

3DPDF: Open Source Solutions for Incorporating 3D Information in PDF Files

Norman A. Graf

Abstract—Experimental science is replete with multi-dimensional information which is often poorly represented by the two dimensions of presentation slides and print media. Past efforts to disseminate such information to a wider audience have failed for a number of reasons, including a lack of standards which are easy to implement and have broad support. Adobe's Portable Document Format (PDF) has in recent years become the de facto standard for secure, dependable electronic information exchange. It has done so by creating an open format, providing support for multiple platforms and being reliable and extensible. By providing support for the ECMA standard Universal 3D (U3D) and the ISO PRC file format in its free Adobe Reader software, Adobe has made it easy to distribute and interact with 3D content. Until recently, Adobe's Acrobat software was also capable of incorporating 3D content into PDF files from a variety of 3D file formats, including proprietary CAD formats. However, this functionality is no longer available in Acrobat X, having been spun off to a separate company. Incorporating 3D content now requires the additional purchase of a separate plug-in. In this talk we present alternatives based on open source libraries which allow the programmatic creation of 3D content in PDF format. While not providing the same level of access to CAD files as the commercial software, it does provide physicists with an alternative path to incorporate 3D content into PDF files from such disparate applications as detector geometries from Geant4, 3D data sets, mathematical surfaces or tessellated volumes.

I. MOVING BEYOND PRINT MEDIA

PUBLICATION is undergoing a revolution due to the 21st Century multimedia explosion. In academia, education, and in the public in general, the trend is towards enhanced media. Digital publication of research journals, such as these IEEE Conference Proceedings, is becoming the standard. Rich internet applications such as animations, movies and interactive demonstrations are becoming the norm. Yet, despite the fact that PDF files are capable of supporting fully interactive 3D spatial and temporal content, the use of those features within the scientific community in general, and HEP in particular, has been almost non-existent. To-date, the utilization of 3D capabilities in PDF has been largely relegated to engineering applications incorporating CAD designs. The purpose of this presentation is to encourage wider use of these capabilities within the scientific community.

II. ADOBE PDF

Adobe's Portable Document Format (PDF) has in recent years become the *de facto* standard for secure, dependable

Manuscript received November 19, 2012. This work was supported in part by the U.S. Department of Energy.

N. A. Graf is with the SLAC National Accelerator Laboratory, Menlo Park, CA 94025 USA (telephone: 650-926-5317, e-mail: Norman.Graf@slac.stanford.edu).

electronic information exchange. Adobe has succeeded in making PDF a standard by creating an open, reliable and extensible format, supporting multiple operating systems and computing platforms, and, perhaps most importantly, by providing a free program to open and interact with PDF documents. By providing support for the ECMA standard U3D and the ISO standard PRC file formats in its free Adobe Reader software [3], Adobe has made it easy to distribute and interact with 3D content. By providing support for scripting and animation, temporal data can also be easily distributed to a wide audience.

III. U3D AND PRC: YET ANOTHER "KILLER 3D APP"?

The Universal 3D (U3D) format [1] is an ECMA standard originally designed by a consortium of businesses led by Intel. It aims to "simplify the transformation of complex 3D data into a format that can be streamed, compressed and viewed on affordable, nonproprietary software/hardware platforms while providing a high quality viewing experience". PRC (Product Representation Compact) [2] is an ISO standard capable of storing 3D product structure, geometry, and product manufacturing information. The geometry can be stored in either a tessellated or precise boundary representation (BREP) form and can be stored in a very compressed format. The goal of providing easy access to 3D content is not new, there have been many proposals in the past, such as VRML, and current proposals under development, such as X3D. For a number of reasons, none have managed to become widely accepted outside their specific application domain. What sets U3D and PRC apart is that both formats are fully supported within the PDF standard and access to the 3D content is provided by Adobe's free Reader software.

IV. INTERACTING WITH 3D CONTENT IN PDF

Figure 1 shows a number of ways to present a 3D object in 2D images. The author must usually decide on the style and viewpoint which best illustrates the point he is trying to make with the image. Often, one image does not suffice, or the reader is left wondering what the rest of the object looks like.

