, when boundaries become more fuzzy, some of the original structures are still detected.



**Fig. 7.24** Distributions of modularity values and number of modules for an ensemble of 100 partitions computed for each value of the bootstrapping parameter  $\gamma$ . The dashed line corresponds to 78.2%, the amount of flux ignored if all links shorter than 400 km would be removed



**Fig. 7.25** Linear superposition of 100 partitions for four different values of the bootstrapping parameter  $\gamma$ , color-coded according to the fraction of partitions they appear in

### 7.7.2 Measuring Overlap of Two Boundary Networks

In this section, we describe how to compare boundary networks defined on a planar graph, in our case the county network of the continental US excluding Alaska.

A *boundary network*  $b$  is simply given by assigning a nonnegative number  $w$  to each edge between adjacent counties: If the two counties are not divided by  $b$ , then

$w = 0$ . Otherwise,  $w > 0$  implies that the border shared between the two counties has the strength  $w$ . In Sect. 7.4, we described how to generate such a boundary network by superposition of many partitions of the Where’s George money travel network. We denote this boundary network by the modularity boundaries  $b_M$ .

We want to quantify how much information the modularity boundaries  $b_M$  shares with e.g. a state network, a random network, or a boundary network generated with another method.

For this we essentially need to determine the cross-correlation between two boundary networks  $b$  and  $b'$ . However, cross-correlation itself is not well-suited for dealing with the non-negativity of the edge weightings, so we calculate a non-centered version of it. The *absolute cross-correlation* of the two boundary networks  $b$  and  $b'$  is then given by the normalized scalar product of their edge weightings, i.e. by

$$a(b, b') := \frac{(1/|E|) \sum_{e \in E} b(e)b'(e)}{\sqrt{(1/|E|) \sum_{e \in E} b(e)^2} \sqrt{(1/|E|) \sum_{e \in E} b'(e)^2}},$$

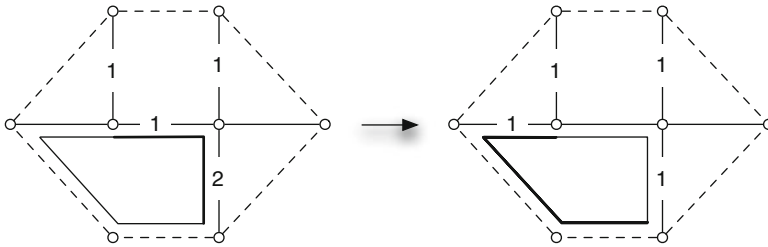
where  $E$  denotes the set of edges connecting adjacent counties. This quantity lies between 0 and 1 and equals 1 if and only if the two boundaries are identical up to scaling.

Apart from the upper bound, this quantity however is difficult to judge. In particular, we cannot compare right away two cross-correlations between different networks since  $a(\cdot, \cdot)$  might depend on the number of clusters and inhomogeneity of weights etc. We avoid finding a direct interpretation of the absolute cross-correlation by instead considering deviation of observed values against cross-correlations with a null model.

Such null models are used to tell random occurrences of structures from true information. One typically wants to keep some statistics of the network fixed while at the same time randomly sampling from its representational class. This results in the notion of random graphs with certain additional properties such as Erdős–Rényi [18] or Barabási–Albert [4]. The key idea is to generate a random network preserving planarity and possible additional information by using the original structure and iteratively changing it by a random local modification. For instance for unweighted networks, a random graph can be generated by “rewiring”: two distinct edges and two different vertices contained in either of the two are randomly selected and then swapped. Clearly this operation keeps both degree distributions fixed. After a certain number of iterations, the thus-generated Markov chain produces independent samples of the underlying random graph with given degree distributions [33]. This concept has been generalized to weighted graphs [43]; in this case, it is debatable whether to swap the whole weighted edge or to split up the weight.

