

IRIS: an Intelligent Tool Supporting Visual Exploration of Spatially Referenced Data

Gennady L. Andrienko, Natalia V. Andrienko

GMD - German National Research Centre for Information Technology
FIT.KI - Research Division Artificial Intelligence
Schloss Birlinghoven, Sankt-Augustin, D-53754 Germany
Tel +49-2241-142329 Fax +49-2241-142072
E-mail: gennady.andrienko@gmd.de
WWW: <http://allanon.gmd.de/and/and.html>

Abstract. To obtain a visual data presentation with a conventional geographic information system (GIS) or another graphic tool, the user should have special knowledge on graphics design for making correct selection of visualisation techniques in accord with the characteristics of data to be presented and relations between data components. Including such knowledge into a visualisation system allows to find correct solutions automatically and to save user's time and efforts for data exploration. A very important role in data exploration belongs to interactive manipulations with data and graphics such as querying, filtering, transformations of graphics, obtaining supplementary presentations. An attempt to provide extensive intelligent support to users in data investigation is made in the knowledge-based system IRIS. It automatically generates cartographic presentations of spatially referenced data and enables a wide range of interactive manipulations with maps and data. The peculiarities and restrictions of cartographic presentation are accounted for in the approach to visualisation design applied in the system. The recent implementation of the system runs in the World Wide Web and allows work of remote users from all over the world.

Keywords: visualisation, data exploration, dynamic data displays, geographical information systems (GIS), knowledge-based systems, Internet, World Wide Web (WWW), Java programming.

1. Introduction.

Working with conventional software that enables graphical presentation of data, the users should themselves care about the consistency of visualisation techniques selected for some data with characteristics of these data. However, this requires special knowledge on graphics design that cannot be regarded as a typical attribute of any user. In particular, the users of the systems dealing with spatially referenced data (geographic information systems, or GIS) are expected to know the principles of cartographic data presentation. In addition, the users should master rather complicated GIS operations. Even a user with sufficient competence can be annoyed by the necessity to interfere solving her/his problem with the process of designing data presentations and implementing the design by means of available operations.

Solving problems with the use of graphics can be facilitated by a system that incorporates knowledge on graphics design and on this basis automatically performs intelligent visualisation of available data. Research on the use of basic knowledge on data visualisation for automated graphics generation is carried on for at least a decade. A number of systems were created in this area (see the survey [Murray 94]). The system APT developed by Mackinlay [Mackinlay 86] is recognised as classical. Mackinlay offered the approach to data presentation design that was followed in all later research works. In this approach the system

partitions the data set to be presented into subsets with less numbers of domains (fields) so that it becomes possible to select some *visualisation primitive* for each subset. The visualisation primitives used in APT include seven *visual variables* considered by Bertin [Bertin 83]: *position*, *size*, *value* (the degree of darkness), *colour*, *texture*, *orientation*, and *shape*. Besides, more complex graphical encoding techniques, such as bar charts, pie charts, plots, graphs, are also included into the basic set of primitive graphical languages. The visualisation primitives selected for the partitions are combined with the use of *composition operators*. Composition is done by merging parts of two designs that encode the same information. A planner capable of backtracking and revision of commitments made earlier performs the design process.

Selection of primitives is done depending on characteristics of data sets: number of domains, cardinality (number of different values) of a domain, type of values in a domain (nominal, ordinal or quantitative), kind of dependency between domains (functional or one-to-many). It is governed by the basic principle of graphical presentation stated by Bertin: features of data components should be matched by properties of visual variables selected to represent them. For example, a quantitative data component can be represented only by a visual variable that can express quantities, that is, by *position* or *size*.

The Mackinlay's work showed the feasibility of automated graphics design on the basis of generic, domain-independent principles and rules of visualisation. The suggested approach was developed in further research on data visualisation. In the system VISTA [Senay 94] more visualisation techniques and composition operators are available that allow, in particular, to build 3D graphics. In the project SAGE [Roth 90] the inventory of data characteristics to be accounted for in graphics design is significantly extended. The design is sensitive to the user's information seeking goals expressed by primitive operations of the kind "accurate lookup of separate data values" or "comparison of values of two attributes".

Casner [Casner 91] tries to consider more complex information-processing tasks performed by the user. He assumes that a user's task can be specified as a sequence of predefined primitive *logical operators* such as a query for a property of an object or arithmetic operations over values. The Casner's system BOZ receives such a sequence and substitutes logical operators by primitive *perceptual operators* such as *search* for an object with a given graphical property or *lookup* for a specific property of a given graphical object. Then the system selects visualisation primitives that allow realising these perceptual operators.

An attempt to approach the problem of cartographic visualisation is described in [Zhan 95]. The authors built the expert system that consults a GIS user what visualisation technique should be selected for a given data field. The system does not plan presentations for several fields and does not perform visualisation. On the basis of this work and the approach of Mackinlay the system VIZARD was created that automates the presentation of spatially referenced data on maps [Jung 95]. This system designs visualisations on the basis of the techniques conventionally applied in cartography.

Though the research on automated design of graphical presentations has yielded rather impressive results, we should admit that these results refer mainly to the design algorithms and the factors that govern the design commitments. In all information presentation systems obtaining a graphic is seen as the final goal of the work. In contrast, in our work we try to regard the automated data visualisation as an integral part of a set of facilities designed to support **data exploration**. Other facilities required are tools to manipulate with data and presentations.

