University of Groningen

Visualizing Metrics on Areas of Interest in Software Architecture Diagrams

Byelas, Heorhiy; Telea, Alexandru

Published in:
IEEE PACIFIC VISUALIZATION SYMPOSIUM 2009, PROCEEDINGS

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2009

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Byelas, H., & Telea, A. (2009). Visualizing Metrics on Areas of Interest in Software Architecture Diagrams. In P. Eades, T. Ertl, & HW. Shen (Eds.), *IEEE PACIFIC VISUALIZATION SYMPOSIUM 2009, PROCEEDINGS* (pp. 33-40). IEEE (The Institute of Electrical and Electronics Engineers).

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Visualizing Metrics on Areas of Interest in Software Architecture Diagrams

Heorhiy Byelas*

Alexandru Telea†

Institute of Mathematics and Computer Science
University of Groningen, the Netherlands

ABSTRACT

We present a new method for the combined visualization of software architecture diagrams, such as UML class diagrams or component diagrams, and software metrics defined on groups of diagram elements. Our method extends an existing rendering technique for the so-called areas of interest in system architecture diagrams to visualize several metrics, possibly having missing values, defined on overlapping areas of interest. For this, we use a solution that combines texturing, blending, and smooth scattered-data point interpolation. Our new method simplifies the task of visually correlating the distribution and outlier values of a multivariate metric dataset with a system's structure. We demonstrate the application of our method on component and class diagrams extracted from real-world systems.

Index Terms: I.3.4 [Graphics Utilities]: Graphics editors—Paint systems; D.2.2 [Design tools and techniques]: Modules and interfaces—Computer-aided software engineering

1 INTRODUCTION

Software architecture diagrams are widely used in software engineering. They describe the structural and functional relations between different system elements, such as objects, interfaces, and components. Besides diagrams, software metrics, which describe the quality attributes of the elements present in the diagram, play an essential role in forward and reverse engineering, reengineering, and maintenance. In all such activities, software architects need to easily correlate several metrics, computed on a system, with the system structure, captured by its architecture diagram.

In UML class diagrams [8], metrics can be associated with classes, relations between classes, groups of classes, or class members, *e.g.* methods. To correlate metrics with each other and with the system architecture, we must find effective ways to combine the presentation of several metrics and the architecture in a single picture.

We present here a new way to visualize several software metrics, modeled as a multivariate dataset, together with a system's structure, modeled as an architecture diagram. We specifically consider metrics defined on software components (*e.g.* classes), or groups thereof, also called areas of interest (AOIs) [3]. We call these class-level metrics. We extend an existing technique for rendering AOIs by using blending and texturing to render several metrics, defined as a multivariate dataset with potentially missing values, atop of such areas, so that users can spot metric-metric and metric-area correlations. Our technique removes the need of drawing metric icons atop of the diagram elements, so we use this space to show other information, such as class member metrics or text annotations.

*e-mail: h.v.byelas@rug.nl

†e-mail: a.c.telea@rug.nl

This paper is structured as follows. Section 2 reviews related work in visualizing the combination of system structure and software metrics. Section 3 presents our new technique for rendering several metrics atop of AOIs. Section 4 presents two case studies of using our new rendering technique on three UML diagrams. Section 5 discusses the obtained results. Section 6 concludes the paper.

2 RELATED WORK

Software system structure is typically visualized using diagrams, such as UML diagrams for object-oriented systems [8] or similar metaphors for component-based systems. Metrics can be computed by static analysis tools [16], simulation tools [1] or dynamic analysis tools such as profilers or debuggers. Traditionally, metrics are shown as separate numerical tables, making their correlation with large diagrams difficult and time-consuming.

Several attempts have been made to combine metrics and UML-like diagrams. Lanza *et al.* render class-level metrics by mapping them to the class size and/or color [7]. In this way, two metrics can be shown simultaneously. Similar techniques are used by many software visualization tools, such as the well-known Rigi toolkit [13]. An extensive overview of such tools is given by Diehl [5]. Often, however, diagram element sizes may be constrained (fixed) to a predefined layout, so they cannot be used to show a metric. This happens when UML diagrams are created in the design process: designers carefully craft the layout of diagrams and would not accept the elements' sizes or positions to be modified, as this destroys their 'mental map'. Element background colors may also be constrained *e.g.* when we want to draw method names or other text annotations inside each class. Termeer *et al.* show UML class-level metrics with icons scaled, colored, and drawn atop classes [12]. This technique can show individual metric values and helps spotting outliers. However, correlating *several* metrics on a *large* diagram is difficult, as one has to memorize the metric values while visually scanning the diagram.