In our case, we search for a randomization of a boundary network i.e. of a planar, weighted graph. Rewiring as above is not possible since it would destroy planarity. Instead, we propose to locally modify the graph at a random county: select a subpath





**Fig. 7.26** Local modification of a planar graph. We select the bottom left county to modify. The selected path to modify is shown in bold in the left figure. Its minimal weight is 1. This is subtracted in the right hand figure, where the complementary path is shown

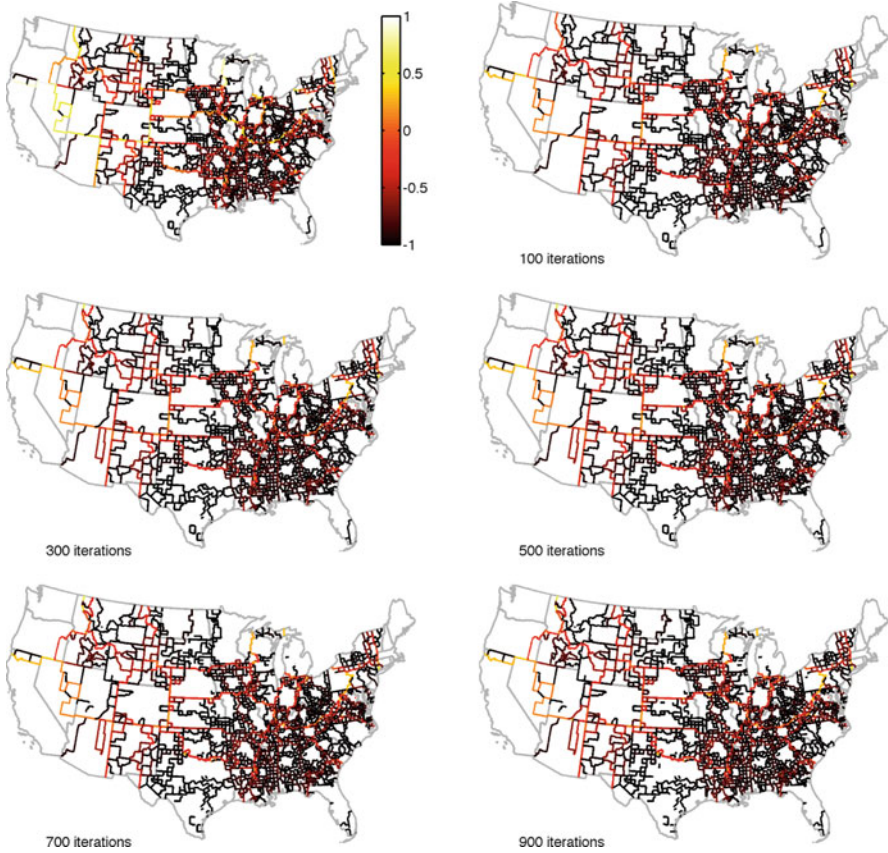
of its boundary and flip it to its complement. In the case of non-trivial weights, we reassign a random number between 0 and the minimal edge weight on the subpath. We have illustrated this procedure on an example in Fig. 7.26.

This procedure is now repeated multiple times until sufficiently de-correlated samples from the original network are produced. In practice, it is common to choose iterations in the range of the number of edges in the network or more.

### 7.7.3 *Randomization of the Mean Partition Boundary of the Where's George Network*

In order to test for significances of calculated similarities, we build a random model of the mean partition boundary by generating 1,000 random networks using the above algorithm with >15,000 successful iterations for each random network. The corresponding maps for the first 900 iterations are shown in Fig. 7.27. Clearly, the original structure in the boundary network is increasingly diluted, and after >10,000 iterations becomes stably random.

This can be seen by calculating the absolute cross-correlation  $a(b_M, b_R)$  of the modularity boundary network  $b_M$  with the random networks  $b_R$ , when increasing the number of iterations, see Fig. 7.28. We observe convergence to roughly 0.5 after about 10,000 steps. This lies well in the range of random correlation with a mean of 0.49 and a standard deviation of 0.028, see histogram in Fig. 7.29(a). This implies that the randomization procedure converges to a set of random boundary networks, which can now be used to put calculated autocorrelations into perspective against this null model.



