Technical Disclosure Commons

May 01, 2019

Encoding Auxiliary Data into Raster Images

Sascha Brawer

Encoding Auxiliary Data into Raster Images

Sascha Brawer — April 30, 2019

Abstract

In distributed Geographic Information Systems such as web maps, it is common to split large images into raster image tiles. Traditionally, clients request auxiliary data for an image region (such as a pixel position on which the user has clicked, tapped or hovered) by issuing a separate network request for additional information. However, such network round-trips slow down the user experience, particularly on slow or unreliable connections such as mobile networks. What is disclosed is a system and method for encoding auxiliary data into raster images so that remote clients can retrieve small amounts of encoded data without issuing separate network requests.

Keywords

- Geographic Information Systems (GIS)
- web maps
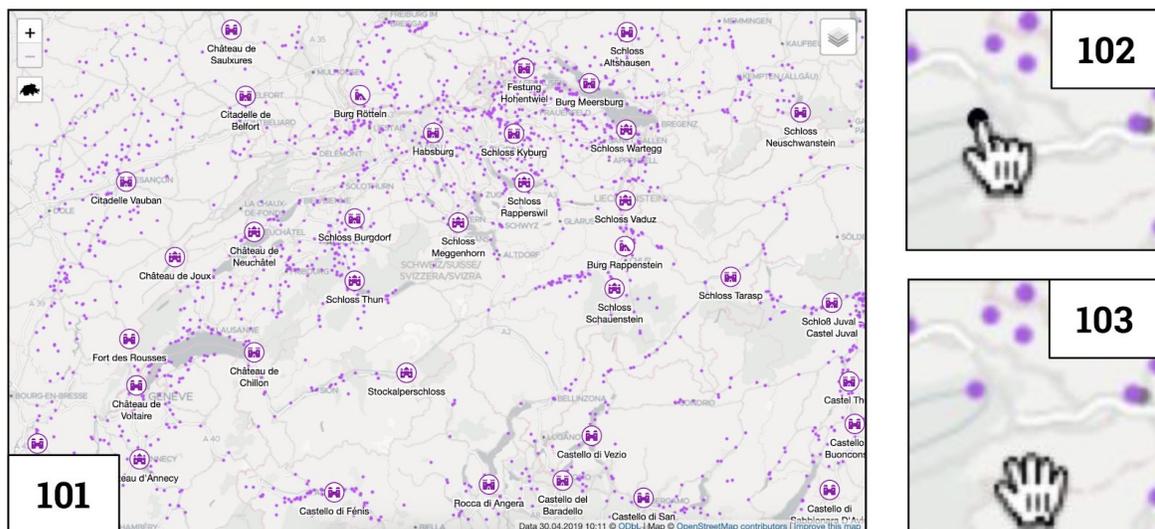- latency
- raster tiles
- data encoding
- side channel

Fig. 1: An example application of the disclosed method is a clickable geographic map with many features (101), transmitted from server to client as a raster image. When the user moves their mouse pointer, the client-side part of the system needs to quickly adjust the pointer depending on whether the hovered-over pixel can be clicked (102) or not (103). The disclosed method allows making such decisions without separate network requests. Additional data can also be embedded, such as the center of the depicted feature which is needed for placing a marker (in 102, a black circle underneath the mouse pointer).

Background

In distributed Geographic Information Systems such as web maps, it is common to split large images into raster image tiles. Each tile consists of pixels with a color value, typically in a color space such as RGB or sRGB where colors gets encoded as numbers corresponding to red, green, blue and (if needed) transparency. Typical image formats use eight bits for each color channel, leading to $2^{3\times8} = 2^{24}$ different colors with 2^8 transparency values. A common technique for reducing the file size of an image is to restrict the colors to a limited palette with 16 or 256 colors; however, in this case each single palette entry can still have an arbitrary color value.

0	0	128	192	128	0	0	0
0	128	255	255	255	128	0	0
0	192	255	255	255	192	0	0
0	128	255	255	255	128	0	0
0	0	128	192	128	0	0	0
0	0	0	0	0	0	0	0

Fig. 2: In traditional raster tiles, anti-aliased shapes are painted in different transparency values ranging from 0 to 255.

Disclosed method

When representing colors and transparency in raster images (either directly or indirectly via palette entries), the least significant bits of each channel can be used for encoding auxiliary data. Although this modifies the image, small differences in color or transparency are difficult to perceive for humans. Therefore, colors and transparency can be used as side channel for transmitting auxiliary information.

For example, the server component in a client-server or web geographic information system can use this method for encoding the center points of painted shapes; the client component can then handle user events (such as changing the shape of the mouse pointer upon hovering; or positioning a highlight marker) just by interpreting the numeric color and transparency values in the raster image tile that has already been transmitted to the client. Traditionally, client-server geographic information systems would have to issue additional network requests for retrieving such auxiliary information; due to the network latency, this hurts user experience, especially on slow and unreliable networks. By using the disclosed method, systems can improve their user experience.

0	0	128	192	128	0	0	0
0	128	255	255	255	128	0	0
0	192	255	254	255	192	0	0
0	128	255	255	255	128	0	0
0	0	128	192	128	0	0	0
0	0	0	0	0	0	0	0

Fig. 3: Small variations in color or transparency are difficult to perceive for humans. Therefore, they can be used for encoding auxiliary data into an image. In this example, the transparency value 254, which looks almost identical to the fully opaque 255, marks the center point in a geometric shape. The method can be used for encoding small amounts of arbitrary other auxiliary data into color channels, transparency, or both.