Incorporating Road Networks into Territory Design

Nitin Ahuja, Matthias Bender, Peter Sanders, Christian Schulz and Andreas Wagner [a,b,c,c,a]

[a] PTV AG, {nitin.ahuja, andreas.wagner}@ptvgroup.com
[b] FZI Research Centre for Information Technology, matthias.bender@fzi.de
[c] Karlsruhe Institute of Technology, {sanders, christian.schulz}@kit.edu

Abstract. Given a set of *basic areas*, the territory design problem asks to create a predefined number of territories, each containing at least one basic area, such that an objective function is optimized. Desired properties of territories often include a reasonable balance, compact form, contiguity and small average journey times which are usually encoded in the objective function or formulated as constraints. We address the territory design problem by developing graph theoretic models that also consider the underlying road network. The derived graph models enable us to tackle the territory design problem by modifying graph partitioning algorithms and mixed integer programming formulations so that the objective of the planning problem is taken into account. We test and compare the algorithms on several real world instances.

Keywords: Territory Design, Graph Partitioning, Road Networks

1 Practical Motivation

The starting point for this research was a real territory design problem at the PTV AG; design and plan service territories for companies that send agents to customers on a regular basis. In this case, the number of territories to be planned is predetermined by the number of service agents. Required properties of territories are manifold. For example, the obtained territories should be reasonably balanced with respect to an activity index of the customers, they should have "good" accessibility (non-accessible areas have to be accounted for), have to be contiguous and compact; and most importantly the expected journey time for visiting a typical set of customers within the territory should be small. Even small improvements in the objective save fuel and time that is needed to visit the customers. In turn, the costs of the service company are reduced and the usage of valuable resources is optimized. This optimization can also lead to the company serving even more customers using the same amount of employees or increasing the face time with its customers.

More formally, given a set of *basic areas* such as zip code areas or customer sites, the territory design problem asks to create a predefined number of larger territories such that an objective function encoding the desired properties is optimized and territory constraints are fulfilled. Apart from the introduced example, other important applications include political districting, electrical power restricting or planning territories for social facilities like hospitals or administrative units.

We organize the paper as follows. We begin in Section 2 by introducing basic concepts and by summarizing related work. The main parts of the paper are Section 3 and Section 4. The former introduces our rationale, graph models and the approach employing graph partitioning algorithms to perform territory design. The latter uses a mixed integer programming approach to tackle the problem. We report on carefully designed experiments performed on several real world instances in Section 5. Experiments indicate that our algorithms compute high quality territories that have small average journey times in particular. Finally, we conclude with Section 6.

2 Preliminaries

2.1 Basic concepts

Consider a set of *basic areas* $B = \{b_1, \ldots, b_n\}$ with an *activity index* $a : B \to \mathbb{R}_{\geq 0}$. We use the short notation a_i for $a(b_i)$ and extend the activity index to sets, i.e. , for a set $B' \subseteq B$, its activity $a(B') := \sum_{b_i \in B'} a(b_i)$. A territory T is defined as a subset of the set of basic areas. The activity index of a territory T is defined as $a(T) = \sum_{b_i \in T} a(b_i)$. The territory design problem is to assign each basic area to one of the k territories, where $k > 1$ is the number of desired territories given in advance. A solution to the territory design problem is a set of territories $\mathcal{T} = \{T_1, \ldots, T_k\}$ such that for each pair of territories, $T_i \cap T_j = \emptyset$ and $T_1 \cup \cdots \cup T_k = B$. A *balance constraint* on the activity index demands that $a(T) \leq (1+\epsilon)\lceil \frac{a(B)}{k} \rceil \; \forall T \in \mathcal{T}$ for some *imbalance tolerance* parameter $\epsilon \geq 0$. A solution is called *feasible* if it obeys the balance constraint and other constraints placed on the *contiguity* and *compactness* of territories. The objective is to minimize the sum of pair-wise travel distances within the territories. Note, for small values of ϵ a feasible solution may not exist.