We recognise that complex problems usually cannot be solved by mere viewing static pictures; different transformations of the graphics and building auxiliary presentations are required. This conforms to Bertin's [Bertin 83] and Tufte's [Tufte 83] recommendations to design graphics for exploration so as to allow easy transformation. Bertin offers some techniques of building dynamic graphics with the use of paper. Computer screen as a medium for presenting graphics gives wide opportunities for interactive manipulations with graphics. *Interactivity* and *transformability* of graphic displays can make visual data exploration much more effective.

According to this approach, we have developed the system IRIS intended to support exploration of spatially referenced data by automated presentation of these data in the form of thematic maps and enabling interactive manipulations with the data displays.

2. IRIS from the user's perspective.

IRIS operated with data stored in the *table* format. A table is a collection of uniform records (rows) composed from fields; the latter form table columns. The data should refer to some *geographical objects* listed in one of the columns. The system requires certain metainformation about the data: notions associated with each column and their relationships. The table selected for exploration is shown to the user as may be seen in Fig.1.

Viewing the table, the user selects columns with data s/he would like to analyse visually. The columns are selected or deselected by direct manipulation: clicking at cells of the table caption. To receive visual presentation of the data, the user simply activates the function "Visualise" available in the menu of the table window. In response s/he receives one or more maps representing the data from the selected columns. The maps are built fully automatically. To generate an appropriate presentation, the system analyses types of the data in the columns and relations between the notions associated with the columns, and on this basis selects permissible visualisation techniques and combinations. All possible designs are realised; therefore the user typically receives several maps. Having several different presentations of the same information is beneficial from the data exploration perspective: each of them gives different opportunities for analysis. For example, pie charts allow seeing proportions while bar charts make differences more easily estimated.

When the user selects another set of columns from the table, s/he receives another family of maps. All currently open map windows are closed to avoid screen overcrowding, but the maps remain available: the system maintains a list of all generated maps, and the user can open any of them when needed. Therefore there is an opportunity to compare maps presenting different data subsets, possibly, selected from different tables.

All maps are scaleable and clickable. When the user points with the cursor to some geographical object, s/he receives the exact data values referring to this object. The data values appear in a separate window, as may be seen in Fig.1.

All maps are provided with legends. A legend is intended not only to give a key to correct map interpretation but also to display the summary statistics of data variation. This information may be valuable for the analysis. In Fig.2 one can see a sample map with its legend. Besides necessary comments about the data presented and the scheme of mapping of data values into colour intensity scale, the legend contains Tukey's "box and whiskers" plot [Tukey 77] that allows to judge about distribution of data values within the series from relative positions of minimum, maximum, quartiles, and median. When a map represents




Fig.1. An example screenshot demonstrates the work of IRIS.

several numeric data components (table columns), its legend shows variation of each of them. If the components are comparable, the corresponding “box and whiskers” plots are also shown together in a row to enable comparison of variations.

When IRIS visualises columns with numeric values together with a column with qualitative information, it assumes the qualitative column to determine grouping of geographical objects and provides summary statistics calculated for the groups as well as for the whole set of objects. For example, the map in Fig.3 presents qualitative information about dominant religions in the countries of Europe together with quantitative data, percentages of children and old people in population. Different dominant religions are shown by painting contours of countries: each religion is assigned a specific colour. The quantitative data are shown by bar charts. Attached to the map is a supplementary window that shows the averaged “portraits”,

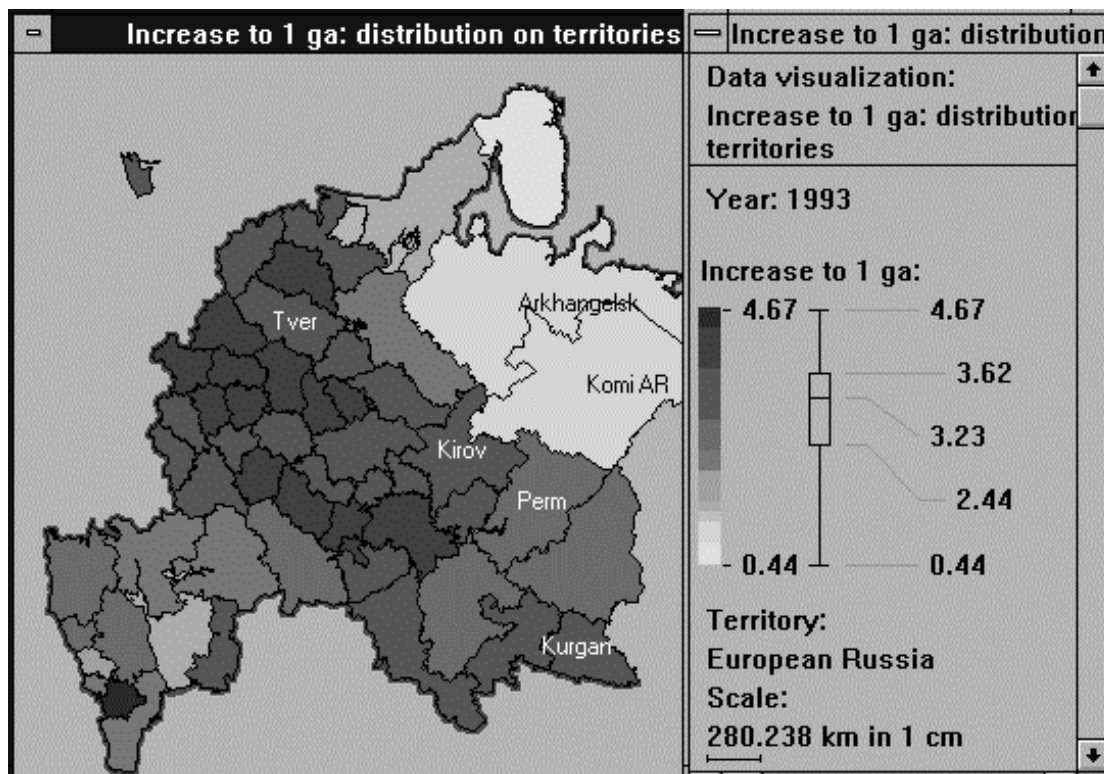


Fig.2. A map with presentation of numeric information by area shading.