By directly incorporating the full 3D information into these proceedings, as in Figure 2, the reader is free to interact with the object directly, choosing not only the viewing style, but also the viewing angle. By incorporating the geometrical hierarchy, the reader can either highlight certain features, or make them invisible.

Simply by hovering the mouse over the image, the navigation toolbar will become visible. Or the reader may simply

*Presented at IEEE 2012 Nuclear Science Symposium, Medical Imaging Conference
Anaheim, California, October 29 - November 3, 2012*

Work supported by US Department of Energy contract DE-AC02-76SF00515.

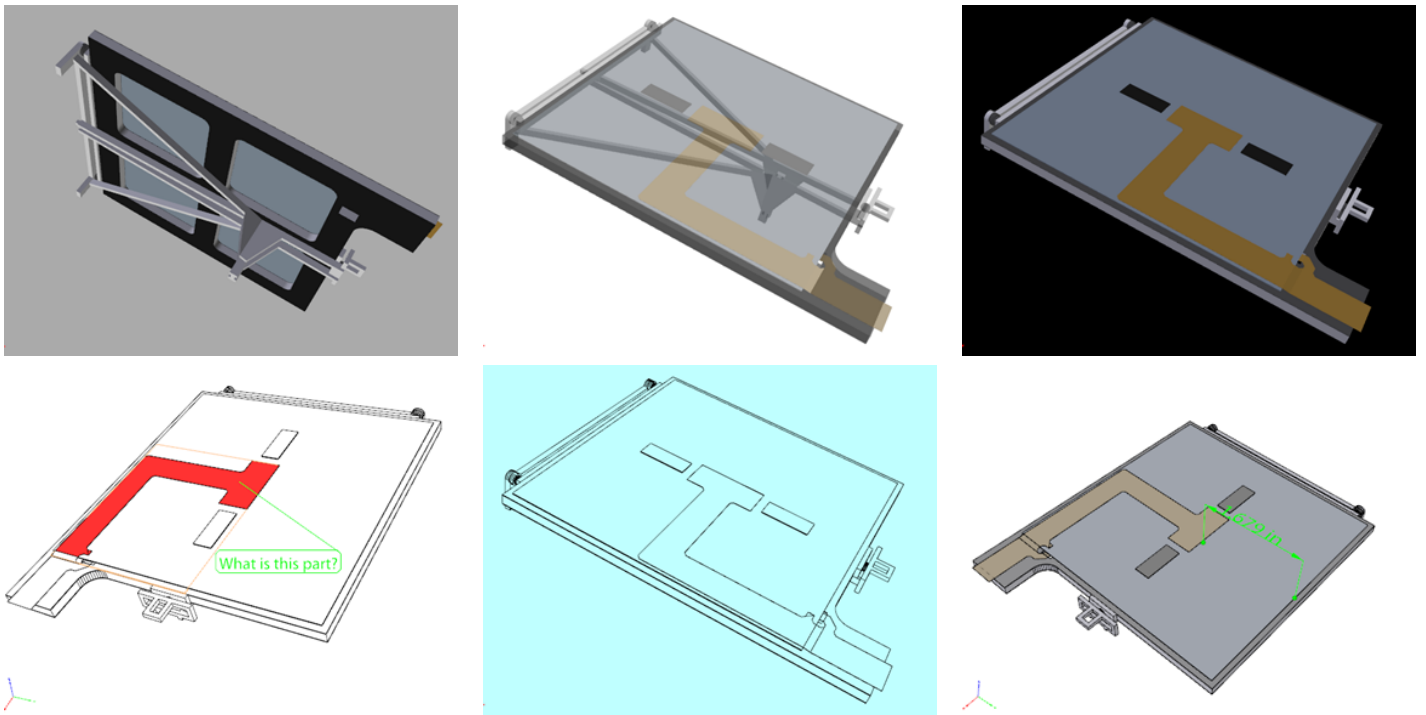


Fig. 1. A number of viewing styles and orientations for an object.




Fig. 2. The object, embedded as a fully interactive 3D model.




