

Rendering Mathematics for the Web using Madoko

Daan Leijen
Microsoft Research
daan@microsoft.com

1. INTRODUCTION

Madoko [6–8] is a novel authoring system for writing complex documents. It is especially well suited for complex academic or industrial documents, like scientific articles, reference manuals, or math-heavy presentations. One particularly important aspect of Madoko is to write a document in high-level Markdown [5] with a focus on semantic content. From this document specification we can generate both high-quality PDF output (via L^AT_EX) but also generate high-quality HTML that can re-scale and re-flow dynamically. Styling is done through standard CSS attributes and can be done orthogonal to the content.

Madoko provides extensive support for mathematics rendering. All math is rendered using L^AT_EX with full compatibility with any L^AT_EX packages and commands. Rendering to PDF comes this way for free but a high quality rendering of the math in the resulting HTML is more involved. This application note article describes in detail how Madoko deals with the various technical challenges. Moreover we show how other mechanisms, like replacement rules, help with creating mini domain-specific extensions to cleanly express complex math.

Since this article is about the rendering of math to HTML, it is highly recommended to read this article as an HTML page instead of PDF! It can be found at <http://tinyurl.com/madokomath>.

2. AN OVERVIEW OF MADOKO

Madoko is based on *Markdown* [5] as its input format. The main design goal is to enable light-weight creation of high-quality scholarly and technical documents for the web and print, while maintaining John Gruber’s Markdown philosophy of simplicity and focus on plain text readability. Since the Markdown input format is well-structured, this allows Madoko to generate both high quality HTML and PDF (through L^AT_EX and B^IB_TE_X). There has been a lot of effort in Madoko to make the L^AT_EX generation robust and cus-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

DocEng '16, September 12-16, 2016, Vienna, Austria

© 2016 ACM. ISBN 978-1-4503-4438-8/16/09...\$15.00

DOI: <http://dx.doi.org/10.1145/2960811.2967168>

tomizable while integrating well with the various academic document- and bibliography styles.

On modern devices like tablets and phones it is generally much more pleasant to read a paper or technical document as HTML instead of PDF since HTML can scale and reflow dynamically. A study by Franze et al. [2] showed that the most desired features when reading papers is being able to change the font size, alter margins, or have a single column layout; all of these are trivial in a web browser. Of course, this article itself was written in Madoko, and the HTML version can be viewed at <http://tinyurl.com/madokomath>. Others have tried to create re-scalable and re-flowable content from paginated PDF [11], or the other way around, paginating dynamic content [3], but we believe starting from a more high-level structured input format is a better way of approaching this problem.

The move to Markdown makes the the documents *structured*, *readable*, and output *independent*. The final ingredients that Madoko adds are to make the documents *styleable* through standard CSS rules, and *programmable* through transformation rules. These additions also makes it easy to add custom domain specific document elements, like *exercise* or *answer*, that can be transformed, numbered, and styled in a declarative manner.

Finally, the online version at madoko.net integrates seamlessly with Dropbox, GitHub, and OneDrive, making documents available anywhere on any device. Madoko synchronizes automatically and multiple authors can work concurrently on the same document using robust three-way merges on concurrent updates. This means that updates by others are not quite real-time as in other collaborative environments (although they are performed frequently), but anyone can now work off-line and still reliably merge when connecting again. Madoko.net is itself a HTML5 web application and the editor continues to work in the browser even when offline. Of course, you can always use the plain command line version of Madoko locally (`npm install -g madoko`).

3. SCALABLE MATH ON THE WEB

Madoko uses regular L^AT_EX for describing math formulas since T_EX is still the gold standard for rendering and describing mathematics. Any formula can be directly embedded in a Madoko document. For example:

```
A famous formula is  $e^{i\pi} + 1 = 0$ , but the
following one is also well-known:
~ Equation { #eq-gaussian }
 $\int_{-\infty}^{\infty} e^{-a x^2} dx$ 
```

```
= \sqrt{\frac{\pi}{a}}
```

A famous formula is $e^{i\pi} + 1 = 0$, but the following one is also well-known:

$$\int_{-\infty}^{\infty} e^{-ax^2} dx = \sqrt{\frac{\pi}{a}} \quad (1)$$

Here we use `$` to start inline math as in \LaTeX . For the equation we used a so-called *custom block* of Madoko. The standard prelude of Madoko defines `~ equation` for numbered equations, `~math` for plain display math, and `~mathpre` for pre-formatted math discussed in Section 4. In the example, we also give the equation a name so we can refer to it using links in Markdown where `[#eq-gaussian]` expands to the equation number, e.g. Equation (1).