Drawing metrics using icons or modifying the size or color of diagram elements emphasizes the relation of a metric with a single element. However, we may want to visualize metrics defined on *groups* of elements, *e.g.* safety of all multithreaded components, speed of all performance-critical components, and so on. Byelas and Telea partially addressed this: They surround groups of related elements, also called *areas of interest* (AOIs), with a smooth contour, similarly to the way humans draw such groups on paper diagrams [3]. This shows which elements are in which AOI, *e.g.* all multithreaded elements in our example. However, AOIs themselves cannot show metric values (*e.g.* the safety metric in our example) so correlating structure with metrics is still a problem.

3 PROPOSED METHOD

First, we introduce our data model. Consider a system diagram with n areas of interest $A_1 \dots A_n$ defined over its elements, where $e_{ij}, j \in [1, |A_i|]$ are the elements in area A_i and $|A_i|$ is the number of elements in area i . For each A_i , we have a metric $m_i : [1, |A_i|] \rightarrow \mathbb{R} \cup \text{None}$ defined over its elements¹. m_{ij} , the value of m_i on e_j ,

¹Several metrics defined on the same area are handled analogously

Figure 1: Smooth interpolation of element metrics over an area-of-interest

can have either a numerical value, or *None*, if that value is missing. Missing metric values are frequent in software analysis, *e.g.* due to various limitations of the analysis tools [16]. The set of metrics m_i can be seen as a multivariate scattered-point dataset [11], with elements j as data points and the metric values i as variables. We want to show all metric values for all areas in one image, so that

- we can compare the metric values of each element,
- we can visually follow how a metric varies over an area,
- we see the elements having missing values,
- we do not draw metrics on the elements themselves.
- we use, but do not change, a given diagram layout

The original AOI method [3] constructs a contour that encloses the elements located in an area of interest (see *e.g.* Fig. 1 b). We show next how to render several metrics so that metric values, elements, and areas of interest can be easily correlated. We use a two step solution. First, we render the values of a single metric m_i over a given area A_i (Sec. 3.1). Next we combine all metrics m_i for all areas A_i in a single image (Sec. 3.2). Finally, we add shading to the areas to further emphasize their structure (Sec. 3.3).

3.1 Rendering a single metric

Termeer *et al.* show element metrics using icons scaled and colored by metric values, drawn atop of the elements [12]. This has several drawbacks. Consider a diagram with five metrics over five areas of interest (Fig. 5). First, icon sizes are constrained by the element sizes (which can be small), so it is hard to see the specific metric values. Second, we want to keep the element surfaces free to draw other data, such as method names and annotations. Third, correlating metrics with areas of interest, *e.g.* seeing how metric values change over one or several areas, is difficult, since there is no explicit visual correspondence (mapping) from metrics to areas.

We address these issues by rendering metric values *outside* the diagram elements. Denote by $\{e_i\}$ the elements in area A , with metric values m_i - we drop area-indexes here since we consider a single area. We encode missing metric values in a separate dataset $p_i : [1..|A|] \rightarrow \{0, 1\}$, *i.e.* set p_i to 0 if m_i is missing, else set p_i to 1.

Our key idea is to produce an interpolation function $\mathcal{M}$ of the values m_i over area A . $\mathcal{M}(x)$ should equal the given metric values m_i for points x inside or close to the elements e_i , and vary smoothly in-between. We compute $\mathcal{M}$ as follows. First, we compute the Delaunay triangulation of A using the Triangle library [10]. Next, we initialize $\mathcal{M}$ at each triangulation vertex x with the metric value $m(e_{closest})$ of the element

$$e_{closest} = \underset{i \in [1..|A|], m_i \neq None}{\operatorname{argmin}} (||e_i - x||)$$

