Visual Evaluation of Cloud Infrastructure Performance Predictions

Holger Stitz, Samuel Gratzl, Harald Rogner, and Marc Streit

JOHANNES KEPLER UNIVERSITY LINZ

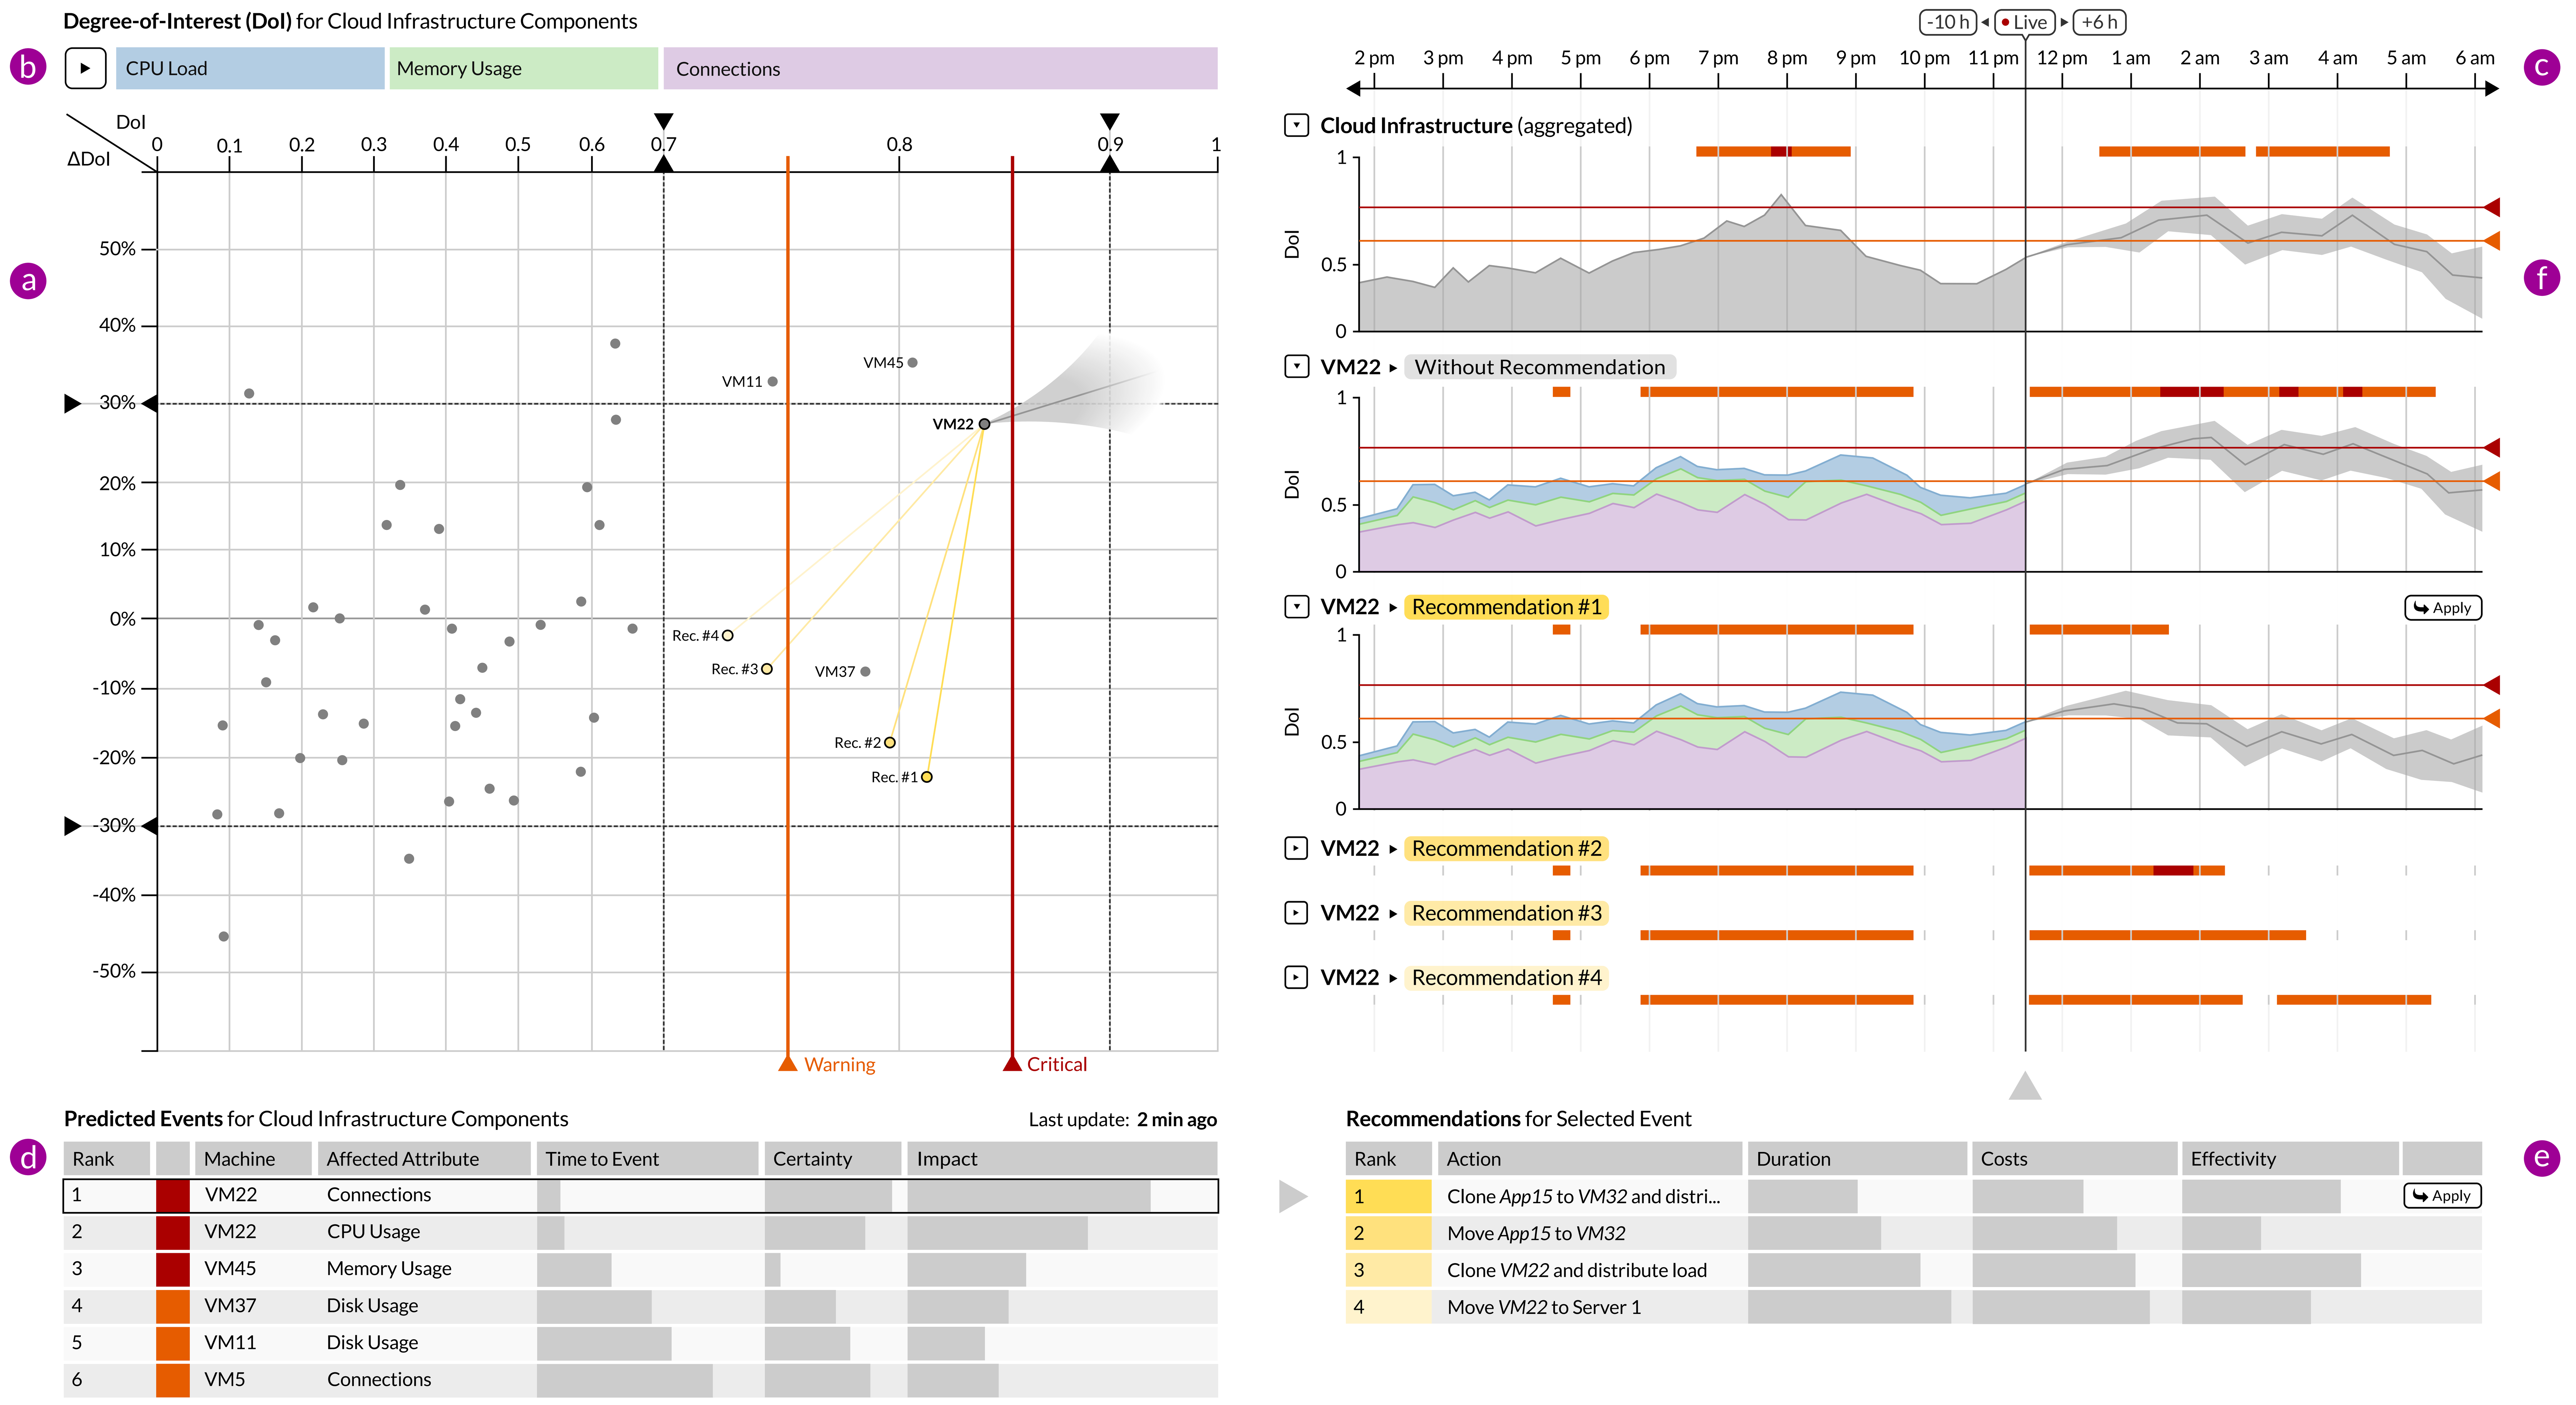

Figure 1: Design sketch of the proposed visualization concept. (a) The administrator can configure the degree-of-interest (DoI) that determines how critical the status of a component currently is. (b) The ThermalPlot shows the current as well as predicted positions of components in the DoI space. (c) Ranked list of predicted performance bottleneck events. (d) Selected time span with past and future. (e) Ranked list of possible countermeasures for avoiding the future selected bottleneck event. (f) Detail view visualizing the past and predicted item performance using peak bars (collapsed) or as streamgraphs (expanded).

The workload of an active cloud infrastructure is continuously changing. Administrators need to find the balance between quality of service and cost efficiency. Proactively avoiding performance bottlenecks is a challenging endeavor due to the number and heterogeneity of components in the network, the relationships among them, and the associated attributes. While pattern detection and simulation methods can be used to predict the performance of the network, efficient tools for exploring and evaluating those predictions are missing. In our design study we combine established visualization techniques (*CloudGazer* [3], *ThermalPlot* [2], and *LineUp* [1]) for exploring large item collections with custom techniques for investigating predicted time-series data, allowing the administrator to effectively monitor, evaluate, and optimize cloud infrastructure.

(a) *ThermalPlot* View

The *ThermalPlot* space maps the criticality of components to the x-axis and the positive and negative change of the criticality to the y-axis (see Fig. 2). Consequently, the more critical a component is, the further on the right it will appear in the *ThermalPlot* space. The position is calculated by the configurable DoI function (see Fig. 1(b)). Two vertical threshold lines discretize the criticality of components into two states: 'warning' and 'critical'. The administrator can freely configure the thresholds of the states by changing the vertical position of the lines via drag and drop.