When creating PDF output, Madoko can simply include the literal formula in the generated \LaTeX with full compatibility with any \LaTeX package. Unfortunately, for HTML output the process is more involved as we need to render math separate from the rest of the HTML. There exist various tools that use JavaScript to interpret \LaTeX math commands directly and generate a rendering on the client. One of the most well-known libraries to do that is MathJax [1].

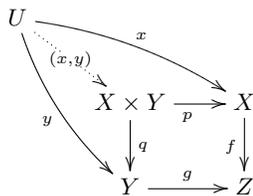
This works well for simple mathematics but one of the great advantages of using \LaTeX for math is that it comes with many many packages to render advanced mathematics, or packages that render math in the style required by a journal. Such packages are generally not supported by tools like MathJax. Even though Madoko has an option to use MathJax for dynamic math rendering, the default mode is to invoke \LaTeX and render all math at compile time.

In Madoko any \LaTeX package can be used through a metadata key at the start of the document, for example:

```
Package: [curve]xypic
```

after which we can use the `\xymatrix` command to render category theory diagrams:

```
~ Math
\xymatrix @-0.5em{
U \ar@/_/[ddr]_y \ar@/^/[drr]^x \ar@{.}>[dr] |-(x,y) \\\
& X \times Y \ar[d]^q \ar[r]_p & X \ar[d]_f \\\
& Y \ar[r]_g & Z
```



Generating good looking HTML from \LaTeX rendered formulas is a challenging problem though and we describe here various solutions adopted by Madoko.

Hashing of formulas

When rendering a document, Madoko first collects all math formulas and assigns a unique MD5 hash to each individual formula. This ensures that each formula is only rendered

it follows that

$$e^{i\pi} = -1 + 0i,$$

which yields Euler's identity:

$$e^{i\pi} + 1 = 0.$$

it follows that

$$e^{i\pi} = -1 + 0i,$$

which yields Euler's identity:

$$e^{i\pi} + 1 = 0.$$

Figure 1. Screenshots of different math renderings in the browser: the left image is rendered by Madoko using SVG graphics, while the right image is a rendering by Wikipedia using PNG images (https://en.wikipedia.org/wiki/Euler%27s_identity)

once which is important since many short formulas are usually often repeated. Madoko generates a special \LaTeX math file that contains 'snippet' entries for each formula. For example, for this document, one of the entries is:

```
%mdk-data-line={138}
\begin{mdInlineSnippet}[f2d2e607c3e99d5c34bc0aad01893a0d]
\math{i\pi} + 1 = 0%
\end{mdInlineSnippet}
```

The initial comment is how Madoko maps back \LaTeX error messages to the correct line in the original Madoko file – this is very important in practice to quickly solve \LaTeX problems. Next, the `mdInlineSnippet` command ensures that each formula gets rendered on its own page in the resulting DVI file. That DVI file is now passed to another tool to extract the rendering.

Madoko uses the excellent `dvivsgm` converter by Martin Gieseke [4] to convert \LaTeX generated DVI files to *scalable vector graphics* (SVG) files. The `dvivsgm` converter automatically extracts an SVG file for each page in the DVI file, numbering them sequentially. Since Madoko maintains a mapping between the MD5 hashes and the page numbers, it can then automatically include the correct SVG images in the generated HTML for each formula.

Many tools extract math formulas as PNG images from a rendered PDF or Postscript file. Unfortunately, this is a non-scalable image and looks generally quite fuzzy on a screen especially for inline formulas surround by text. Figure 1 compares the rendering of Euler's identity in Wikipedia, which uses PNG images, versus the rendering in Madoko which uses SVG. The difference is quite stark and the quality of SVG rendering is excellent even compared to PDF – when demoing Madoko, often people are under the impression of viewing PDF while they are actually seeing the HTML rendering of a Madoko document in the browser.

Baseline alignment

There are still various technical hurdles to overcome though. The most tricky one is proper baseline alignment. In particular, an inline formula should align as $\sum_{i=0}^{\infty} e^i$ with the e aligned with the text baseline. Contrast this with $\sum_{i=0}^{\infty} e^i$ for example where the bottom of the extracted image aligns with the baseline. There are often a lot of small inline formulas and not aligned well with the baseline looks very irregular to the eyes.

To achieve proper baseline alignment, we need to have an exact measurement of the *depth* of the formula, i.e. the bottom vertical distance to the baseline. If we know the depth, we can adjust the vertical alignment of the extracted image by lowering it by its depth.

The `mdInlineSnippet` environment does this by first rendering the formula in a \TeX box. This box can be queried for its rendered height, width, and depth. After figuring out the dimensions the box is rendered to the page. For each formula, we write out the measured dimensions together with its hash (which is an argument to `mdInlineSnippet`) to a separate text file. After the \LaTeX run, Madoko reads this dimension file to determine the precise baseline alignment for each formula in the HTML.