with respect to percents of children and old people, of the groups formed by countries with the same religion. Differences between groups are noticeable. Below the averaged bar charts “box and whiskers” plots characterise variations of values for each numeric data component. Exact values of extremes, medians, quartiles are easily accessible: the user receives them pointing with the cursor to the rectangles representing groups.

Data exploration is supported not only by building graphical presentations but also by some interactive facilities for operating with data and maps. Data available in a table may be *filtered*: the user sets some restrictions on values in one or several columns (lower and/or upper limits for numeric columns and lists of selected values for qualitative columns), and only the rows that satisfy the restrictions remain in the table. The user may construct complex filters with the use of logical “and”, “or” and “not” operations. After setting a filter, all further visualisations are applied only to the data remaining in the table, i.e. shown in the maps are only objects satisfying the filter, and all statistics are counted for these objects. Reversing the filter by “not” operation, the user can obtain similar maps for the rest of the objects. Another opportunity is saving the filtering result in a new column of the table. The new column will have the *logical* type with the values ‘+’ (standing for *true*) designating filter satisfaction and ‘-’ (*false*) for the opposite; all table rows become visible again after filter saving. This logical column may be visualised, like any qualitative column, separately or together with other attributes. In the latter case the user will be also provided with summary statistics for the two object classes derived from the filtering.

Calculations can also help in data analysis. The user may construct any arithmetic formula from identifiers of numeric columns and numeric constants, and the system will make calculations by this formula for each table row. The obtained results will be placed in a new column. This derived information is further treated equally to the data initially present in the table: it can be visualised, used for filtering, in calculations, etc.

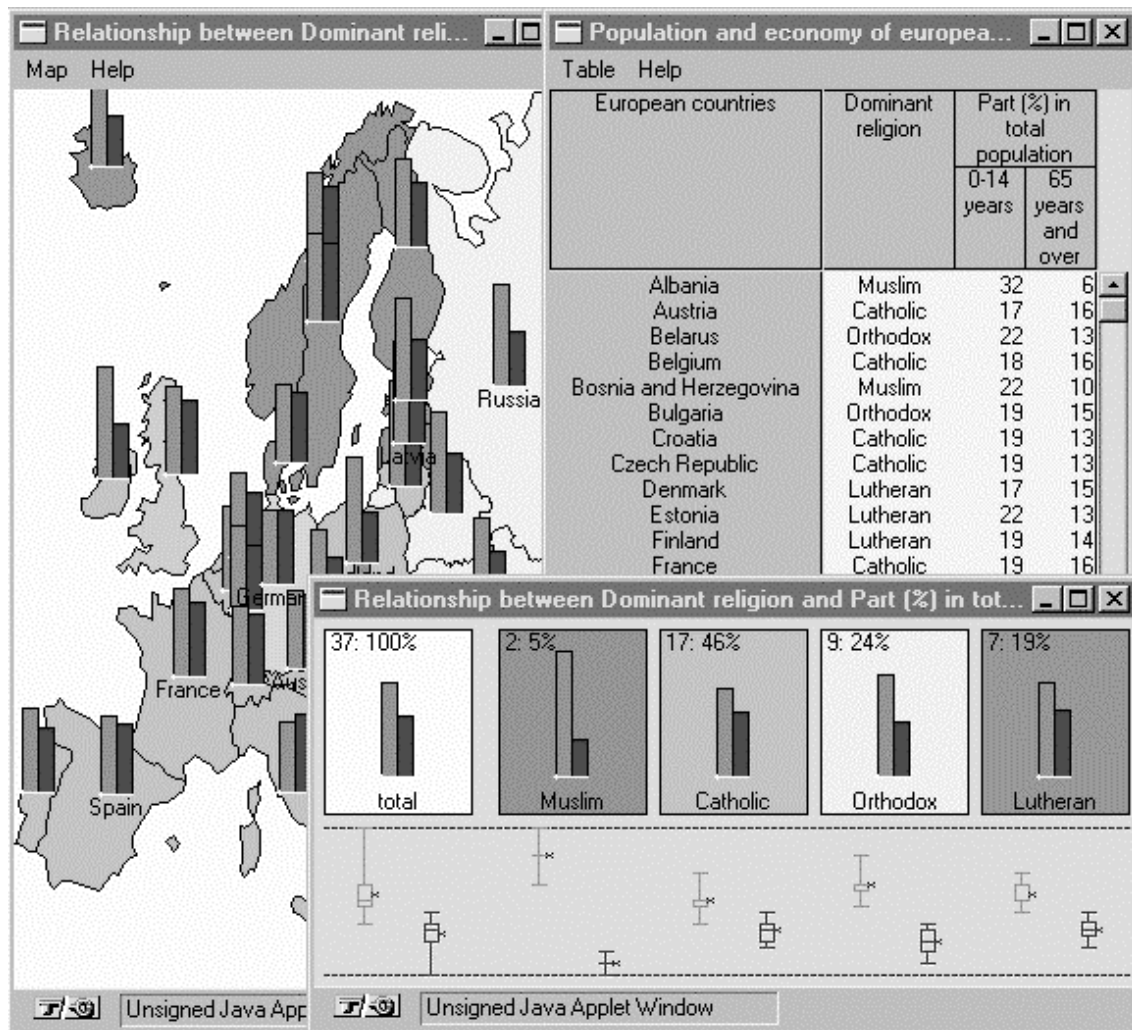


Fig.3. Averaged “portraits” of groups of geographical objects. Here European countries are grouped according to dominant religions. Both the map and the summary window show dominant religion by colouring and percentages of young and old people by bar charts.