Consider an undirected graph $G = (V, E, c, \omega)$ with edge weights $\omega : E \to \mathbb{R}_{>0}$ and node weights $c : V \to \mathbb{R}_{\geq 0}$. Again, we extend c and ω to sets, i.e. , for a set $V' \subseteq V$ its node weight $c(V') := \sum_{v \in V'} c(v)$ and for a set $E' \subseteq E$ its edge weight $\omega(E') := \sum_{e \in E'} \omega(e)$. The set $\Gamma(v) := \{u : \{v, u\} \in E\}$ denotes the neighbors of v. A maximal connected component is a maximal subgraph of G in which every pair of nodes are connected by a path. The *graph partitioning problem* is looking for k *blocks* of nodes $V_1, \ldots, V_k$ that partition V, i.e., $V_1 \cup \cdots \cup V_k = V$ and $V_i \cap V_j = \emptyset$ for $i \neq j$. Here, a *balancing constraint* demands that $\forall i \in \{1, \ldots, k\} : c(V_i) \leq L_{\max} := (1 + \epsilon)\lceil c(V)/k \rceil$ for some imbalance parameter $\epsilon \geq 0$. Often the objective of graph partitioning problems is to minimize the total *cut* $\sum_{i<j} w(E_{ij})$ where $E_{ij} := \{\{u, v\} \in E : u \in V_i, v \in V_j\}$. However, during the course of the paper we will modify this objective.

2.2 Related Work

There has been a *huge* amount of research on territory design so that we refer the reader to the surveys [4,8,16] for most of the material. For material related to graph partitioning, we refer the reader to the survey [2]. Here, we focus on issues closely related to our main contributions.

The basic formulation of territory design problems is due to Hess and Samuels [6] who used relaxations of integer linear programming models to tackle the problem. Ronen [10] presents a mixed integer formulation to minimize the total driving distance of salesmen. However, it is assumed that each trip to a customer requires a separate trip which includes that the salesmen returns home, so that the average journey time is not measured.

Using the same assumption, Zoltner and Sinha [15] develop a linear optimization model incorporating road networks for the distance computation to minimize the travel time. Additionally, the networks are used to represent connectivity between the territories which helped to compute contiguous territories with improved accessibility. Note that in both cases the objective also implicitly encodes compactness of the territories.

Forman and Yue [5] use a genetic algorithm to compute congressional territories. The basic idea is based on an encoding and on genetic operators that where originally used to solve the Traveling Salesman Problem. Desired properties of the territories are integrated into the fitness function of

the individuals. In this paper, we also use an evolutionary algorithm which instead encodes the individuals as partitions of a graph model.

Hess et al. [7] describe a mixed integer program based on the location-allocation approach. In each iteration basic areas are assigned to territory centers and afterwards the centers are updated. Compact territories are achieved by minimizing the sum of the distances of the basic areas to the territory centers. Since we also implement and extend this approach we go into more detail later.

Recently, Butsch et al. [3] used a recursive geometric bipartitioning approach to assign basic areas to territories. The approach recursively computes bipartitions of the basic areas using their coordinates. However, this approach may compute longish, non-compact territories and has problems incorporating geographic obstacles such as rivers or mountains. We compare our algorithm against this approach in Section 5.

2.3   KaHIP

Within this work, we use the open source multilevel graph partitioning framework KaHIP [11,13] (Karlsruhe High Quality Partitioning)[1]. More precisely, we modify the distributed evolutionary algorithm KaFFPaE contained therein to create partitions of our graph models which in turn yield solutions to the territory design problem. Hence, we shortly outline the main components of KaHIP.

Besides the evolutionary algorithm, KaHIP implements many different algorithms, for example flow-based methods and more-localized local searches within a multilevel framework called KaFFPa, as well as several coarse-grained parallel and sequential meta-heuristics. The algorithms in KaHIP have been able to improve the best known partitioning results in the Walshaw Benchmark [14] for many inputs using a short amount of time to create the partitions.

Evolutionary Algorithm Outline. We now roughly outline the general structure of KaFFPaE since one of our algorithms employs a modified version of the evolutionary algorithm. KaFFPaE starts with a population of individuals (in our case partitions of the graph) and evolves the population into different populations over several rounds. In each round, the evolutionary algorithm uses a selection rule based on the fitness of the individuals of the population to select good individuals and combines them to obtain improved

---

Algorithm 1 General Stucture of the Evolutionary Graph Partitioning Algorithm KaFFPaE.

create initial population P
while stopping criterion not fulfilled
    select parents p_1, p_2 from P (using fitness function)
    combine p_1 with p_2 to create offspring o (using a graph partitioner)
    mutate offspring o
    evict individual in population using o
return the fittest individual that occurred

---

[1] available at `http://algo2.iti.kit.edu/documents/kahip/`

3

offspring. More precisely, KaFFPaE uses a tournament selection rule to select individuals for combination. In KaFFPaE, the fitness function is set to the number of edges cut, however, we will modify the fitness function for the purpose of territory design. A combine operation then employs a graph partitioning framework to obtain an offspring having the good cuts of both input partitions. We refer the reader to [12] for more details. Our algorithm generates one offspring per generation. The general structure of the evolutionary algorithm is depicted in Algorithm 1.