The final height of the math image should be determined by the relative font size of the surrounding text used in the HTML. This means that the height and baseline adjustment must be made in font-relative `em` units instead of absolute units. Madoko renders mathematics in a 10pt font size when taking measurements. For output to HTML we read the measurements from the dimension file (in `pt`) and divide by 10 to get the relative `em` units. We use the CSS `vertical-align` attribute to lower the math image by its measured depth. In practice we also scale the math image by 105% in order to look more natural with most web fonts. For example, our initial example, $\sum_{i=0}^{\infty} e^i$ is positioned in the HTML output as:

```
<svg style="vertical-align:-0.3502em;height:1.2355em"
  viewBox="88.467 53.397 33.929 11.767"
  class="math-inline math-render-svg math">
<desc>\sum_{i=0}^{\infty}e^i</desc>
<g id="page26">
  <use x="88.667" y="54.364" xlink:href="#g14-80"
    xmlns:xlink="http://www.w3.org/1999/xlink"></use>
  ...
</g></svg>
```

The `<use>` element puts the glyph `#g14-80` (the Σ) at a specific position. That glyph is defined separately to enable sharing of graphical elements between different formulas.

Sharing glyph paths

Math heavy documents can easily contain thousands of formulas. Madoko already shares representations for equal formulas through hashing but more is needed. For example, in one example math-heavy article [10] the math formulas generate 2242kb of SVG images. It turns out though that many formulas contain similar glyphs, like e , or x . Each of these glyphs is (usually) rendered as a `path` in the SVG image. For example, the formula x is described in SVG as:

```
<defs>
  <path d="M3.328 -3.009C3.387 -3.268 3.616 -4.184 ...
    -0.986 2.879 -1.205 2.989 -1.644L3.328 -3.009Z"
    id="g6-120"></path></defs>
<g><use x="88.667" y="61.836" xlink:href="#g6-120"
  xmlns:xlink="http://www.w3.org/1999/xlink">
  </use></g>
```

Here we see how the image places the path element `#g6-120` at a specific position using an `xlink`. The path element though just traces a specific glyph, in this case the x , independently of its position. As the shapes are independent of the position, we can share all the glyph paths between different formulas. Madoko will collect all equivalent paths in a separate definition block and all formulas reference these shared paths. This can lead to significant space savings in practice – in the example article the space usage went down 79% from 2242kb to 467kb.

More significant space savings can be made by not describ-

$$\int_{-\infty}^{\infty} e^{-ax^2} dx = \sqrt{\frac{\pi}{a}} \quad \text{vs.} \quad \int_{-\infty}^{\infty} e^{-ax^2} dx = \sqrt{\frac{\pi}{a}}$$

Figure 2. A browser screenshot of two math SVG images generated by Madoko. The left image used SVG path elements to trace glyphs, while the right image uses direct font elements.

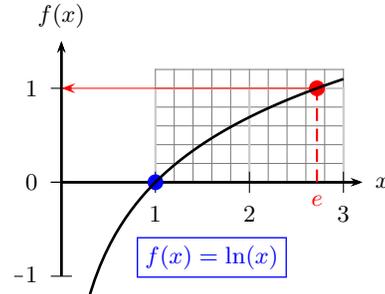


Figure 3. SVG image of a log graph rendered using the `pstricks` and `pst-plot` packages which emit Postscript specials.

ing glyphs with paths at all but using direct font entries and text elements in the SVG description. Unfortunately, font support in SVG is very spotty across browsers and most formulas do not render faithfully when using fonts directly¹. Figure 2 shows two browser screenshots where one formula is rendered using traces while the other uses font elements.

Rendering of DVI specials

Some \LaTeX commands depend on specific output drivers. For example, the advanced `TikZ` package draws vector graphics using specific PDF primitives which are not directly supported in DVI files. In many cases, we can still extract correct SVG images from the DVI since `dvissvgm` supports many extensions. Ultimately, if that fails Madoko can also generate PNG files from a PDF or Postscript rendering although such image will no longer be scalable. However, currently even large packages like `pstricks` and `TikZ` work with DVI output so in practice this is almost never necessary. Figure 3 shows the SVG output of a log graph using the `pstricks` package which issues Postscript specials.

4. PRE-FORMATTED MATHEMATICS

Mathematics mode in \TeX can be surprising in its handling of whitespace and identifiers. In general, whitespace in the text is not relevant and a sequence of letters is *not* seen as a single identifier. Look for example at the following formula:

```
$function sqrt( x : int)$
~>
functionsqrt(x : int)
```

We can see that there is no whitespace between `function` and `sqrt`, and how they are rendered as a sequence of letters instead of two identifiers; note in particular the whitespace between the `f` and `u` for example.

This behavior may be good for general mathematics, but in many fields, like computer science, this is often cumber-

¹Madoko supports this option though through a metadata flag.