(b) Degree-of-Interest (DoI) Function Editor

The DoI function is a weighted sum of multiple performance attributes, such as CPU load, RAM usage, and the number of currently opened connections. The exact selection of performance attributes depends on the use case.

(c) Timeline

The timeline shows the selected time range (e.g., 10 hours into the past and 6 hours into the future) for the DoI computation (past to live) and prediction (future). Adjusting the selected time range triggers a re-computation of all DoI values and events.

(d) Predicted Events Ranking

Ranked list of predicted performance bottleneck events. Events are the result of a prediction model that uses the historical and live data collection to discover potential bottlenecks. Each event consists of:

- cloud infrastructure component
- specific attribute that is expected to be critical
- estimated timespan until the event occurs (the shorter the closer)
- certainty of the prediction (the longer the more likely)
- impact value on the perspective or infrastructure (the longer the bigger)

Selected events are visualized in the *ThermalPlot* space with the predicted position and a funnel that encodes the certainty (see Fig. 1 (a)).

(e) Recommendations Ranking

Ranked list of possible countermeasures for avoiding the future selected bottleneck event (see Fig. 1 (d)). Each recommendation consists of:

- description for the action
- predicted duration that is needed to apply this recommendation (the shorter the faster)
- cost estimation (e.g., bandwidth; the longer the more expensive)
- effectiveness estimation to reduce the event's impact (the longer the more effective)

Recommendations are also shown (with their respective color coding) in the *ThermalPlot* space (see Fig. 1 (a)) and in the detail view (see Fig. 1 (f)).

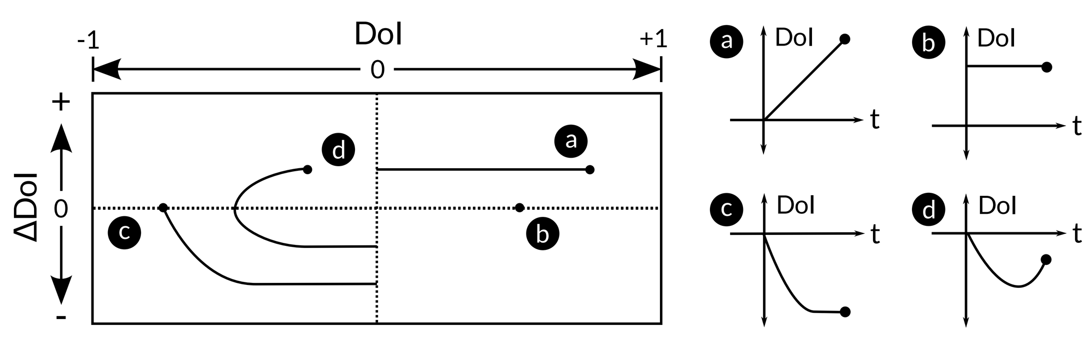

Figure 2: DoI values that change over time result in distinctive positions and trajectories of items in the *ThermalPlot* space.

(f) Detail View

Detail view visualizing the past and predicted item performance using peak bars or as streamgraphs.

Peak bars (collapsed state)
Visualizes DoI values that are above a certain threshold in the corresponding threshold color. The predicted performance changes that are expected for the different recommendations are also visualized using peak bars.

DoI streamgraph (expanded state)
Visualizes the full time-series data within the selected time span. Available data in the past is visualized as a stacked DoI streamgraph that represents how much each attribute contributes to the aggregated DoI value over time. For predicted future performance values, we show the predicted maximum, expected, and minimum DoI value.

[1] S. Gratzl, A. Lex, N. Gehlenborg, H. Pfister, and M. Streit. LineUp: Visual Analysis of Multi-Attribute Rankings. IEEE Transactions on Visualization and Computer Graphics (InfoVis '13), 19(12):2277–2286, 2013.
[2] H. Stitz, S. Gratzl, W. Aigner, and M. Streit. ThermalPlot: Visualizing Multi-Attribute Time-Series Data Using a Thermal Metaphor. IEEE Transactions on Visualization and Computer Graphics, 2016.
[3] H. Stitz, S. Gratzl, M. Krieger, and M. Streit. CloudGazer: A Divide-and-Conquer Approach for Monitoring and Optimizing Cloud-Based Networks. In Proceedings of the IEEE Pacific Visualization Symposium (PacificVis '15), pages 175–182. IEEE, 2015.

Contact
holger.stitz@jku.at
pipes-vs-dams.at

Acknowledgements
This work was supported by the Austrian Research Promotion Agency (FFG) via the PIPES-VS-DAMES project (840232).

FFG RISC Software GmbH KT-NET mowis