Fig. 3. 3D map of a dipole field.

right-click in the image to bring up a context-sensitive menu of available operations. One can rotate, zoom or translate using either the mouse (left button, right button, both buttons, respectively) or by selecting menu options. The Viewing Options menu allows the model rendering modes to be selected, the lighting to be specified, the background color to be changed, or even cross sections of the model to be calculated and rendered.

It should be noted, however, that to-date Adobe's Reader or Acrobat software are the only programs which correctly render the 3D content. However, as Adobe Reader is freely available for a multitude of operating systems, this should not be an issue.

This functionality is not, of course, restricted to the inclusion of CAD drawings. One can include mathematical functions, field maps or other 3D plots or surfaces, such as the 3D map of a dipole field in Figure 3. Although the file contains 40,000 points and 80,000 faces, the performance is excellent and the file size remains manageable.

One can also include 3D representations of structures which have been elucidated using imaging radiation detectors. For example, Figure 4 contains the structure of a portion of the Steap3 protein that helps red blood cells acquire iron. This structure was determined using the Stanford Synchrotron Radiation Lightsource (SSRL) at the SLAC National Accelerator Laboratory [7]. Dedicated programs to manipulate and interact with such 3D biological and chemical models clearly provide much more functionality to the scientists directly involved in such molecular research, but being able to bring some of this 3D functionality to the general public provides a much richer experience than any static 2D image.

Although not described in this document, one can also include temporally varying information in the form of animations within PDF documents. This can be as simple as an

assembly or disassembly sequence for a detector element, to something as complicated as protein folding. Interaction with the embedded objects can also be customized using Javascript, but a full description of that is also beyond the scope of this writeup.

In the following sections we will describe how to actually embed 3D content into your own PDF files.

V. EMBEDDING 3D CONTENT INTO PDF USING ADOBE ACROBAT

The easiest way to import 3D objects into PDF is to use Adobe Acrobat combined with the 3DPDF Converter plugin developed by tetra4D [5] in a post-processing step. One simply creates the PDF document in the usual way, whether with Microsoft Word, LaTeX, or any other word processor capable of producing PDF, including static 2D images as has been the norm for print media. One then opens this PDF file in Acrobat (not the free Adobe Reader) and inserts the 3D objects directly into the document. Over 50 third-party file formats are fully supported for import, from such commercial CAD formats as AutoCAD, CADDSS, CATIA, I-DEAS, Parasolid, Pro/ENGINEER, Solid Edge and Solid Works, to CAD exchange formats such as 3Dxml, IGES, STEP and STL, to open 3D formats such as VRML and U3D. In addition, many CAD software packages support native PDF export using customizable templates. Defaults are provided for all of the supported formats, but the user is able to fully customize the import process, specifying everything from activation settings (*i.e.* when the 3D content is activated in the viewer), background color, lighting and default views. One can also associate JavaScript to the object to provide enhanced functionality. The level of detail of the imported object is also user-definable, from highly compressed models appropriate for




Fig. 4. Structure of a portion of the Steap3 protein that helps red blood cells acquire iron.

web-quality resolution all the way to CAD/CAM-exportable fidelity. The default conversion settings are designed to balance the quality and size of the PDF document. Anyone with the free Adobe Reader can now interact with the 3D content in your PDF file.

VI. EMBEDDING 3D CONTENT INTO PDF USING L^AT_EX

It is also possible to include 3D objects into PDF directly as part of the normal workflow using L^AT_EX. This requires the use of the media9 [6] package. Invoke the package by putting the line

```
\usepackage[] {media9}
```

in the preamble of the document. Content can be embedded using the

```
\includemedia[<options>]{<poster text>}
    {<main 3D file (PRC | U3D)>}
```

command. The documentation which ships with the media9 package is quite complete and includes a number of examples, so we will only briefly highlight some of the command options relevant to 3D objects here. The options include:

activate

onclick, pageopen or pagevisible : determines when the 3D content becomes active. When inactive, only the content defined by <poster text> is visible.

3Dtoolbar

if set, the 3D toolbar is visible when content is activated.