The algorithm is parallelized by giving each processing element its own population so that combine and mutation operations can be performed independently. This is combined with a scalable communication protocol to exchange high quality solutions between the processing elements over time.

3 Territory Design by Graph Partitioning

The territory design problem and the graph partitioning problem are closely related. Our approach to territory design using graph partitioning consists of two steps. At first we construct a graph that corresponds to the territory design problem. Then, a customized algorithm partitions the constructed graph into k blocks. We develop a one-to-one mapping between basic areas and vertices in the constructed model. Hence, a partition of our model yields a solution of the territory design problem. Edges are defined by a *neighboring* relation. We begin this section by explaining how we construct the graph to be partitioned, then define the fitness function that is used in the evolutionary algorithm and outline how everything is put together.

3.1 Constructing the Graph

We now explain how we construct the graph that will be partitioned by the graph partitioning algorithm which in turn gives us a solution to the territory design problem. Every basic area b_i with activity a_i corresponds to a node v_i in the graph with weight $c_i := a_i$. An edge between a node v_i and v_j exists if and only if the basic areas b_i and b_j are *neighbors* as described below. In general, we set the weights of the edges in our model to one. However, as we will see later, the evolutionary graph partitioning algorithm that is used to partition the model takes the real distances into account to compute the value of the objective function of a solution.

Edges in our Model. The algorithm we use to compute the edges in our model is as follows. Note that we have multiple design goals for our model. First of all, two basic areas that are *close* should be connected by an edge in the model. Additionally, the model should not be too dense, e.g. if the maximum degree of the graph is bounded, the graph should be connected and edges that are too long should also not be contained in the model. Roughly speaking, we perform two iterations of Kruskal's algorithm [9] on a *complete graph* where every node v_i corresponds to a basic area b_i and the edge weights are the distances between the basic areas $\omega(v_i, v_j) = d(b_i, b_j)$. Recall, that Kruskal's algorithm scans the edges of the graph in increasing order of their weight to grow a forest and that it adds the edges joining two trees in the forest.

After the first run of Kruskal's algorithm, *every* edge that is in the minimum spanning tree (MST) computed by the algorithm is inserted into our model. For the second run of Kruskal's

algorithm, we remove the just computed MST edges from the complete graph. Now, while the algorithm scans the edges in increasing order, it adds the current edge to our model if the maximum node degree in the current state of the model does not exceed a user defined parameter $\gamma > 1$ and if the length of the edge is smaller or equal to $\beta \cdot \omega_{\text{avg}}$. Here, β is a user specified factor and ω_{avg} denotes the average edge weight of MST edges of the first iteration of the algorithm.

3.2 Fitness Function

Recall that we are looking for a partition of the just defined graph model into k blocks of nodes $P = \{V_1, \ldots, V_k\}$. Moreover, each subgraph induced by a block V_i should to be *contiguous*. In other words, if $n_{mcc}(V_i)$ is the number of maximal connected components in the subgraph induced by V_i, we want $n_{mcc}(V_i) = 1$ for each block V_i. We define the number of connected components of a partition P as $n_{con}(P) = \sum_{i=1}^{k} n_{mcc}(V_i)$. To guide the evolutionary algorithm towards contiguous blocks, we use a penalty approach. More precisely, we use the factor $(1 + \alpha(n_{con}(P) - k))$ with penalty parameter $\alpha > 0$ in our objective function to ensure that non-contiguous territories are penalized. We then set the objective/fitness function of the evolutionary algorithm that partitions our model to

$$f_{obj}(P) = (1 + \alpha(n_{con}(P) - k)) \sum_{i=1}^{k} \sum_{u,v \in V_i} d(u,v) \ ,$$

where $d(u,v)$ is the average time needed to traverse the shortest route in the original network between the basic areas corresponding to u and v respectively.