i.e. the closest element to point x which has a metric value. This yields an approximation of the Voronoi diagram of the element set $\{e_i\}$, so $\mathcal{M}$ is a piecewise-constant interpolation of $\{m_i\}$ over A . Figure 1 b shows $\mathcal{M}$ for the metric values in Fig. 1 a, using a red-to-blue colormap². Element D has a maximum value, as shown by the surrounding red color. Element A has a minimum value, shown by the blue color. Elements E and F do not belong to the area. Element B , although inside the area, has no value. We show this using a neutral gray hue, as follows. We compute an interpolation $\mathcal{P}$ of the set $\{p_i\}$ over A , just as the interpolation $\mathcal{M}$ of $\{m_i\}$. With $\mathcal{M}$ and $\mathcal{P}$, we now compute the hue-saturation-value color of any point $x \in A$ as

$$h(x) = \text{rainbow}(\mathcal{M}(x)) \quad (1)$$

$$s(x) = \mathcal{P}(x) \quad (2)$$

$$v(x) = 1 \quad (3)$$

where $\text{rainbow}()$ is the chosen colormap (see Fig. 1 d). Hence, points having metric values are rendered with saturated colors, while points with missing values are gray. Finally, we render the area's border using a soft gray texture.

In the final step, we smooth our piecewise-constant interpolation. For this, we apply a Laplacian filter [4] on $\mathcal{M}$ and $\mathcal{P}$, by setting the value of each triangle vertex x to the average value of all vertices connected to it, and repeating the process for 30..50 iterations. The points contained inside the elements e_i are kept fixed to the prescribed metric values m_i , to enforce the interpolation's boundary conditions. The result shows the values m_i close to their elements e_i , smooths values in-between, and grays out colors close to elements without values (see Fig. 1 c).

3.2 Combining several metrics

Now we must combine several metrics defined on possibly overlapping areas. We cannot simply additively blend areas of different colors as in [3], as this would mix the individual colors which show metric values beyond recognition. We use a texture-based solution: For each area A_i , we use a different texture³. We carefully designed a small set of textures (Fig. 2). The overlap of any textures in this set creates a visually different pattern. The textures contain just opacity: black denotes opaque zones, white gaps are fully transparent, gray indicates an alpha value between 0 and 1.

We now render each area A_i by combining its color (showing metrics) computed by interpolation (Sec. 3.1) with its transparency texture (showing the area's identity) using OpenGL's texture modulation. Figure 1 d shows the application of texture c from Fig. 2 on

²The colormap choice is discussed further in Sec. 5

³We can relax this: non-overlapping areas can use the same texture, see Fig. 6.

Figure 2: The proposed set of textures. Gray value denotes opacity

the area in Fig. 1 c. To maximize information visibility, we draw areas starting from the largest to the smallest one, so that small areas appear atop large areas.

Figure 3 shows three overlapping areas defined over four elements. Transparency creates hole-like patterns that let us see which textures, *i.e.* which areas, overlap, since each area has a different texture. The visual ‘weaving’ of the textures also lets us distinguish their different colors, hence correlate metric values. For example, we see that *D* has low values in area 1 and high values in area 3 - blue circles atop red diagonal lines; *B* has high values in area 1 and no value in area 3 - red circles atop gray diagonal lines; and so on.

Transparency acts more like a stencil, so there is little or no actual blending; colors do not mix, but get spatially woven. Color interpolation spreads the metrics information from elements over entire areas, creating large smooth hue spots which are easier to follow than rapid changes. We acknowledge this is a controversial issue: color blending may suggest that there is a continuous metric variation over an AOI, which is not the case. If less blending is perceived as better, one can simply do less smoothing iterations: See *e.g.* Fig. 10 where only a few iterations are done, which yields well-separated color areas around the elements, and almost no interpolated colors. For instance, the transition between blue and green in *A*₂ is sharp and quick. Also, one can use discrete (categorical) colormaps with no change in the method, if these are seen to produce less ambiguous results.

Figure 3: Diagram showing three areas of interest with metrics

3.3 Shading for enhanced area separation

Although each area has its own distinctive texture, this can create confusing overlaps where it is hard to tell where an area exactly stops and another one starts. This happens *e.g.* where contours of different areas run almost tangent.

