

Apache™ FOP Development: Implementation Overview

Following a Document Through Apache# FOP

Version 1298724

by Arved Sandstrom

Table of contents

1 Overview.....	2
2 Startup.....	2
3 Formatting Object Tree.....	2
4 Layout.....	3
5 Area Tree.....	3
6 Rendering.....	3

The purpose of this document is to tie together the Apache™ FOP design (interface) with some of the key points where control is passed within FOP (implementation), so that developers can quickly find the section of code that is relevant to their needs. The process described is for a "typical" command-line document. All classes are in org.apache.fop unless otherwise designated.

1 Overview

The input FO document is sent to the FO tree builder via SAX events. Fragments of an FO Tree are built from this process. As each page-sequence element is completed, it is passed to a layout processor, which in turn converts it into an Area Tree. The Area Tree is then given to the Renderer, which converts it into a stream of data containing the output document. The sections below will provide additional details. Where needed differences between the trunk and maintenance branches are shown in tabular format.

2 Startup

- The job starts in *apps.Fop.main()*.
- Control is passed to *apps.CommandLineStarter.run()*.
- Control is passed to *apps.Driver.render()*. This class fires up a SAX parser, the events from which indirectly control the remaining processing, including building the FO Tree, building the Area Tree, rendering, output and logging.

3 Formatting Object Tree

Trunk	Maintenance
The SAX events that the parser creates are handled by <i>fo.FOTreeBuilder</i> , which uses <code>startElement()</code> , <code>endElement()</code> , and <code>characters()</code> methods to build the FO Tree.	
<code>fo.FOTreeBuilder.endElement()</code> runs the <code>end()</code> method for each node as it is created. The <i>fo.pagination.PageSequence</i> class overrides this <code>end()</code> method to run <code>apps.LayoutHandler.endPageSequence()</code> , which in turn runs <code>fo.pagination.PageSequence.format()</code> .	the end of a <code>PageSequence</code> element causes the <code>PageSequence</code> object to be passed to <code>apps.StreamRenderer.render()</code> , which in turn runs <code>fo.pagination.PageSequence.format()</code> .
<code>fo.pagination.PageSequence.format()</code> creates a <i>layoutmgr.PageLayoutManager</i> , passing the <code>AreaTree</code> and <code>PageSequence</code> objects to it, then calls its <code>run()</code> method.	<code>fo.pagination.PageSequence.addFlow()</code> programatically adds a <code>Flow</code> object to the page sequence.
.	<code>fo.pagination.PageSequence.makePage()</code> creates a <code>BodyArea</code> and passes it to <i>fo.Flow.layout</i>
.	the layout process is then driven from <code>fo.pagination.PageSequence.format()</code> .

4 Layout

There are layout managers for each type of layout decision. They take an FO Tree as input and build a laid-out Area Tree from it. The layout process involves finding out where line breaks and page breaks should be made, then creating the areas on the page. Static areas can then be added for any static regions. As pages are completed, they are added to the Area Tree.

5 Area Tree

The area tree is a data structure designed to hold the page areas. These pages are then filled with the page regions and various areas. The area tree is used primarily as a minimal structure that can be rendered by the renderers.

The area tree is supported by an area tree model. This model handles the adding of pages to the area tree. It also handles page sequence starts, document level extensions, id references and unresolved id areas. This model allows the pages to be handled directly by a renderer or to store the pages for later use.

6 Rendering

The renderer receives pages from the area tree and renders those pages. If a renderer supports out of order rendering then it will either render or prepare a page in the correct order. Otherwise the pages are rendered in order. The task of the renderer is to take the pages and output them to the requested type. In the case of the AWTRenderer it needs to be able to view any page.

When rendering a page it takes the page and renders each page region. The main work for a renderer implementation is to handle the viewports and inline areas. The inline areas need to be drawn on the page in the correct place.