3.3 Overall Algorithm

To tackle the territory design problem, we use the distributed evolutionary graph partitioner KaFF-PaE to partition our model. We *modify* the fitness function of the algorithm to the objective function presented in Section 3.2. Note that the partitioning algorithms within KaFFPaE still optimize the number of cut edges. The amount of allowed imbalance is a constraint of the partition problem in KaFFPaE. Due to the local search algorithms in KaFFPaE no block of a partition will be empty. However, it may be possible that partitions created during the course of the evolutionary algorithm are not contiguous. Hence, whenever we create an individuum/partition, we try to make it contiguous. This is done by grouping neighboring connected components of the partition so that each block becomes contiguous. More precisely, excess connected components are assigned to the neighboring block with the least activity. Doing this may result in imbalanced partitions so that a rebalancing step is performed afterwards to ensure the balance constraint. We call the overall algorithm to tackle the territory design problem KaTeD (Karlsruhe Territory Design).

4 Territory Design by Mixed Integer Programming

Besides the graph partitioning approach, we use a location-allocation method for territory design. In this section, we first describe the location-allocation method introduced by Hess et al. [7] and then present the details of our modified and extended location-allocation approach.

4.1 The location-allocation approach by Hess et al.

The idea of solving territory design problems by means of a location-allocation approach goes back to Hess et al. [7]. They apply the approach to the design of legislative districts. In this application, territories are supposed to be compact, contiguous, and balanced in terms of population. Due to the complexity of the problem, Hess et al. [7] decompose it into two subproblems which are solved in an iterative manner:

- The **location** subproblem seeks to find (virtual) territory centers which are used to calculate the compactness measure. In the first iteration, the centers are a guess; in all subsequent iterations they are obtained by computing the center of gravity in each territory.
- In the **allocation** subproblem the basic areas are assigned to territory centers. To this purpose, Hess et al. [7] formulate an integer linear program. Let C denote the set of territory centers and $d_{b_i c_j}$ the distance between basic area $b_i \in B$ and center $c_j \in C$. Furthermore, define the following decision variables:

$$x_{b_i c_j} = \begin{cases} 1 & \text{if basic area } b_i \text{ is assigned to center } c_j \\ 0 & \text{otherwise} \end{cases}$$

Then, the model of Hess et al. [7] can be stated as follows:

$$\min \sum_{b_i \in B} \sum_{c_j \in C} d_{b_i c_j}^2 a_i x_{b_i c_j} \tag{1}$$

$$\sum_{c_j \in C} x_{b_i c_j} = 1 \qquad\qquad b_i \in B \tag{2}$$

$$(1 - \epsilon)\frac{a(B)}{k} \leq \sum_{b_i \in B} a_i x_{b_i c_j} \leq (1 + \epsilon)\frac{a(B)}{k} \qquad\qquad c_j \in C \tag{3}$$

$$x_{b_i c_j} \in \{0, 1\} \qquad\qquad b_i \in B, c_j \in C \tag{4}$$

The objective function (1) optimizes compactness which is measured as the sum of the squared distances between basic areas and associated territory centers, weighted by the basic areas' activity index. Constraints (2) in combination with the integrality conditions (4) ensure that each basic area is assigned to exactly one territory center. Balance is achieved by the constraints in (3) which guarantee that the activity index of a territory deviates by at most $\epsilon \cdot 100$ percent from the mean activity index. Instead of solving the integer program, Hess et al. [7] set ϵ to zero, solve the linear programming relaxation and then resolve all fractional assignments.

Hess et al. [7] perform location and allocation alternately until the solution converges. The algorithm repeats if multiple initial guesses for the territory centers are available.

4.2 Modifications and Extensions

We now outline our modifications and extensions of the approach by Hess et al. [7]. Roughly speaking, we use a well-known procedure to determine good initial centers, modify the balance constraint, solve the integer program directly and apply a multi-start procedure. We call our version of that approach KaLocAlloc (Karlsruhe Location Allocation).