To alleviate this, we emphasize each area *A* by *shading*, as follows. We construct a signal $\mathcal{S}$ over *A* that is zero on the contour ∂A of *A*, one further from the contour, and varies smoothly with the distance within a narrow band of thickness δ along the contour. We compute $\mathcal{S}$ on the same triangle mesh as $\mathcal{M}$ and $\mathcal{P}$ used for the color interpolation (Sec. 3.1), as follows. First, we set $\mathcal{S}$ to 0 on the contour vertices and 1 elsewhere. Next, we use the same Laplacian filter as for color smoothing, keeping $\mathcal{S}$ fixed to 0 on the contour points, for 10..30 iterations. More iterations increase the thickness of the shading effect. After each iteration, we renormalize $\mathcal{S}$ to the range $[0, 1]$.

We now use $\mathcal{S}$ as luminance by setting $v(x) = \mathcal{S}$ in Eqn. 3. This darkens areas close to their borders, but keeps them bright in the middle. Normalization ensures that shading is always bright in the middle of an area and dark on the contour. A direct application of shading would only affect the texture stripes (non-transparent) but would not show up in the texture ‘holes’. This would create a broken, distracting shading effect. We prevent this by increasing the holes’ opacity in the texture patterns from 0 (fully transparent) to 0.2 (slightly opaque).

Overall, we obtain the effect of convex, shaded 3D shapes - compare Fig. 3 (no shading) with Fig. 4 (with shading). At overlaps, the shaded shapes get woven by blending. The darkened borders help to visually separate areas (see the images in Sec. 4). The slight opacity of the texture pattern holes is able to show the shading close to the areas’ contours and also a faint hue of the interpolated colors, *i.e.* metrics, in the pattern holes. This further strengthens the visual cohesion of all elements within an area and limits the breaking effect of the holes, but still allows pattern weaving to take place. When using textures to show metrics, as users noted on several occasions, textures seem to complicate the visual tracking of an area’s contour, so shading has a stronger value for areas textured to show metrics.

Figure 4: Enhanced areas using shading (compare with Fig. 3)

4 APPLICATIONS

We now illustrate the use of our multivariate metric-and-structure visualization in two different case studies.

4.1 Case Study - JPEG Decoder Architecture

We consider a real-world software project: the architecture of a component-based JPEG decoder [2, 14]. The system model was built and its operation numerically simulated using the CARAT toolkit [1]. This delivered several run-time performance metrics. We next show two such metrics:

- μ_{CPU} : CPU usage for active components (each active component has its own process)
- μ_{mem} : memory usage for passive components (a passive component is used by active processes)

Given the actual architecture of the JPEG decoder, not all components have both memory and CPU metric values.

The decoder performs five tasks ($T1 \dots T5$): JPEG stream starter ($T1$), inverse discrete cosine transform (IDCT), IDCT column process ($T2$), IDCT row process ($T3$), rasterization ($T4$), and rendering ($T5$). For a detailed description, we refer to [2]. We consider six areas: $A_1 \dots A_5$ contain the components in tasks $T1 \dots T5$. Each component has a memory usage metric μ_{mem} for each task area it is part of. The sixth area A_{CPU} holds all active components, which also have a CPU usage metric μ_{CPU} . We now address two goals which were named as important by the developers:

- understanding the distribution of tasks over the system structure and the memory usage of passive components
- understanding the CPU utilization over different tasks

To illustrate the advantage of our method, we first use metric icons [12]) to show the memory usage metric. First, we draw the areas $T1 \dots T5$. Next, we draw pie and height-bar icons colored by task and scaled to show memory usage μ_{mem} (Fig. 5). The metric legend shows the tasks' colors and also shows where each icon from each task-area is placed within each element (see [12]). However, in Fig. 5 it is hard to tell the metric values of each component for each area it belongs to. We cannot increase icon sizes, as each icon already takes one-sixth of a component's size. It is hard to visually correlate metric values over large areas. Also, a missing icon has an ambiguous meaning: does it show a zero $\mu_{mem} = 0$ or missing metric value or a missing metric value $\mu_{mem} = None$?