Very powerful means of exploration are the facilities for *manipulating with maps* built by the system: the user can, to certain reasonable limits, vary the volume of the information present on a map and map appearance. These operations are sensitive to visualisation techniques used and therefore will be described later.

3. Visualisation in IRIS.

Implementing knowledge-based data mapping in IRIS, we were aware that the general principles of graphics design are certainly valid for maps but the latter have their own peculiarities. Geographical objects must be depicted on a map so as to reflect actual geographical positions, sizes and outlines of these objects. This requirement prescribes the way of organising other kinds of information in a map: graphical elements selected to represent data should be placed at the points of locations of the geographical objects the data refer to or inside their outlines. The visual variable *position* cannot be used to encode data. Besides, the graphical elements should be well visible against the background. The latter includes outlines of the geographical objects the data refer to and, possibly, other geographical layers such as rivers, forests, mountains and so on. This geographical

information provides familiar appearance of the territory and easy identification and location of the objects under analysis.

An arsenal of techniques appropriate for data visualisation on maps has been developed in cartography. Most of these techniques were included into IRIS. Among the visualisation techniques we distinguish those based on painting and those based on the use of signs. Signs may be either simple or structured. Painting can be applied when outlines of geographical objects are given to the system (the outlines are specified as polygons with co-ordinates of each vertex indicated). In this case we have *area* objects. Another variant is when geographical objects are specified by pairs of co-ordinates (*point* objects). Signs may be applied both to area and to point objects. For area objects, signs are placed near centres of their contours. Placement of signs for point objects is determined by their co-ordinates.

Each visualisation technique exploits some visual variable. Painting techniques may use *colour hue* or *colour value*¹ (further referred to as hue and value). Simple signs may vary in *size*, *shape*, *hue*, and *value*; combinations of these variables are limited to *size + hue*, *size + value*, *shape + hue*, *shape + value*. Size and shape are not combined because of the difficulties in comparing sizes when shapes differ. The variable *size* is used in two forms: *length* (bars) and *area* (circles). Area is better in case of large differences in data values. Structured signs (further called *diagrams*) are composed from graphical elements varying in *size* (length, area, angle), *hue*, or *value*; combinations of variables are not used. To ensure that graphical elements are distinguished within a diagram, they are given different colours (hues) or textures or/and fixed positions.

Consistency of properties of visual variables with data types is the main principle governing selection of visualisation techniques. Quantitative data can be shown by *size* or *value*, ordered qualitative data - by *value*, unordered qualitative data - by *shape* or *colour*. When several data components are to be visualised, presentation of each of them should be consistent with its type, and these presentations should be in some way combined.

The opportunities for technique combination in cartography are severely restricted: one can combine painting with signs but is not encouraged to place signs of different types on the same map. However, simple signs may combine two visual variables each assigned to a distinct data component, and structured signs themselves are means for combining several data components in one presentation. It is convenient for our purposes to classify diagrams depending on the methods of their construction from elementary graphical components: *juxtaposition*, *divergence*, *segmentation*, and *inclusion* (Fig.4). The basic rule of diagram type selection is the following: the construction method must be consistent with relationships among the components of data to be presented. Thus, juxtaposition (bar chart is an example) encourages comparisons between elements of a diagram and therefore may be chosen for *comparable* data components, such as numbers of births and deaths per 10000 population in countries of Europe. Bar charts give a good opportunity to estimate differences between birth and death rates. But this feature is undesirable when we wish to analyse together birth rates and sizes of national product per capita in dollars: differences of values of these two attributes have no sense. For *incomparable* data components diagrams based on divergence are used such as bars radiating from a point in different directions. In this case direct comparisons are obstructed but variations of the shapes of the resulted signs (whether there are typical patterns and, if so, whether they are nearly even or substantially asymmetric) allows to judge about

¹ *Texture* is applicable too, but the literature on graphics perception recommends to limit its use for it can produce the “vibratory effect” uncomfortable for eyes.

relatedness of the attributes.

Examples of segmentation are provided by pie charts and segmented bars (the latter may, in their turn, be juxtaposed or diverge). Segmentation is consistent with sets of data components that together make some wholes, for example, numbers of young, middle-aged and old people in country population. Diagrams based on inclusion are applied when there are inclusion relationships among data components, for example, number of population in total, number of students, and number of university students. Data related in such a way are shown in IRIS by nested squares.




Fig.4. Examples of techniques of diagram construction: *juxtaposition* (a), *divergence* (b), *segmentation* (c, d), *inclusion* (e).