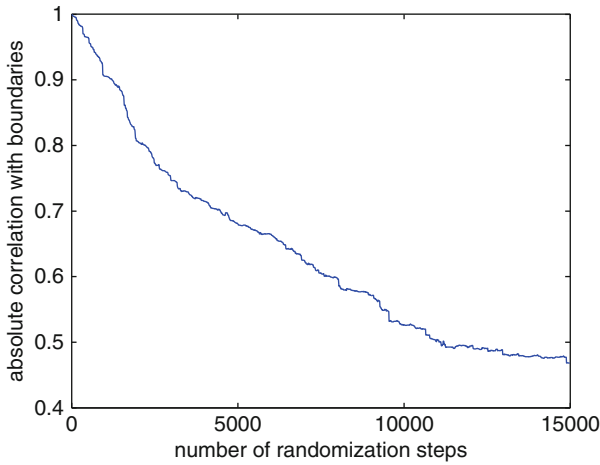
**Fig. 7.27** Randomization of the modularity boundary network  $b_M$ . The original network and the first 900 iterations are shown

#### 7.7.4 Significances when Comparing Boundary Networks with the Null Model

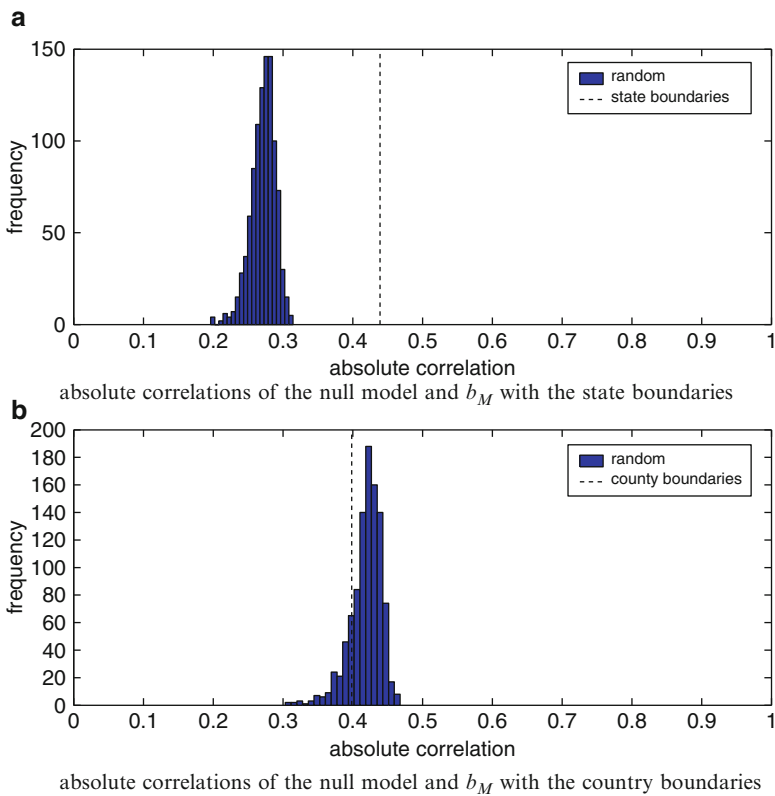
We describe and quantify overlap of the estimated modularity boundaries  $b_M$  with other political or social boundaries. As described before, we can quantify overlap by determining the absolute cross-correlation  $a(b, b_M)$ . In order to determine interpretable numbers, we compare this value to correlations with random boundaries  $b_R$  from a null model.