We now use our new technique. Each area (task) uses a different texture (see legend in Fig. 6 left). Color shows the memory usage μ_{mem} (blue=low, red=high). We now better see which value μ_{mem} each component has in each area, even though the images in Fig. 6 are half the size of the icon-based visualization in Fig. 5. We see, for instance, that components A , C , D , E and F use much more memory than the rest in at least one task they are involved in. Components A and C consume high memory amounts in the tasks they are involved in ($T1$ and $T3$ for A and $T2$ and $T4$ for C). Component C is the main memory consumer of the entire system, as both textures surrounding it are red. Indeed: C implements the decoder's pixel raster buffer, which consumes a lot of memory. Finally, we see that components $D \dots F$ have a similar memory usage pattern: low in task $T4$, high in task $T5$. The results match the design expectations, as rendering ($T5$) is more memory-demanding than rasterization ($T4$).

In our second scenario, we add the CPU utilization metric μ_{CPU} (Fig. 6 right). The area A_{CPU} , containing all active components using CPU cycles ($G \dots K$), intersects the task-areas $T1 \dots T5$. To visually segregate the two aspects (tasks and CPU utilization), we use diagonal stripes for the task-areas $T1 \dots T5$ and vertical stripes for the CPU utilization area A_{CPU} . We see now the CPU-intensive components: J and K . We also see that all components

in A_{CPU} miss memory consumption data: the diagonal stripes textures around all components ($G \dots K$) are gray (Fig. 6 left). This is correct, as the design of this JPEG decoder splits data (passive) components from algorithm (active) components.

Figure 7 shows the effect of adding shading. The left image depicts the six areas with color interpolation (showing metrics) but no textures. We provide this image to emphasize the useful effect of shading to understand area overlaps. The right image shows the six areas and two metrics (memory and CPU usage). Compared to Fig. 6 right, it is easier to tell in the shaded image which components are in which areas.

4.2 Case Study - Large Class Diagrams

In our second application, we extract an UML class diagram from the source code of a C++ graphics editor in a reverse engineering process, using an ANTLR-based C++ parser [9]. Talking to the system designer, we identified several high-level functional aspects:

- *main*: the application's entry point
- *core*: the application's control code
- *logging*: code involved in logging actions
- *GUI*: user interface code
- *I/O*: code for saving and loading data
- *OpenGL*: rendering code
- *XML*: code for loading 3D models

Each aspect yields an area-of-interest A_i . We now want to see which class participates in which design aspect, and how much. An ideal object-oriented design would require each class strongly involved only in one aspect [7]. We quantify the participation degree p_{ij} of each class j in each aspect A_i as its code percentage specific to A_i . For example, an OpenGL class has $p = 0.5$ if it has 50% OpenGL-specific code. The goal is to understand how the identified aspects map to actual classes, *i.e.* whether the code follows the intended design, and whether we have modularity problems.

The entire system is shown in Fig. 8. The legend shows, for each area A_i , the number of classes it contains, the number of classes having missing values for that area's metric p_i (due to the fact that we were unable to reliably estimate the percentage of code involved in each aspect), and the texture used to show the area. We notice several facts. Few classes participate in two aspects, and none take part in three. This indicates a good functional modularity. The only class strongly involved in two aspects is B , part of the *main* and *core* areas. Since B is actually the system's entry point, this strong involvement is not a problem. Class E participates strongly in *core* (red in A_6) and weakly in *GUI* (blue in A_1). E the main window, so its weak involvement in *core* and strong in *GUI* is correct. Class D is strongly I/O-related (A_7), and also part of the *core* (A_6). However, its code is quite complex, so we were unable to assess how strongly it belongs to the *core* (missing metric of D in A_6).

Figure 9 shows the same diagram, areas, and metrics as in Fig. 8, with shading added. As for the JPEG decoder example, shading helps better seeing which elements are in which areas.

Figure 10 shows another class diagram: a part of our own UML visualizer. We show two functional areas: classes involved in visualization (A_2), and the class hierarchy modeling a UML graphical element, or glyph (A_1). Colors show degrees of participation in the two aspects. Since our metric-rendering does not draw on classes, we can show an additional metric: the lines-of-code (LOC) for all class methods, drawn atop of classes with purple bars. Long bars indicate large methods. Methods are sorted in decreasing LOC from

Figure 7: JPEG decoder architecture. Left: shaded AOIs. Right: Adding shading to textured AOIs (compare with Fig. 6 right)

Figure 8: Large UML class diagram with 7 areas and over 50 classes. Metrics show the participation of classes in two aspects

base includes around 10 professional software engineers involved in creating UML architecture diagrams, such as the JPEG decoder (Sec. 4.1), and computing quality metrics on them.

