BRNO UNIVERSITY OF TECHNOLOGY
FACULTY OF INFORMATION TECHNOLOGY

VYSOKÉ UČENÍ TECHNICKÉ V BRNĚ
FAKULTA INFORMAČNÍCH TECHNOLOGIÍ

DEPARTMENT OF COMPUTER GRAPHICS AND MULTIMEDIA
ÚSTAV POČÍTAČOVÉ GRAFIKY A MULTIMÉDIÍ

ACCELERATED SPARSE MATRIX OPERATIONS IN NONLINEAR LEAST SQUARES SOLVERS

AKCELERACE OPERACÍ NAD ŘÍDKÝMI MATICEMI V NELINEÁRNÍ METODĚ NEJMENŠÍCH ČTVERCŮ

PH.D. THESIS
DISERTAČNÍ PRÁCE

AUTHOR | ING. YOUR NAME
AUTOR PRÁCE

SUPERVISOR | PROF. YOUR PROF NAME, PH.D.
VEDOUCÍ PRÁCE

BRNO 2016 – version 1.0

Ing. Your Name: *Accelerated Sparse Matrix Operations in Nonlinear Least Squares Solvers*, doctoral thesis Brno, Brno University of Technology, Faculty of Information Technology, 2016.

I'd like to dedicate this thesis to my parents who raised
me even despite I am a difficult teenage ingrate.

Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetuer id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.

ABSTRACT

Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetuer id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.

Nam dui ligula, fringilla a, euismod sodales, sollicitudin vel, wisi. Morbi auctor lorem non justo. Nam lacus libero, pretium at, lobortis vitae, ultricies et, tellus. Donec aliquet, tortor sed accumsan bibendum, erat ligula aliquet magna, vitae ornare odio metus a mi. Morbi ac orci et nisl hendrerit mollis. Suspendisse ut massa. Cras nec ante. Pellentesque a nulla. Cum sociis natoque penatibus et magnis dis parturient montes, nascetur ridiculus mus. Aliquam tincidunt urna. Nulla ullamcorper vestibulum turpis. Pellentesque cursus luctus mauris.

Nulla malesuada porttitor diam. Donec felis erat, congue non, volutpat at, tincidunt tristique, libero. Vivamus viverra fermentum felis. Donec nonummy pellentesque ante. Phasellus adipiscing semper elit. Proin fermentum massa ac quam. Sed diam turpis, molestie vitae, placerat a, molestie nec, leo. Maecenas lacinia. Nam ipsum ligula, eleifend at, accumsan nec, suscipit a, ipsum. Morbi blandit ligula feugiat magna. Nunc eleifend consequat lorem. Sed lacinia nulla vitae enim. Pellentesque tincidunt purus vel magna. Integer non enim. Praesent euismod nunc eu purus. Donec bibendum quam in tellus. Nullam cursus pulvinar lectus. Donec et mi. Nam vulputate metus eu enim. Vestibulum pellentesque felis eu massa.

KEYWORDS

Nonlinear least squares; numerical methods; sparse block matrix; GPGPU.

ABSTRAKT

[Here goes abstract in your native language] Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetuer id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.

Nam dui ligula, fringilla a, euismod sodales, sollicitudin vel, wisi. Morbi auctor lorem non justo. Nam lacus libero, pretium at, lobortis vitae, ultricies et, tellus. Donec aliquet, tortor sed accumsan bibendum, erat ligula aliquet magna, vitae ornare odio metus a mi. Morbi ac orci et nisl hendrerit mollis. Suspendisse ut massa. Cras nec ante. Pellentesque a nulla. Cum sociis natoque penatibus et magnis dis parturient montes, nascetur ridiculus mus. Aliquam tincidunt urna. Nulla ullamcorper vestibulum turpis. Pellentesque cursus luctus mauris.