So, in visualisation design IRIS significantly relies upon metainformation about relationships among data components. Limitations of the map form of data presentation, from the one side, and the availability of metainformation, from the other side, allowed implementing in the system a rather simple design algorithm. In the system all components selected for visualisation are grouped by comparability. For each group the system finds an appropriate visualisation technique. For groups with two or more *comparable components* the selection is determined by their *relationships*. In the case of one group nothing else is needed, for two groups the system applies permissible combination operators (area painting + signs or combining visual variables within simple signs), and for more than two groups the system tries to find a divergence-based diagram technique. Of course, presentation of arbitrarily selected data components within a single map is not always possible.

In our experiments with the earlier versions of IRIS we noticed that a powerful instrument to examine relationships between two or more data components is comparison of several maps each presenting one of the components. Area painting technique is especially good for these purposes. This way of data analysis was much easier for the users and gave larger number of interesting findings than that with presentation of the same information in one map by complicated diagrams. This observation inspired adding presentation by multiple maps to the arsenal of combination techniques used in IRIS. Each of the multiple maps presents one of the selected data components by area painting (for contour objects) or by simple signs (for point objects). These maps are united within a single panel, and the system supports their simultaneous scaling so that they always have the same size and show the same territory; this simplifies their joint analysis.

4. Exploration by interactive manipulation of maps.

In IRIS we combine automated visualisation with facilities for dynamic manipulation of maps. We strive to develop such interactive tools that strengthen the potential of a visual presentation, depending on the techniques used. Our approach consists in creating **specialised**

interactive manipulation facilities for each presentation technique. Such specialised interactive tools reinforce the original potential of visualisation in data exploration.

Interactive visual comparison with a number. This tool is designed for maps presenting a single numeric attribute by *value*, i.e. degree of darkness. It allows selection of some number N between minimum and maximum values of the shown attribute. In response the map is repainted so that values greater than N are depicted by shades of green and those less than N are shown by shades of cyan. The greater is the difference between some value and N, the darker is the shade used to represent it. The values exactly equal to N are shown in light yellow. The map is immediately repainted after any change of the basis for comparison. There are several ways to control N: entering an exact number, moving the slider, selecting an object in the map or in the list (the attribute value associated with this object becomes the reference value for the comparison), automatically locating the object with the previous or with the next value of the attribute.

The same tool is applied as well to maps encoding values of a numeric attribute by heights of bars. In this case bar heights are changed so as to be proportional to the difference between values of the attribute and the basis for comparison N. Positive differences are shown by green bars and negative by cyan ones.

Bringing to the common colour scale is designed for multi-map panels that depict comparable numeric data components. Initially each of the maps is generated independently, with its own scheme of correspondence of values and colour shades. However, the user may have these maps brought to the common colour scale: the system ensures that each shade encodes the same numeric value on all maps. For this purpose the system selects the shade for each value of each component depending to the position of this value between minimum and maximum found for all components instead of the “local” minimum and maximum for this component. The operation of colour scheme unification is reversible: on the user’s command the system redraws the maps in their initial state.

Switching on/off depicting of particular qualitative values helps in concentrating user’s attention on certain data values and examining their distribution without being interfered by other values. This operation is possible for maps presenting qualitative data by area colouring or by shapes of simple signs. It differs from *filtering* described earlier in that it is applied to a map instead of a table. The system does not redesign data presentation. It simply dynamically hides some values on a previously built map and shows them again in response to user’s manipulations with switches.

So, according to our approach to the design of dynamic manipulation tools for data exploration with maps, interactive tools are based on capabilities of each presentation technique and enhance these capabilities. The construction of interactive facilities for other presentation techniques applied in map generation is currently under way.

5. An example of data exploration with IRIS.

Cited below is a model data exploration session made with IRIS. The example can help to understand the work of different facilities offered by the system. During the session a hypothetical analyst considered a table with real data on the population structure and the stock

of living houses in 63 districts of Bonn City, Germany². Among others, the table contained the columns “Number of living houses in 1995” and “Number of living houses in 1996”. The analyst was interested how the number of houses changed in 1996 compared to 1995. For this purpose she made the system calculate the difference between these two columns and store the results in a new column. The system immediately offered two different map presentations of the new column. One of them was the choropleth map shown in Fig.5.

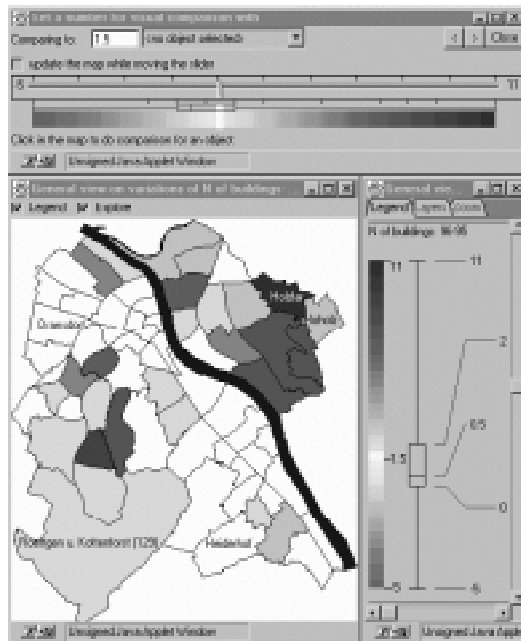


Fig.5. Distribution of values of the calculated attribute “Difference in number of buildings 1996-1995”.