We asked the designers to utilize both the icon-based and texture-based metric visualizations to present their own work (diagrams and metrics) in around 10 project meetings of around 20 participants over a period of about 1 year. In such presentations, important goals are to show metric-metric and metric-structure correlations, as described in the previous sections. We silently observed the presentations and gathered off-line feedback from presenters and participants. Overall, there was a strong positive feedback about using areas of interest: they are simple to understand and effective to show software aspects. Metric icons were accepted only when showing a *single* metric, possibly over several areas. Texture-based metric visualizations were seen as more effective and intuitive when correlating *several* metrics. The overlap of more than three metric textures was, however, very hard to understand. In such cases, presenters would switch metrics on and off to show only three metrics simultaneously. Interestingly, color smoothing was not seen as a problem, even though it generates colors between the diagram elements which do not correspond to actual values in the data. When talking about this issue, we got the impression that users focus pre-

dominantly on the colors *close* to the diagram elements *and* use color smoothing as a visual cue to navigate from element to element over a given area, being aware that in-between colors do not represent data.

5 DISCUSSION

We discuss several aspects of our technique, as follows.

Scalability: we can easily show up to 10 areas of interest, each with its own metric, on diagrams of 20..80 of classes. Larger diagrams occur very rarely in software engineering practice. The Delaunay triangulator and Laplacian filter used are well-known for their fast, subsecond performance on meshes of thousands of triangles. Rendering a metric over an AOI uses a single texture pass over a triangle mesh, which is also very fast on any graphics card.

Understandability: The main limitation is the number of distinct areas that can overlap at one given place. Consider the AOIs $A_1 = (A, B, C, D)$, $A_2 = (A, B, C)$ and $A_3 = (A, B, D)$ in Fig. 11, rendered with textures shown in the legend. From the 'woven' texture pattern we believe it is reasonably easy to see which element is in which area and the colors (metric values) at overlaps. The addition of

Figure 9: Visualization of the UML diagram in Fig. 8, now with area shading and half-transparent elements

shading (Sec. 3.3 further helps in separating areas with complex overlaps. Yet, adding a fourth overlapping area would make this image hard to understand.

Figure 11: Complex intersection of three overlapping areas

Obtaining a good pattern mix constrains the texture parameters. All textures should have similar ratios of opaque-to-transparent pixels, so we can 'see through' at all overlaps. Ratios between 40% and 60% give good results - lower values yield too sparse textures, on which we cannot see colors or shading; higher values yield occlusion at overlaps, so we cannot see more than one texture. Patterns must be chosen so that the overlap of *any* $n - 1$ patterns looks different from the n^{th} pattern, n being the number of overlapping areas. The texture set used here gives good results for $n \leq 3$, as shown in

a different application [15]. Finally, the frequency range (related to the pattern stripe thickness and circle radius) is important. Too thin patterns are hard to distinguish at overlaps; too thick patterns do not let the eye smoothly switch between areas at overlaps. We found an empirically good pattern size in the range of 10..20 pixels (Fig. 2).

Related methods: To our knowledge, there is only one other software visualization that uses textures to show numeric metric values [6]. Our method differs from this as follows. Holten *et al* encode two metrics in the texture frequency and luminance, and use a treemap layout, so their areas are rectangular, cannot overlap, and always contain a single element. We smoothly interpolate metrics over arbitrarily-shaped, overlapping areas. We use a fixed texture-set, use opacity to allow overlaps, and encode metric values in hue and metric availability in saturation. Finally, we use luminance to pseudo-shade the areas to visually emphasize contours rather than encoding data. This is conceptually similar to the cushion treemaps used by Holten *et al*, but generalizes to complex-shaped, overlapping, areas.

The colormap choice can be considerably Here, we demonstrate only a simple blue-to-red continuous colormap, for simplicity and conciseness. However, better choices are available, such as other hue gradients or discrete few-hue colormaps. Such issues need to be further investigated. A second discussion point is our choice to interpolate colors. As mentioned, this creates colors between elements which do not correspond to actual values. However, we believe this is acceptable since users are fully aware that there are no data values except on the diagram elements, and smooth colors help following the contents of a given area as opposed to hard color boundaries between elements. Also, using a (discrete) few-hue colormap would considerably alleviate this problem as there would be less, or no, different hues created between the ones in the colormap. Still, a rigorous user evaluation of the effectiveness and/or limitations of color interpolation is still needed.