We now determine significance of coincidence of the modularity boundary network  $b_M$  and the SPT boundary network  $b_S$  with:

- Modularity boundaries  $b_M$
- State boundaries



**Fig. 7.28** Absolute cross-correlation of the randomized network after the given number of iterations with the modularity boundary network  $b_M$



**Fig. 7.29** Absolute cross-correlation of state and county boundaries when compared with a null model based on the modularity boundary network  $b_M$

**Table 7.4** Comparing boundary overlaps for various boundary networks with the modularity boundaries  $b_M$  and the corresponding null model  $b_R$  using absolute cross-correlation  $a$ 

Boundary network	$a(\cdot, b_M)$	$a(\cdot, b_R)$	$p$ -value	$z$ -score
Modularity boundaries	1.000	$0.495 \pm 0.028$	$<10^{-3}$	18.15
SPT communities	0.552	$0.385 \pm 0.024$	$<10^{-3}$	7.03
State boundaries	0.439	$0.272 \pm 0.018$	$<10^{-3}$	9.46
County boundaries	0.398	$0.419 \pm 0.023$	0.84	0.90
Gravity boundaries	0.260	$0.253 \pm 0.019$	0.35	0.40
Large-range network boundaries	0.198	$0.181 \pm 0.017$	0.14	1.02
Federal reserve district boundaries	0.377	$0.227 \pm 0.019$	$<10^{-3}$	7.91
Economic area boundaries	0.452	$0.307 \pm 0.018$	$<10^{-3}$	8.024

- County boundaries (to test for sensitivity of the method against number of communities)
- Boundaries resulting from the SPT algorithm  $b_S$
- Boundaries determined on the gravity model
- Boundaries determined on long-range distances only
- Federal reserve district boundaries (FRB)
- Economic area boundaries (<http://www.bea.gov>)

The significance is calculated by replacing  $b_M$  and  $b_S$ , respectively, by elements from the corresponding null model.

For illustration we show two histograms and actual values for state and county boundaries in Fig. 7.29. Clearly, the random cross-correlations are quite different, which means that we have to interpret the actual values of 0.439 and 0.398 differently as well. Indeed it turns out that the state value is far from the mean random cross-correlation  $0.272 \pm 0.018$ , whereas the county one is not ( $0.419 \pm 0.023$ ). Indeed, the empirical  $p$ -values, determined as the fraction of random correlations above the observed true one, is 0 in the former and 0.84 in the latter case.

In order to compare cases with large deviation from the distribution, we determine the  $z$ -score that is the distance of the absolute cross-correlation from the mean of the null model normalized by the standard deviation:

$$z(b) := \frac{a(b, b_M) - E(a(b, b_R))}{\text{std}(a(b, b_R))},$$

where  $E$  denotes mean and  $\text{std}$  standard deviation. In the state case, this  $z$ -score is very high, 9.46, which means that the observed correlation is more than 9 standard deviations away from the random mean. In contrast the county  $z$ -score is 0.90, which means that the observation is within one standard deviation and hence not significant.

We summarize the calculated cross-correlations in Tables 7.4 and 7.5 for  $b_M$  and  $b_S$ .

**Table 7.5** Comparing boundary overlaps for various boundary networks with the SPT-based boundary  $b_S$  and the corresponding null model  $b_R$  using absolute cross-correlation  $a$ 

Boundary network	$a(\cdot, b_S)$	$a(\cdot, b_R)$	$p$ -value	$z$ -score
Modularity boundaries	0.552	$0.251 \pm 0.013$	$<10^{-3}$	22.55
SPT communities	1.000	$0.367 \pm 0.0164$	$<10^{-3}$	40.63
State boundaries	0.358	$0.220 \pm 0.0138$	$<10^{-3}$	10.99
County boundaries	0.569	$0.562 \pm 0.016$	0.36	0.44
Gravity boundaries	0.305	$0.260 \pm 0.016$	0.002	2.73
Large-range network boundaries	0.257	$0.199 \pm 0.015$	$<10^{-3}$	3.94
Federal reserve district boundaries	0.307	$0.159 \pm 0.013$	$<10^{-3}$	11.79
Economic area boundaries	0.492	$0.318 \pm 0.013$	$<10^{-3}$	13.29

### 7.7.5 Discussion

For the state and the SPT boundaries we observe a strong deviation from the null model when comparing against the modularity boundaries. So we can conclude that both state boundaries and SPT boundaries are more similar to  $b_M$  than expected by chance with a  $p$ -value  $<10^{-3}$ .