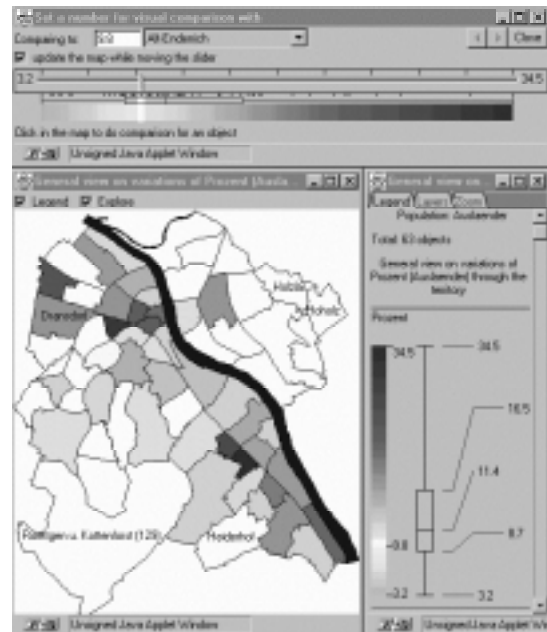


Fig.7. Dynamic comparison makes the similarity to Fig.5 clearer.

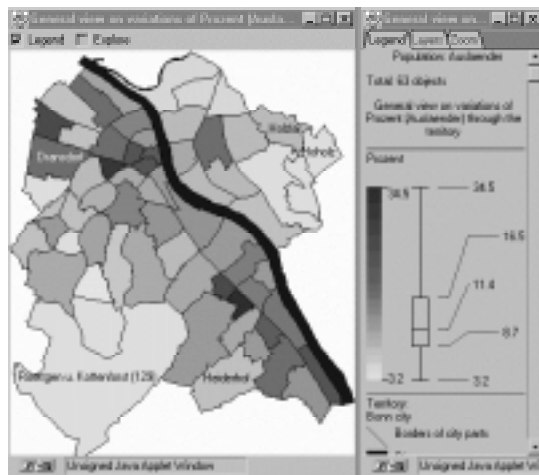


Fig.6. Distribution of percentage of foreigners.

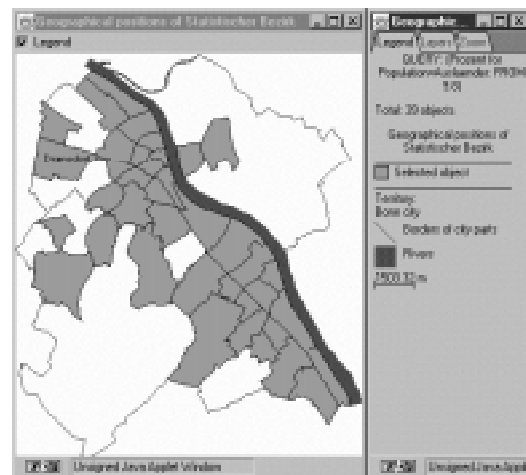


Fig.8. The districts with high percentage of foreigners: the map presents the results of filtering.

This map demonstrates an interesting trend in spatial distribution of values of the new attribute: there is a band stretching from north-west to south-east along the river on its left

² The system and the data described are available at the URL <http://allanon.gmd.de/and/java/iris>. The reader is kindly invited to reproduce this example in colour and dynamics as well as to try other system’s functions and other applications.

side where the number of houses did not increase or even decreased. The trend becomes more clearly visible when the dynamic comparison is applied³.

It was interesting to see which of the other available attributes have a similar distribution trend. The analyst selected different columns and looked at their visual presentation until she noticed a dark band with spatial orientation from north-west to south-east on a map representing the attribute “Percentage of foreigners” (Fig.6). Application of the dynamic comparison exposed a green band along the left bank of the river (Fig.7 shows it as grey whereas cyan was changed to white, as in Fig.5). It is similar to the cyan band in the map showing the difference in the number of houses. So, it seems that the two attributes inversely correlate: new houses are built mostly in the districts with low percentage of foreigners, and in the districts with many foreigners the number of buildings remains the same or decreases.

To check the hypothesis about the correlation between these two attributes, the analyst decides to divide the districts into two groups: districts with high percentage of foreigners and those with low percentage. For this purpose she used the filtering function of the system: she set the restriction “Number of foreigners > 9.8” (this number was found in the course of dynamic comparison), and the system selected and showed on a map the objects satisfying this condition (Fig.8).

The results of filtering were “saved”: on the analyst’s request the system produced a new attribute with logical values showing whether or not a district fits the condition. Then the analyst selected this new table column and the column with differences in number of houses and let the system visualise them together. The system proposed three visualisations. In Fig.9 one can see how the system offers new visualisations to the user. Icons denote what visualisation techniques are exploited in the maps, and map names give hints as to which analytical tasks the maps is suitable for.

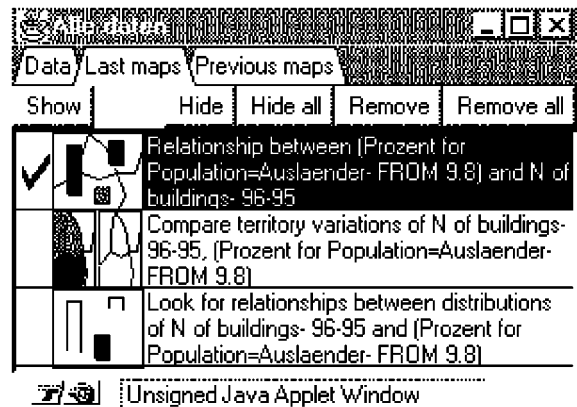


Fig.9. Several different presentations of the same data are differentiated by map names and icons.

