


Synthesizing Adaptive Layout Engines



Nate Yazdani and Ras Bodik


In the age of big data, humans need new tools to visually and interactively explore their ever-growing data sets.

Adaptive layout engines will enable this by adjusting runtime behavior based on available resources and properties of the problem at hand.




SANDCAT

ex: *treemap visualizations*



500 data points

versus



95,000 data points

Using a declarative domain-specific language, visualization designers specify layout problems as *attribute grammars*.

With a fine-tuned constraint model, an SMT solver takes less than a minute (rather than 8+ hours) to find valid *traversal schedules*, which compile to layout engines.

```

...
13 interface H VBox {
14     var width : int;
15     var height : int;
16     var right : int;
17     var bottom : int;
18 }
19 class VBox : H VBox {
20     children {
21         upper : H VBox;
22         lower : H VBox;
23     }
24     actions {
25         self.height := lower.height + upper.height;
26         self.width := max(lower.width, upper.width);
27         upper.right := self.right;
28         lower.right := self.right;
29         lower.bottom := self.bottom;
30         upper.bottom := self.bottom + lower.height;
31     }
32 }
...

```

ex: *sequential schedule*

```
post{height, width} ;; pre{bottom, right}
```

ex: *parallel schedule*

```
(post{height};;pre{bottom}) || (post{width};;pre{right})
```

Adaptation happens both...

Ahead of time

By augmenting synthesis with a *cost model*, the search can find a set of schedules, each optimal under different circumstances.

Just in time

Given an input tree, a *cost predictor* will dispatch to one of the layout engines based on properties such as tree size, shape, number of cores, cache size, etc.