Nulla malesuada porttitor diam. Donec felis erat, congue non, volutpat at, tincidunt tristique, libero. Vivamus viverra fermentum felis. Donec nonummy pellentesque ante. Phasellus adipiscing semper elit. Proin fermentum massa ac quam. Sed diam turpis, molestie vitae, placerat a, molestie nec, leo. Maecenas lacinia. Nam ipsum ligula, eleifend at, accumsan nec, suscipit a, ipsum. Morbi blandit ligula feugiat magna. Nunc eleifend consequat lorem. Sed lacinia nulla vitae enim. Pellentesque tincidunt purus vel magna. Integer non enim. Praesent euismod nunc eu purus. Donec bibendum quam in tellus. Nullam cursus pulvinar lectus. Donec et mi. Nam vulputate metus eu enim. Vestibulum pellentesque felis eu massa.

KLÍČOVÁ SLOVA

Nelineární metoda nejmenších čtverců; numerické metody; řídké blokové matice; GPGPU.

BIBLIOGRAPHIC CITATION

Ing. Your Name: *Accelerated Sparse Matrix Operations in Nonlinear Least Squares Solvers*, doctoral thesis Brno, Brno University of Technology, Faculty of Information Technology, 2016.

DECLARATION

I declare that this dissertation thesis is my original work and that I have written it under the guidance of Prof. your Prof Name, Ph.D.. All sources and literature that I have used during my work on the thesis are correctly cited with complete reference to the respective sources.

Brno, 2016

Ing. Your Name, November 14, 2016

ACKNOWLEDGMENTS

Many thanks to my supervisor Pavel Smrž, who was being patient and supportive throughout my whole study and who got me on research projects where I was given free rein and could develop my ideas. I would also like to thank my colleagues at the Graph group, especially Viorela Ila, Pavel Zemčík, Adam Herout, Marek Šolony and Pavel Svoboda for helpful and motivational comments, proofreading as well as lots of help with implementation and benchmarking.

Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetuer id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.

Last but not least, I'd like to thank the guys who kept the department's coffee machine going. God knows I would not be able to do anything without it.

This is derived from `classicthesis`. I hereby apologize to its authors, who did a great job preparing this beautiful template, for hacking it and publishing my fork in a total ignorance of their efforts and of the consequences. The name of this modified template (`classicfeces`) is supposed to point to this fact and discourage everyone from using it.

PUBLICATIONS

Some ideas and figures have appeared previously in the following publications:

[PIS13a] Lukáš Polok, Viorela Ila, and Pavel Smrž. Cache efficient implementation for block matrix operations. In *Proc. of the High Performance Computing Symp.*, pages 698–706. ACM, 2013.

[PIŠ+13b] Lukáš Polok, Viorela Ila, Marek Šolony, Pavel Smrž, and Pavel Zemčík. Incremental block Cholesky factorization for nonlinear least squares in robotics. In *Robotics: Science and Systems (RSS)*, 2013.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF LISTINGS

LIST OF ALGORITHMS

LIST OF ACRONYMS

1 INTRODUCTION: TEMPLATE TUTORIAL

This is a modified classicthesis template, with countless improvements added[1]. There are two ways to go from here. Either you take this template where everything is pretty much set up and sort of do the things my way. If you don't like that, you are probably better off just grabbing a blank copy of the original classicthesis and go from there.

If you go with classicfeces (this template), you probably want to cruise over to classicthesis-config.tex and tweak things. To disable the annoying "draft" message, just remove the drafting option in the first \PassOptionsToPackage command. If, on the other hand, you want to make it bigger to remind you every day, remove the unobtrusivedraft option. Feel free to explore other options as well.

Under that, there is the title of the thesis, your name, your supervisor name and so on. Change that and it will change throughout the text. When you are ready to print, set printingmonochrome to true. Otherwise you will have light blue links in the text, which seems nicer for reading on screen but would be uncool on paper. In case you have very complex plots or vector figures, you may want to use enabledraftvectorgraphicsrasterization (affects only the draft mode). You can also check out Ela's math macros, around line 390 in classicthesis-config.tex. For the sake of your professor, who's going to read countless versions of these dissertations, please rename the output at the first line of main.tex to include your name (or just rename the pdf using a script if that doesn't work), to make his/her work easier.