This is not the case for the gravity model, the county boundaries and the long-range model. In these cases, the cross-correlation with  $b_M$  is not larger than with a random model ( $p$ -value  $\approx 0.44$ ,  $\approx 0.84$  and  $\approx 0.14$ ). This means that they do not significantly coincide with  $b_M$ .

The absolute cross-correlation of the FRB boundaries with  $b_M$  is  $a(b_F, b_M) = 0.38$ , which is significantly high when compared with the null model, which exhibits cross-correlations of only  $a(b_F, b_R) = 0.23 \pm 0.019$ . We observe a strong deviation from the null model and can therefore conclude that the FRB boundaries are more similar to  $b_M$  than expected by chance with a  $p$ -value  $<10^{-3}$ .

The corresponding  $z$ -score equals 7.91, which is lower than the one for states (9.46). This implies that the modularity boundaries' overlap with the states is larger than the one with the FRB boundaries.

We interpret the results on the FRB boundaries when compared with  $b_M$  as follows:

- The structure of  $b_M$  may be (partially) due to political structure i.e. result from  $b_S$  or due to additional money transport within FRB districts i.e. correlate with  $b_F$ . Since both  $b_S$  and  $b_F$  share strong similarities, in each of the two situations, we would see overlap with both boundaries, so we can only judge strength of overlap with respect to the other boundary.
- We quantified strength of overlap by deviation from the null model, and the corresponding  $z$ -score was more than 1.5 standard deviations higher for the state model. This stronger overlap of states with  $b_F$  therefore favors the first hypothesis i.e. the situation that political boundaries are a stronger factor for the pattern observed in  $b_M$ . In the case of dominance of the second hypothesis, we would instead expect to still see overlap with state boundaries, but less overlap than with the FRB ones.