From the proposed presentation the analyst selected the combination of colour painting showing condition satisfaction and bars encoding changes in number of houses. Simultaneously with displaying the selected map the system counted and presented the statistics of variation of the numeric attribute for the groups of districts formed according to the values of the logical attribute. The clearly seen differences in variations demonstrated by “box and whiskers” plots support the hypothesis about inverse correlation between the two attributes (Fig.10). This finding was unexpected and non-trivial. It was allowed by the ease of obtaining maps obeying the principles of graphical presentation and by the variety of tools for manipulating data and visual presentations provided by IRIS.

³ To enable seeing the result of dynamic comparison in a grey-scale picture, we had to repaint all cyan-coloured map objects in white in the initial colour screenshot.

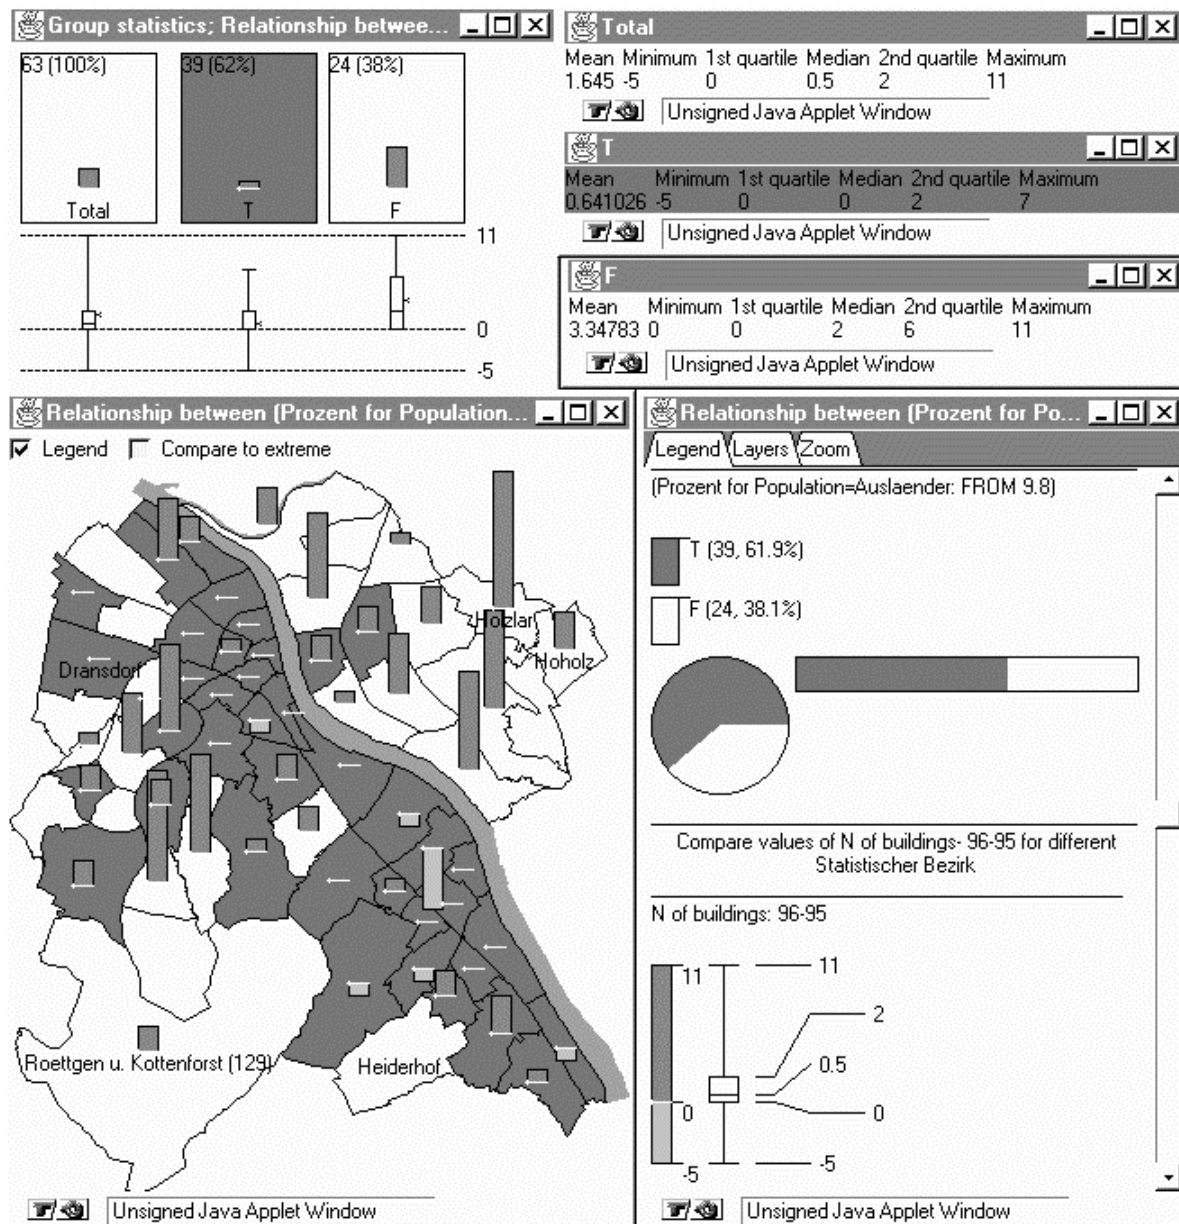


Fig.10. Group statistics supports the hypothesis about inverse correlation between the number of foreigners and the increase in the stock of living houses.