We can use *emphasis* in *opportune places* – and for a good reason. We can also use acronyms such as Central Processing Unit (CPU) or in plural form CPUs, or explicitly short CPU / CPUs and full Central Processing Unit (CPU) / Central Processing Units

The original template also lets you write behind the margins. I haven't actually used that for anything, I prefer the oldschool footnotes. This is pretty narrow.

[1] As far as I remember, the PDF TOC was broken, the lists of figures / tables / algorithms / listings were inconsistent and broken, the less than and greater than symbols were slanted which I didn't like, the page margins were broken at times, there was no support for acronyms and setting up own citations was pure hell. The lists of citations were not aligned nicely (the left margin was varying, according to the number of digits in the [123] brackets). I have added the indispensable `% !TeX root = ../main.tex` at the beginning of each file so that you can build any open document correctly in TeXworks (otherwise it has no idea that all those files belong to the same document). There were also issues with the page headings in the "List of ..." chapters. The original classicthesis is built entirely without bold fonts; this template does not follow that completely, the Listings and Algorithms are typeset with bold in them (if you are a typographic nazi, you should hack \spacedlowcaps{} in there somewhere).

(CPUs). You shouls also cite papers [3], making sure to use non-break space. You can also mention the authors of the papers in case it makes sense, e.g. Goto and Van De Geijn [6]. Look, a footnote[2]!

This template is set up so that you can refer to your own publications in a different style using the `\citeown{something}` command, e.g. [PIŠ+13b]. This will appear at the beginning of the manuscript, in the list of publications. If you want to make sure that all your publications are there, just list them all in a `\nociteown{something}` command (can be anywhere in the text, maybe in the chapter that refers to those papers) – that will put them on the list, but otherwise will not actually display anything.

Here are some more tips for writing (it would have been easier if you could adopt those when you published your first paper). Some of those are a matter of personal taste and the list is by no means ordered by priority (it is ordered by time I discovered the particular shortcomings[3]). You can actually keep those in your manuscript and comment those out as you go through them:

- Check page headings, especially in the list of figures, tables, algorithms, acronyms, your own publications (if you have it) and so on (those are set using the `\markboth` command).
- Mention that on Xeons (or other CPUs that have it), the benchmarks are with turbo boost off (if you keep it on, then the results depend on weather – don't do that).
- Learn writing semantic LaTeX. Use e.g. `\emph` instead of `\textit`, define `\vec` and `\mat` with styles for vectors and matrices instead of putting `\mathbf` everywhere, ...
- Remove "best viewed in color", you will be printing the thesis with colors and the pdf will be with colors. right?
- Get rid of `\{ \}` in equations where it appears in place of `( )` (except when defining sets perhaps).
- Unify how Cholmod is spelled.
- Unify how you use `\parallel`, `\Vert`, `\lVert`, `\rVert` and `\|`.
- Unify how you use `\mid`, `\vert`, `\lvert`, `\rvert` and `|`.
- Replace `|` as binary op by `\mid` (that should add the missing operator spaces).
- Decide how to write CPU-GPU and CPU+GPU (with or without spaces?, short dash?).
- An algorithms, there is `\Return`, no need for `\textbf{return}`.

[2] I'm pretty.

[3] Or when preparing the manuscript for print ... so there are a few things high on the list, which should be done last.

