

The pbalance (poor man’s balance) package*

Nelson Lago
lago@ime.usp.br

<https://gitlab.com/lago/pbalance>

December 17, 2020

Abstract

This package tries to *safely* make the columns on the last page of a two-column document have approximately the same height. It should “just work” without user intervention, which is particularly useful for class authors, but also offers the user a command to adjust the height of the columns. There are, however, three caveats:

1. Results are adequate, but often not optimal;
2. In some very rare cases, the package may give up (the document is generated correctly but the last page is not balanced);
3. The package demands additional \LaTeX passes.

1 Introduction

First things first: this package is a **hack**. It is also **beta** quality. However, I believe it is *safe*, i.e., it should not generate incorrect output.

In a two-column document, it is desirable for the columns in the last page to have the same height (the columns should be “balanced”). Two packages (`balance` and `flushend`) try to provide this in \LaTeX , but both may generate defective output in some circumstances. This package tries to solve this problem.

Balancing is often not possible: imagine a document with an odd number of lines on the last page, or with ten lines of text together with an image that takes the height of fifteen lines of text, or a document where making columns the same height would leave a sectioning command at the bottom of the first column etc. Since such cases happen often, we consider it enough for the columns to be “not too unequal”, even when equal heights might be possible (the algorithm cannot deliver better results and, in fact, striving to make columns exactly the same height with it may sometimes yield poor results, so we do not even try).

*This document corresponds to pbalance v1.0.1, dated 2020/12/16.

2 Usage

To use the package, add `\usepackage{pbalance}` to the preamble and things should “just work”. However, the document will take more \LaTeX passes to compile, and minute changes to it that would normally need one additional \LaTeX pass will likely demand three. If you load the package with the `draft` option (or pass the `draft` option to the document class), balancing is disabled, alleviating the need for extra passes during document preparation.

If the last page does not have floats, footnotes, or marginpars, the package simply uses `\balance`, from the `balance` package; otherwise it uses its own algorithm. You may instead use `\shrinkLastPage{some-measurement}` to manually define how much shorter the first column of the last page should be. This, however, will not eliminate the need for extra passes (the package needs to find out which page is the last).

If there are many floats near the end of the document, the \LaTeX float placement parameters (`\topfraction`, `topnumber`, etc.) may interfere negatively with the algorithm. In some rare cases¹, this may force the package to give up on balancing (the document is still generated correctly in this case). To minimize the chances of this outcome, it is probably a good idea to make the parameters a little more liberal than the default, for example:

```
\setcounter{totalnumber}{5} % default 3
\setcounter{topnumber}{3} % default 2
\renewcommand{\topfraction}{.85} % default .7
\renewcommand{\textfraction}{.15} % default .2
\renewcommand{\floatpagefraction}{.75} % default .5, must be < \topfraction
```

3 How does it work

This package works on two fronts:

1. It uses a \LaTeX pass to gather information: which is the last page? Does it have floats? Footnotes? How much free space in each column?
2. It uses this information in a subsequent pass to balance the last page. If there are no floats or footnotes, it uses the `balance` package; otherwise it uses the measurements collected to shrink the first column (using `\enlargethispage`) by a “reasonable” amount.

If measurements change between passes, either the document was edited or \LaTeX is still adding crossrefs, citations, etc. In both cases, the previously gathered data is bogus,

¹For example, imagine a document in which page 5 is a float page that happens to be the last one. We will balance page 4, which is the last text page. Balancing page 4 may move a float from the left to the right column, but this may in turn exceed the maximum number of floats in the right column. As a result, one or more floats may be deferred to page 5. At the same time, a float of a different kind from page 5 may be pulled back to the extra space that just opened up in page 4. Depending on the size of this float, the text may no longer fit in page 4, making some text spill onto page 5. This means page 5 is now the last text page and should then be balanced instead of page 4. Oops!

so we start over. However, the package cannot detect changes that happen between the “collect data” pass and the “balance for the first time” pass: it has to assume that any change at this point is caused by the fact that we started balancing (it can, however, detect whether the last text page changed; if this happens, we give up balancing). If a small change (such as an updated crossref adding or removing a line somewhere) slips through, that is acceptable: we are not aiming at balancing perfectly anyway. However, if the change is big (for example, the addition of the bibliography block after running bibtex/biber), results will suffer.

To prevent this, it would be ideal to only start the process after all passes needed to stabilize the document have run. In practice, however, this would be too hard to detect (and demand even more passes). What we do instead is simply wait for two consecutive passes to result in columns of the same size in the last page. When this happens, we proceed to balance in the next pass and assume (quite reasonably) that, as long as the document is not modified, the effect of concurrent changes during that pass in the last page is small, so ignoring it disturbs the balancing only slightly².

4 TODO / wish list

Some useful stuff we should consider implementing:

- add `\if@twocolumn` tests to prevent crashes and other problems;
- add options to (1) allow using the balance package on pages with floats and (2) completely disable the balance package;
- make the mechanism able to work with any page, not just the last, and with more than one page at once. This is useful for the last page of each chapter of a two-column book, for example;
- on top of that:
 1. create command to manually indicate the page numbers to balance
 2. create command similar to `\balance` from the `balance` package
 3. allow user to set manual balancing for each page independently

Change History

v1.0 General: Initial version 1	v1.0.1 General: Improvements to documentation 1
--	---

²Like I said on the introduction: this package is a hack.