Finally, using full-saturation hues on the areas is sometimes seen as distracting. Fortunately, this is easy to tune: we provide a global opacity control that allows users to set the overall opacity of all

Figure 10: UML class diagram with two areas, class-level participation metrics, and method-level lines-of-code metrics

AOIs, thus smoothly navigating between 'bare' diagrams and diagrams with full-saturation textured areas. In practice, using a global area opacity of 0.4..0.6 gives good results - the actual value used depending on one's taste and type of color screen.

6 CONCLUSION

We proposed a method to render a multivariate set of metrics, with potential missing values, on elements of areas of interest on UML diagrams, so that metric-area correlations and distributions of metrics over areas are easy to distinguish. We use a combination of texturing, blending, and smooth spatial data interpolation techniques. Texture patterns encode aspects in a diagram and also allow textures to visually interlace and show the colors of several metrics. Additional shading further visually separates complex overlapping areas. Although our interest is in software architecture diagrams, our method can be used in many other contexts, such as organization diagrams or spatial maps.

We would next like to study how interaction can help understanding the metrics correlation. A second essential direction is to study the practical effectiveness of the proposed methods by means of user studies involving actual software engineers in the industry.

ACKNOWLEDGEMENTS

We would like to thank Egor Bondarev (TU Eindhoven) for supporting us with the JPEG case study and the ITEA Trust4All consortium for supporting part of our research.

REFERENCES

- [1] E. Bondarev, M. Chaudron, H. Byelas, and P. de With. A toolkit for design and performance analysis of real-time component-based software systems. In *Proc. Intl. Conf. in Software Eng. Advances*, pages 4–8, 2006.
- [2] E. Bondarev, M. Chaudron, and E. de Kock. Exploring performance trade-offs of a JPEG decoder using the DeepCompass framework. In *Proc. Intl. Workshop on Software and Performance*, pages 153–163, 2007.
- [3] H. Byelas and A. Telea. Visualization of areas of interest on software architecture diagrams. In *Proc. ACM SoftVis*, pages 105–114, 2006.
- [4] L. da Fontoura Costa and R. M. Cesar. *Shape Analysis and Classification: Theory and Practice*. CRC Press, 2004.
- [5] S. Diehl. *Software Visualization - Visualizing the Structure, Behaviour, and Evolution of Software*. Springer, 2007.
- [6] D. Holten, R. Vliegen, and J. J. van Wijk. Visual realism for the visualization of software metrics. In *Proc. VisSoft*, pages 27–32. IEEE, 2005.
- [7] M. Lanza and R. Marinescu. *Object-Oriented Metrics in Practice - Using Software Metrics to Characterize, Evaluate, and Improve the Design of Object-Oriented Systems*. Springer, 2006.
- [8] OMG. *The Unified Modeling Language*. 2008. <http://www.uml.org>.
- [9] T. Parr and R. Quong. ANTLR: A predicated-LL(k) parser generator. *Software - Practice and Experience*, 25(7):789–810, 1995.
- [10] J. R. Shewchuk. Triangle: Engineering a 2D quality mesh generator and delaunay triangulator. In *Proc. Applied Computational Geometry*, pages 124–133. ACM Press, 1996.
- [11] R. Spence. *Information Visualization: Design for Interaction (2nd ed.)*. Prentice Hall, 2007.
- [12] M. Termeer, C. Lange, A. Telea, and M. Chaudron. Visual exploration of combined architectural and metric information. In *Proc. VISSOFT*, pages 21–26. IEEE Press, 2005.
- [13] S. R. Tilley, K. Wong, M.-A. D. Storey, and H. A. Miller. Programmable reverse engineering. *Intl. J. of Software Eng. and Knowledge Eng.*, pages 501–520, 1994.
- [14] Trust4All. The Trust4All project, 2005. www.win.tue.nl/trust4all.
- [15] L. Voinea and A. Telea. Multiscale and multivariate visualizations of software evolution. In *Proc. SoftVis*, pages 115–124. ACM, 2006.
- [16] J. Wust. SDMETRICS: The software design metrics tool for UML. 2006. www.sdmetrics.com.