Location Step. The quality of the solutions obtained by the location-allocation approach strongly depends on the selection of the initial centers, i.e. , the centers used in the first iteration of the algorithm. A good initial set of centers should be well-distributed across the region under study. To this end, we adopt the initialization procedure from the k-means++ algorithm [1]: Among all basic areas we pick the first center uniformly at random. Among all remaining basic areas the next center is picked with a probability which is proportional to the basic area's squared distance to the nearest center already chosen. This is repeated until k centers have been selected.

Allocation Step. For the allocation of basic areas to centers we use model (1) - (4) with one modification. In order to limit only the *maximum* activity index of the territories, we drop the lower bound in constraints (3):

$$\sum_{b_i \in B} a_i x_{b_i c_j} \leq (1 + \epsilon)\frac{a(B)}{k} \qquad \forall\, c_j \in C \tag{5}$$

Although we drop the lower bound, solutions cannot contain empty territories. This is because the centers are picked among the basic areas and, therefore, at least the basic areas corresponding to centers are assigned to the respective center due to a distance of 0. Another difference to Hess et al. [7] is that we do not solve the linear programming relaxation of the model, but the integer program.

Multi-start Procedure. As already mentioned, the selection of good initial centers is important to achieve solutions of high quality. Therefore, we apply a multi-start procedure in which the problem is solved multiple times with different initial centers at each start. The centers at each start are selected randomly according to the k-means++ scheme described above. After a user-defined number of starts the approach returns the best solution across all starts. Since in our case the objective is to minimize the sum of pairwise travel times between all basic areas of the same territory, the multi-start algorithm returns the best solution according to this criterion.

5 Experimental Evaluation

Methodology. All experiments were performed on real-world data provided by PTV Group. The test data comprises 15 instances whose sizes range from approximately 300 to 5,000 basic areas. The number of territories to be planned is given in the test data and varies from 3 to 46. Furthermore, road distances and travel times for the shortest path between each pair of basic areas were available and have been used for compactness evaluation. We also compare our algorithms against the recursive partitioning approach by Butsch et al. [3] (BKNS) which has been provided by the authors. However, we perform only one repetition since the algorithm is deterministic.

We evaluated our approaches on all 15 instances with a time limit of 300 seconds. The imbalance tolerance parameter ϵ was set to 0.05 in all experiments. From the practical experience of PTV Group the chosen values for the time limit and for the imbalance tolerance are very acceptable values for human planners.

Parametrization. Parameters which are specific to territory design by graph partitioning were set as follows: penalty parameter to increase connectedness $\alpha = 0.1$; edge length factor in our model $\beta = 5$; node degree bound in our model $\gamma = 20$. We repeated both approaches 40 times using different random seeds for initialization and the average was taken. Parameters which are specific to the location-allocation approach were set as follows: For the allocation step, the relative MIP optimality gap was set to 0.001. The maximum time spent on the allocation step was limited to 15 seconds, and the number of multi-starts was unrestricted.

System. All experiments have been done on an Intel Xeon CPU E3-1245 at 3,4GHz having 16 GB RAM and Microsoft Windows 7. The location-allocation approach was implemented in Java. We used Gurobi 5.6 to solve the integer linear program within the location-allocation approach. The number of threads was set to 4 for KaTeD and KaLocAlloc to ensure comparability.

5.1 Computational Results

We shortly summarize the main results and present detailed per instance results in Table 2. First of all, the balance constraint is satisfied in all cases for all approaches. In 8 of 15 instances, the average solution quality of KaTeD outperforms KaLocAlloc. Considering the other 7 instances, the value of the objective of the territories computed by KaLocAlloc is 1.6 percent smaller than the ones computed by the evolutionary approach. The recursive partitioning approach BKNS always yields worse results than both of our approaches, but is faster since we could perform only one repetition of the algorithm due to the fact that the algorithm is deterministic. On the largest instance, G02, the BKNS algorithm needed 203 seconds to compute a result. On average, BKNS yields 3.7 percent larger objectives than KaTeD and 4.3 percent larger objectives than KaLocAlloc. The largest improvement over BKNS is obtained on instance BL16 (Thüringen) and amounts to 10.3 percent compared to the result of KaTeD which computes the best objective on that instance. Figures 1 to 6 compare the visual result computed by the different algorithms on two exemplary instances, Thüringen and Baden-Würtemberg in Germany.