3Dnavpane

if set, the model hierarchy navigation pane is visible when content is activated.

3Dbg

<r> <g> rgb background color with values from 0 to 1.

3Dlights

sets the lighting on the model. Default is that defined in the 3D model. Options include ‘CAD’ and ‘Headlight’, among others.

3Drender

sets the default rendering style. Options are Solid, SolidWireframe, Transparent, TransparentWireframe, BoundingBox, TransparentBoundingBox, TransparentBoundingBoxOutline, Wireframe, ShadedWireframe, HiddenWireframe, Vertices, ShadedVertices, SolidOutline, Illustration, ShadedIllustration. Some styles are shown in Figure 1.

3Dviews

This allows predefined views to be included in the model. Views may also contain a cross-section of the object. Details on how to generate and store views can be found in the media9 documentation.

add3Dscript

Associates a JavaScript file with the 3D object. Upon activation the scripts are executed in the order in which they are added.

These are followed by either text, or a graphics file to be used as the 2D poster image when the 3D object is not activated. Finally, one specifies the 3D model, in either U3D or PRC format. One can also create predefined views of the object to guide the reader through an exploration of the model.

For instance, the dipole field shown in Figure 3 was created using the following snippet of code:

```
\begin{figure}[htb]
\centering
\includemedia[width=.5\textwidth,
             height=.5\textwidth,
             activate=onclick,
             3Dlights=Headlamp]{}
             {dipoleField.u3d}
\caption{3D map of a dipole field.}
\label{fig:dipole}
\end{figure}
```

However, it should be noted that these simply define the default settings which are in place when the 3D content is activated. The user can invoke a context-sensitive menu by right-clicking in the image. These menus can be used to customize the content.

VII. CREATING 3D CONTENT

For 3D content originating in CAD programs, one can either include the proprietary formats into pdf using the commercial software described above, or export the model to U3D or PRC format and use the media9 package in L^AT_EX. If, however, one is generating 3D content outside a CAD environment, such as displaying surfaces, mathematical functions, field maps, 3D histograms, "Lego" plots, etc. one is faced with the need to create U3D or PRC files directly. The VCG library [8] is a portable C++ templated library for the manipulation and processing of triangle and tetrahedral meshes. It is able to import meshes in a variety of formats and has implemented a U3D exporter. Meshlab [9] is an open source, portable, and extensible system for the processing and editing of unstructured 3D triangular meshes which is heavily based on the VCG library. Both projects provide direct conversion of 3D meshes into the U3D format, which can then be imported into a final PDF file. The author is currently engaged in developing an open-source project to facilitate the creation of U3D and PRC files from a number of programs used within the HEP community.

VIII. SUMMARY

Publication is undergoing a revolution due to the 21st Century multimedia explosion. In academia, education, and in public, the trend is towards enhanced media, with rich internet applications such as animations, movies and interactive demonstrations becoming the norm. The digital publication of research journals is also becoming standard, but the vast majority of publications continue to follow the 2D norm of print publications. To-date, the utilization of the 3D capabilities of PDF has been largely relegated to engineering applications incorporating CAD designs. I would like to encourage wider

use of this functionality within the HEP community. The ability to embed 3D objects and enable the time evolution of information in standard PDF documents opens up whole new dimensions of information transfer. Users have the ability to interact with the content, customize the view, and obtain more information than is possible using static 2D images.

REFERENCES

- [1] <http://www.ecma-international.org/publications/standards/Ecma-363.htm>
- [2] http://www.aiim.org/documents/standards/PDF-E/PRC_FAQ_final.pdf
- [3] <http://www.adobe.com/products/reader/>
- [4] <http://www.adobe.com/products/acrobat.html>
- [5] <http://www.tetra4d.com/3d-pdf-converter.html>
- [6] <http://www.ctan.org/tex-archive/macros/latex/contrib/media9>
- [7] http://www-ssrl.slac.stanford.edu/research/highlights_archive/steap3.html
- [8] <http://vcg.sourceforge.net>
- [9] <http://meshlab.sourceforge.net/>