## References

1. Altmann, E.G., Pierrehumbert, J.B., Motter, A.E.: Beyond word frequency: Bursts, lulls, and scaling in the temporal distributions of words. *PLoS ONE* **4**(11), e7678 (2009). DOI 10.1371/journal.pone.0007678. URL <http://dx.doi.org/10.1371%2Fjournal.pone.0007678>
2. Anderson, J.: A theoretical foundation for the gravity equation. *The American Economic Review* **69**(1), 106–116 (1979). URL <http://www.jstor.org/stable/1802501>
3. Balcan, D., Colizza, V., Gonçalves, B., Hu, H., Ramasco, J.J., Vespignani, A.: Multiscale mobility networks and the spatial spreading of infectious diseases. *P Natl Acad Sci USA* **106**(51), 21,484–21,489 (2009). DOI 10.1073/pnas.0906910106
4. Barabási, A., Albert, R.: Emergence of scaling in random networks. *Science* **286**, 509–512 (1999)
5. Barabási, A.L.: The origin of bursts and heavy tails in human dynamics. *Nature* **435**, 207 (2005)
6. Brandes, U., Delling, D., Gaertler, M., Goerke, R., Hofer, M., Nikoloski, Z., Wagner, D.: On modularity clustering. *IEEE T Knowl Data En* **20**(2), 172–188 (2008). DOI 10.1109/TKDE.2007.190689. URL <http://ieeexplore.ieee.org/search/wrapper.jsp?arnumber=4358966&tag=1>
7. Brockmann, D., Hufnagel, L., Geisel, T.: The scaling laws of human travel. *Nature* **439**(7075), 462–465 (2006). DOI 10.1038/nature04292
8. Brockmann, D., Theis, F.: Money circulation, trackable items, and the emergence of universal human mobility patterns. *IEEE Pervas Comput* **7**(4), 28–35 (2008)
9. Buzna, L., Peters, K., Ammoser, H., Kühnert, C., Helbing, D.: Efficient response to cascading disaster spreading. *Phys. Rev. E* **75**(5), 056,107–056,114 (2007). URL <http://pre.aps.org/abstract/PRE/v75/i5/e056107>
10. Cochrane, R.: A possible economic basis for the gravity model. *Journal of Transport Economics and Policy* **9**(1), 34–49 (1975). URL <http://www.jstor.org/stable/20052391>
11. Colizza, V., Barrat, A., Barthélemy, M., Valleron, A.J., Vespignani, A.: Modeling the worldwide spread of pandemic influenza: Baseline case and containment interventions. *PLoS Med* **4**(1), 95–110 (2007). DOI 10.1371/journal.pmed.0040013
12. Colizza, V., Barrat, A., Barthélemy, M., Vespignani, A.: The modeling of global epidemics: Stochastic dynamics and predictability. *B Math Biol* **68**(8), 1893–1921 (2006). DOI 10.1007/s11538-006-9077-9
13. Danon, L., Díaz-Guilera, A., Duch, J., Arenas, A.: Comparing community structure identification. *J Stat Mech-Theory E* p. P09008 (2005). DOI 10.1088/1742-5468/2005/09/P09008. URL <http://www.iop.org/EJ/abstract/1742-5468/2005/09/P09008/>
14. Dijkstra, E.: A note on two problems in connexion with graphs. *Numerische Mathematik* **1**(1), 269–271 (1959). URL <http://www.springerlink.com/index/UU8608UU27K7256.pdf>
15. Eagle, N., Macy, M., Claxton, R.: Network diversity and economic development. *Science* **328**(5981), 1029–1031 (2010). DOI 10.1126/science.1186605
16. Eagle, N., Pentland, A.S., Lazer, D.: Inferring friendship network structure by using mobile phone data. *P Natl Acad Sci USA* **106**(36), 15,274–15,278 (2009). DOI 10.1073/pnas.0900282106. URL <http://www.pnas.org/content/106/36/15274>
17. Eaton, J., Kortum, S.: Technology, geography, and trade. *Econometrica* **70**(5), 1741–1779 (2002)
18. Erdos, P., Rényi, A.: On random graphs. i. *Publicationes Mathematicae* **6**, 290–297 (1959)
19. Ernst, D., Kim, L.: Global production networks, knowledge diffusion, and local capability formation. *Research Policy* **31**(8–9), 1417–1429 (2002). URL <http://linkinghub.elsevier.com/retrieve/pii/S0048733302000720>
20. Everitt, B., Landau, S., Leese, M.: Cluster analysis p. 237 (2001). URL <http://books.google.com/books?id=htZzDGICnQYC&printsec=frontcover>
21. Ferguson, N., Cummings, D., Cauchemez, S., Fraser, C., Riley, S., Meeyai, A., Iamsirithaworn, S., Burke, D.: Strategies for containing an emerging influenza pandemic in southeast asia. *Nature* **437**(7056), 209–214 (2005). DOI 10.1038/nature04017