6 Conclusion and Future Work

In this paper we addressed the territory design problem by developing graph theoretic models that also consider the underlying road network. The derived graph models enabled us to tackle the territory design problem by reducing it to a graph partitioning problem. The resulting graph partitioning problem is then solved by using a modified evolutionary graph partitioning algorithm takes the objective function of the territory design problem into account. On the other hand we extended an existing mixed integer programming formulation. We tested and compared the algorithms on several real world instances.

Important future work includes the integration of the objective function of the territory design problem directly into a multi-level graph partitioning algorithm. In particular, it would be interesting to define local search algorithms for our objective. On the other hand, it would be good to have a graph partitioning algorithm that can ensure connectedness of blocks.

References

1. D. Arthur and S. Vassilvitskii. k-means++: The Advantages of Careful Seeding. In *Proc. of 18th ACM-SIAM Symposium on Discrete Algorithms*, pages 1027–1035. SIAM, 2007.
2. A. Buluç, H. Meyerhenke, I. Safro, P. Sanders, and C. Schulz. Recent Advances in Graph Partitioning. In *Algorithm Engineering – Selected Topics,* to app.*, ArXiv:1311.3144*, 2014.
3. A. Butsch, J Kalcsics, S. Nickel, and M. Schröder. Geometric Approaches to Districting Problems. Technical report, 2013. Working Paper.
4. J. C. Duque, R. Ramos, and J. Suriñach. Supervised Regionalization Methods: A Survey. *International Regional Science Review*, 30(3):195–220, 2007.
5. S. L. Forman and Y. Yue. Congressional Districting Using a TSP-based Genetic Algorithm. In *Proc. of the 2003 International Conference on Genetic and Evolutionary Computation: PartII*, GECCO'03, pages 2072–2083. Springer, 2003.
6. S. W. Hess and S. A. Samuels. Experiences with a Sales Districting Model: Criteria and Implementation. *Management Science*, 18(4-part-ii):P–41, 1971.
7. S. W. Hess, J. B. Weaver, H. J. Siegfeldt, J. N. Whelan, and P. A. Zitlau. Nonpartisan Political Redistricting by Computer. *Operations Research*, 13(6):998–1006, 1965.
8. J. Kalcsics, S. Nickel, and M. Schröder. Towards a Unified Territorial Design Approach—Applications, Algorithms and GIS Integration. *Top*, 13(1):1–56, 2005.
9. J. B. Kruskal. On the Shortest Spanning Subtree of a Graph and the Traveling Salesman Problem. *Proc. of the American Mathematical society*, 7(1):48–50, 1956.
10. D. Ronen. Sales Territory Alignment for Sparse Accounts. *Omega*, 11(5):501–505, 1983.
11. P. Sanders and C. Schulz. Engineering Multilevel Graph Partitioning Algorithms. In *Proc. of the 19th European Symposium on Algorithms*, volume 6942 of *LNCS*, pages 469–480. Springer, 2011.
12. P. Sanders and C. Schulz. Distributed Evolutionary Graph Partitioning. In *Proc. of the 12th Workshop on Algorithm Engineering and Experimentation (ALENEX'12)*, pages 16–29, 2012.
13. P. Sanders and C. Schulz. Think Locally, Act Globally: Highly Balanced Graph Partitioning. In *Proc. of the 12th International Symposium on Experimental Algorithms (SEA'13)*, LNCS. Springer, 2013.
14. A. J. Soper, C. Walshaw, and M. Cross. A Combined Evolutionary Search and Multilevel Optimisation Approach to Graph-Partitioning. *Journal of Global Optimization*, 29(2):225–241, 2004.
15. A.A. Zoltners and P. Sinha. Sales Territory Alignment: A Review and Model. *Management Science*, 29(11):1237–1256, 1983.
16. A.A. Zoltners and P. Sinha. Sales Territory Design: Thirty Years of Modeling and Implementation. *Marketing Science*, 24(3):313–331, 2005.