6. Realisation notes.

The system IRIS is implemented in two variants. One is written in C++ and runs on IBM PC and compatibles under Windows 3.x or 95. Another variant is designed for running in the WWW. It has the client-server architecture. The server part is written in C++ and operates on a Unix-based (Solaris or Linux) WWW server. It communicates with the client part written in Java and running on any WWW client computer under Netscape, Internet Explorer or other WWW browser enabling Java. The server performs all operations over data and presentation design. The results are sent to the interface part via a socket connection for displaying on the screen. A result of map design has the form of map specification that is realised by the interface part.

This approach differentiates IRIS from other tools providing cartographic presentations in the WWW [Grossley 96]: they are based on CGI interface, that is, building on a server and

transferring via the Internet of a static raster picture. The latter approach cannot provide such wide range of interactive actions that we considered essential to implement in IRIS. A little more interactivity can be provided by the use of the Tcl/Tk plug-in software working with WWW browsers, as in the project Argus [Dykes 97]. However, this is still not sufficient for data exploration.

Implementation of the user's interface in Java provided the opportunity of access to the system with the use of standard software - Internet browsers. At the same time platform independence was obtained: at the present moment Java programs can work on all widely used hardware and software platforms. Realisation in Java allowed wide access to the system via the WWW and brought interesting material for analysis due to the registration of actions of people that tried to work with the system. The analysis gave us some ideas that allowed improving the user's interface of the system. The independent Java applet rating service⁴ included our system into the top 1% of Java applets.

The WWW realisation of the system allowed revealing some shortcomings of the present state of Internet and Java technology:

- Insufficiently developed are graphical libraries of Java: the current version has no means for hatching and using other textures for filling areas.
- There are many bugs in language interpreters and Internet browsers. In fact, each implementation of Java has its own peculiarities in behaviour. Testing on different software and hardware platforms was needed to ensure satisfactory work of the system.
- Relatively slow and unstable work of Internet obstructs remote access to large programs and work of the programs in client-server mode.

7. Conclusions.

The system was tested in various domains such as accounting of forest resources of Russia, analysis of economic situation in the regions of Russia, population and economic indices for European countries, demographic information for Bonn City. In all cases the maps built by the system were consistent with characteristics of data under analysis. Correctness of reflection of the data was proved by the fact that the maps built by IRIS made apparent the peculiarities of the analysed phenomena previously known to the domain experts. In some cases the maps allowed the experts to reveal interesting facts that were formerly unknown to them. We are sure that such potential capabilities are of primary value in exploratory graphics.

The goal of IRIS development was the study of opportunities of providing intelligent support to users in data analysis by presenting the data in a way that facilitates and stimulates their exploration. From this perspective we regard as merits of the systems the following features:

- The user is completely released from the necessity to think about data presentation, s/he is required only to select a data subset for analysis. This gives the user the opportunity to concentrate directly on problem solving.
- The system, when possible, provides the user with several maps showing the same data. Each of these visualisations gives somewhat different view on the data and allows seeing

⁴ URL <http://www.jars.com/>

different properties of them. In this way conditions are created for comprehensive study of the data.

- Making a picture is not seen as the final goal of the system's work. Contrary, interactive facilities support further work of the user on data exploration.

REFERENCES

- [Bertin 83] J. Bertin, *Semiology of graphics. Diagrams, networks, maps*. The University of Wisconsin Press, Madison WI, 1983.
- [Casner 91] S.M. Casner, A Task-Analytic Approach to the Automated Design of Graphic Presentations, *ACM Transactions on Graphics*, 10(2), 1991, pp 111-151.
- [Dykes 97] J.A. Dykes, Exploring Spatial Data Representation with Dynamic Graphics, *Computers & Geosciences*, 23(4), 1997, pp 345-370.
- [Grossley 96] D. Grossley, T. Boston, A Generic Map Interface to Query Geographical Information Using the World Wide Web, *World Wide Web Journal*, 1, 1996, pp 723-735.
- [Jung 95] V. Jung, Knowledge-Based Visualization Design for Geographic Information Systems, in *Proceedings of the 3rd ACM International Workshop on Advances in Geographic Information Systems*, Baltimor, USA, 1995, ACM Press, pp 101-108.
- [Mackinlay 86] J. Mackinlay, Automating the Design of Graphical Presentation of Relational Information, *ACM Transactions on Graphics*, 5(2), 1986, pp 110-141.
- [Murray 94] B.S. Murray, Intelligent Information Presentation Systems, *The Knowledge Engineering Review*, 9(3), 1994, pp 269-286.
- [Roth 90] S.M. Roth, J. Mattis, Data Characterization for Intelligent Graphics Presentation, in *Proceedings of SIGCHI'90: Human factors in computing systems conference*, Seattle WA, USA, 1990, ACM Press, pp 193-200.
- [Senay 94] H. Senay, E. Ignatius, A Knowledge-Based System for Visualization Design, *IEEE Computer Graphics and Applications*, 14(6), 1994, pp 36-47.
- [Tufte 83] E.R. Tufte, *The Visual Display of Quantitative Information*, Cheshire, Connecticut, Graphics Press, 1983/
- [Tukey 77] J.W. Tukey, *Exploratory Data Analysis*. Reading, MA: Addison-Wesley Publishing Company, 1977.
- [Zhan 95] F.B. Zhan, B.P. Buttenfield, Object-Oriented Knowledge-Based Symbol Selection for Visualizing Statistical Information, *International Journal of Geographical Information Systems*, 9(3), 1995, pp 293-315.