22. Fortunato, S.: Community detection in graphs. *Physics Reports* **486**(3-5), 75–174 (2010). DOI doi:10.1016/j.physrep.2009.11.002. URL <http://dx.doi.org/10.1016/j.physrep.2009.11.002>
23. Fortunato, S., Barthélemy, M.: Resolution limit in community detection. *P Natl Acad Sci USA* **104**(1), 36–41 (2007). DOI 10.1073/pnas.0605965104. URL <http://www.pnas.org/content/104/1/36>
24. Girvan, M., Newman, M.: Community structure in social and biological networks. *P Natl Acad Sci USA* **99**(12), 7821–7826 (2002). DOI 10.1073/pnas.1226539799. URL <http://www.google.com/search?client=safari&rls=en-us&q=Community+structure+in+social+and+biological+networks&ie=UTF-8&oe=UTF-8>
25. González, M.C., Hidalgo, C.A., Barabási, A.L.: Understanding individual human mobility patterns. *Nature* **453**(7196), 779–782 (2008). DOI 10.1038/nature06958
26. Good, B.H., de Montjoye, Y.A., Clauset, A.: The performance of modularity maximization in practical contexts. *arXiv Preprint Server* p. 0910.0165v1 (2009)
27. Guimerà, R., Amaral, L.A.N.: Cartography of complex networks: modules and universal roles. *Journal of Statistical Mechanics* **2**, P02,001 (2005). DOI 10.1088/1742-5468/2005/02/P02001
28. Guimerà, R., Mossa, S., Turtschi, A., Amaral, L.A.N.: The worldwide air transportation network: Anomalous centrality, community structure, and cities' global roles. *P Natl Acad Sci USA* **102**(22), 7794–7799 (2005). DOI 10.1073/pnas.0407994102
29. Keller, W.: International technology diffusion. *Journal of Economic Literature* **42**(3), 752–782 (2004). URL <http://www.atypon-link.com/AEAP/doi/abs/10.1257/0022051042177685>
30. Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A.L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., Jebara, T., King, G., Macy, M., Roy, D., Alsthyne, M.V.: *Computational social science* (2009). DOI 10.1126/science.1167742. URL <http://www.sciencemag.org/cgi/content/summary/323/5915/721>
31. Liben-Nowell, D., Novak, J., Kumar, R., Raghavan, P., Tomkins, A.: Geographic routing in social networks. *P Natl Acad Sci USA* **102**(33), 11,623–11,628 (2005). DOI 10.1073/pnas.0503018102
32. Lloyd, A.L., Jansen, V.A.A.: Spatiotemporal dynamics of epidemics: synchrony in metapopulation models. *Mathematical Biosciences* **188**(1-2), 1–16 (2004). DOI DOI:10.1016/j.mbs.2003.09.003. URL <http://www.sciencedirect.com/science/article/B6VHX-4B9K8TY-4/2/f79d91db6e0f3711169a9079e3cb9ca2>
33. Maslov, S., Sneppen, K.: Specificity and stability in topology of protein networks. *Science* **296**(5569), 910–3 (2002). DOI 10.1126/science.1065103
34. Newman, D.: The lines that continue to separate us: borders in our 'borderless' world. *Progress in Human Geography* **30**(2), 143–161 (2006). DOI 10.1191/0309132506ph599xx. URL <http://phg.sagepub.com/cgi/content/abstract/30/2/143>
35. Newman, M., Girvan, M.: Finding and evaluating community structure in networks. *Phys. Rev. E* **69**(2), 026,113 (2004). DOI 10.1103/PhysRevE.69.026113. URL <http://prola.aps.org/abstract/PRE/v69/i2/e026113>
36. Sales-Pardo, M., Guimerà, R., Moreira, A.A., Amaral, L.A.N.: Extracting the hierarchical organization of complex systems. *P Natl Acad Sci USA* **104**(39), 15,224–15,229 (2007). DOI 10.1073/pnas.0703740104
37. Sokal, R., Rohlf, F.: The comparison of dendrograms by objective methods. *Taxon* **11**(2), 33–40 (1962). URL <http://www.jstor.org/stable/1217208>
38. Song, C., Qu, Z., Blumm, N., Barabási, A.L.: Limits of predictability in human mobility. *Science* **327**(5968), 1018–1021 (2010). DOI 10.1126/science.1177170. URL <http://www.sciencemag.org/cgi/content/abstract/327/5968/1018>
39. United States Bureau of Economic Analysis: Economic areas. <http://www.bea.gov/regional/docs/econlist.cfm> (2010). URL <http://www.bea.gov/regional/docs/econlist.cfm>
40. Vespignani, A.: Predicting the behavior of techno-social systems. *Science* **325**(5939), 425–428 (2009). DOI 10.1126/science.1171990

41. Watts, D., Strogatz, S.: Collective dynamics of 'small-world' networks. *Nature* **393**(6684), 440–442 (1998)
42. Xia, Y., Bjornstad, O., Grenfell, B.: Measles metapopulation dynamics: A gravity model for epidemiological coupling and dynamics. *Am Nat* **164**(2), 267–281 (2004)
43. Zlatic, V., Bianconi, G., Díaz-Guilera, A., Garlaschelli, D., Rao, F., Caldarelli, G.: On the rich-club effect in dense and weighted networks. *European Physical Journal B* **67**, 271–275 (2009)