A Tables and Pictures

Table 1. Sum of Pairwise Travel Times – Average

Instance	#Basic Areas	#Territories	KaTeD avg	KaLocAlloc avg	BKNS avg
BL1	520	6	56 080 176	**54 157 925**	56 943 136
BL3	1 265	13	**143 300 619**	144 435 173	146 538 257
BL5	2 440	20	239 489 621	**238 965 327**	250 332 974
BL6	514	6	**109 625 833**	110 561 195	118 085 932
BL7	1 038	9	**48 151 371**	48 456 559	50 664 649
BL8	1 252	11	**159 771 590**	160 239 521	172 659 265
BL9	2 019	18	260 762 258	**258 027 628**	271 895 453
BL11	427	3	28 990 949	**28 796 550**	29 745 445
BL13	289	3	50 774 924	**50 326 699**	50 301 370
BL14	493	5	**55 345 856**	55 723 702	55 844 700
BL15	284	3	41 953 610	**40 997 882**	41 426 146
BL16	428	5	**44 396 659**	44 564 384	48 977 328
G01	4 472	41	469 430 937	**459 630 225**	479 442 432
G02	4 971	45	**597 515 028**	591 514 229	632 623 485
G03	2 241	20	**333 612 833**	329 314 024	343 170 613

Table 2. Sum of Pairwise Travel Times – Minimum and Maximum

| Instance	KaTeD min	KaLocAlloc min	KaTeD max	KaLocAlloc max	BKNS min/max
BL1	54 918 613	**53 663 743**	56 276 617	**54 819 237**	56 943 136
BL3	**142 464 954**	143 406 994	**143 688 813**	145 504 176	146 538 257
BL5	235 752 526	**233 110 996**	**244 426 840**	245 775 302	250 332 974
BL6	**109 140 977**	109 548 283	**109 893 677**	111 336 978	118 085 932
BL7	**47 118 442**	48 293 411	48 591 642	**48 569 098**	50 664 649
BL8	158 613 738	**158 257 253**	**160 899 255**	161 325 525	172 659 265
BL9	257 234 587	**253 678 835**	262 697 498	**261 954 880**	271 895 453
BL11	28 973 547	**28 796 550**	28 992 672	**28 796 550**	29 745 445
BL13	50 774 924	**50 217 885**	50 774 924	**50 446 097**	50 301 370
BL14	**54 102 563**	55 456 757	57 347 598	**55 810 459**	55 844 700
BL15	41 255 274	**40 987 865**	41 971 516	**41 011 058**	41 426 146
BL16	**44 083 530**	44 455 849	44 609 070	**44 631 114**	48 977 328
G01	459 200 690	**447 119 646**	473 914 701	**473 030 108**	479 442 432
G02	581 718 616	**575 445 420**	612 582 152	**606 800 497**	632 623 485
G03	329 746 198	**326 456 733**	336 518 355	**333 657 431**	343 170 613

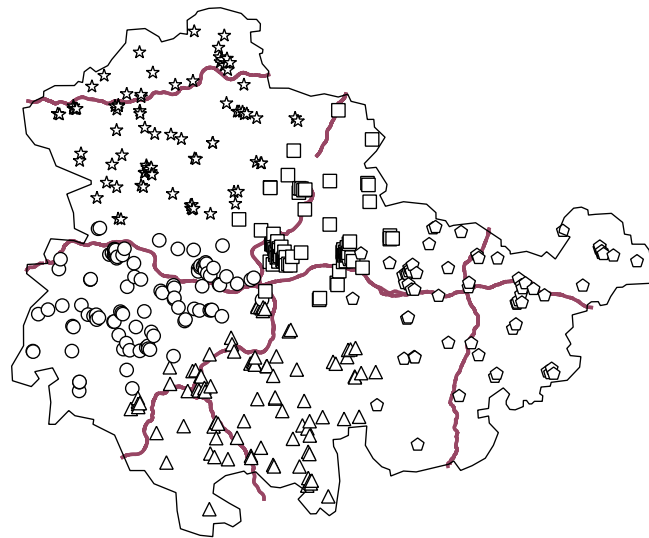**Fig. 1.** Visualization of the KaTeD result of instance BL16 (Thüringen, Germany).

**Fig. 2.** Visualization of the KaLocAlloc result of instance BL16 (Thüringen, Germany).

**Fig. 3.** Visualization of the BKNS result of instance BL16 (Thüringen, Germany).

Fig. 4. Visualization of the KaTeD result of instance BL8 (Baden-Wuerttemberg, Germany)

12

**Fig. 5.** Visualization of the KaLocAlloc result of instance BL8 (Baden-Wuerttemberg, Germany).

Fig. 6. Visualization of the BKNS result of instance BL8 (Baden-Wuerttemberg, Germany).