- Multiple equations on top of each other do have commas on the right, the last one has a comma or a full stop.
- Put `\mbox`es around inline equations and names such as e.g. "SLAM ++".
- Unify how the datasets are mentioned (in `\emph` or not).
- Decide whether comments in algorithms should be full sentences.
- Decide if require statements in algorithms should be full sentences.
- Add `\hspace{\arraycolsep}` around the second = on the row with `&=&` inside eqnarray or stop using eqnarray.
- Cholesky in own citations is spelled cholesky (lowercase), try to do something about it, HPC symposia <-> symposium.
- Use `\setminus` rather than `\;\backslash\;` for linear solving or better yet make a macro?
- Put labels of algorithms to the top rather than on the bottom, otherwise the line number references will point elsewhere.
- Include some minimal math notation chapter?
- Redraw figures using a single color scheme.
- Make a list of algorithmic pseudo-functions, make sure those are named consistently.
- Align table contents to the left / right, centering does not help in most cases.
- Replace `\hfill\hfill` by `\hfill\null` and see what happens (the `\null` prevents the right-most `\hfill` from being gobbled).
- Check your PDF bookmarks (PDF TOC).
- Zeros vs. zeroes.
- Replace all simple "ASCII parentheses" (") by "calligraphic ones" (`` ` ` `` and `''`)
- Make sure that wide floats overflow towards the edge of the page (the edge is the wider margin in this template)!
- Use `\autoref` (package hypperref) instead of `\ref` for tables, figures (not always for subfigures though), sections (possibly `\nameref` for sections or chapters where suitable) but use `\eqref` (package amsmath) for equations, keep `\ref` for algorithm lines. use e.g. `{\tableautorefname}s` to get the plural of "Tables" or better yet declare `\tableautorefnamep` in the intro in case someone changes `\tableautorefname` from "Table" to "Tab.".
- Unify hyphenation of element-wise, blockwise, pose-SLAM, re-linearization, re-ordering, re-*, worst-case, metaprogramming, nonzero, above mentioned, nonzeros, so called, prefetch, upper-triangular, lower-triangular, submatrix, forward substitution, backsubstitution [add your favorite hyphenations here].

- Add shorter captions for list of figures/algorithms/listings/tables (e.g. "Common research datasets, top line synthetic datasets, bottom line real datasets, sphere and garage are 3D pose sets." -> "Common research datasets"). See e.g. Listing 3.3 for how that's done.
- Translate Czech (or your native language) abstract.
- Use acronyms for matrix formats, i.e. `\ac{CSC}` rather than `\texttt{CSC}`, see if it is possible to use `\texttt` with `\ac`, maybe make a frankenmacro for that.
- Use macros for acronyms with explanations (such as `\acf{something}` rather than writing e.g. `central processing unit (\ac{CPU})`).
- Resolve fixed block size, constant block size, single block size.
- "et al." is correctly written "et al." (only one full stop, not both), should be done by the use of `\citet`.
- CSC vs. CCS.
- Replace all `^T`, `^{-1}` and so on by Ela's math mactos (`\tr` and `^T` give different font in the "T" symbol. `\tr` actually uses a symbol other than the letter "T" – compare $R^\top$, $R^\top$, R^T, R^T and $R^\mathbf{T}$).
- Can use `\DeclareMathOperator{\foo}{foo}` rather than `\text{foo}`.
- Replace `{\bf x}` and `{\mathbf x}` by `\mathbf{x}` (extra tricky, as `{\mathbf abc}` will only typeset `a` in bold and `bc` in regular!).
- Use `\hat` for updated variables instead of `\tilde`.
- Gather acronyms (look for acronyms which are not inside `\ac{}`).
- Expand all "Fig." to "Figure", we want to have rather more text than less (mostly done by the use of `\autoref`).
- Capitalize titles correctly.
- `\nocite` (or rather `\nociteown`) own papers before the beginning of each respective section that uses them.
- Replace "..." ellipsis by `\ldots` or Unicode, can use the ellipsis package.
- Differentiate between text comma and math comma, e.g., type "for `$x=a$, $b$, or~$c$`" instead of "for `$x=a,b,\text{ or }c$`".
- Use `\centering` instead of `\begin{center}` / `\end{center}` as it allegedly takes less vertical space.
- Use `\begin{align}` instead of `\begin{eqnarray}` as eqnarray is not supported by amsmath; this requires. *single* ampersands *before* the assignment (&= rather than &=&).
- Put a tilde before refs or citations, e.g.,`the technique in~\cite{something}`.
- Make sure that all the tables have captions on top (or bottom if thats your thing).

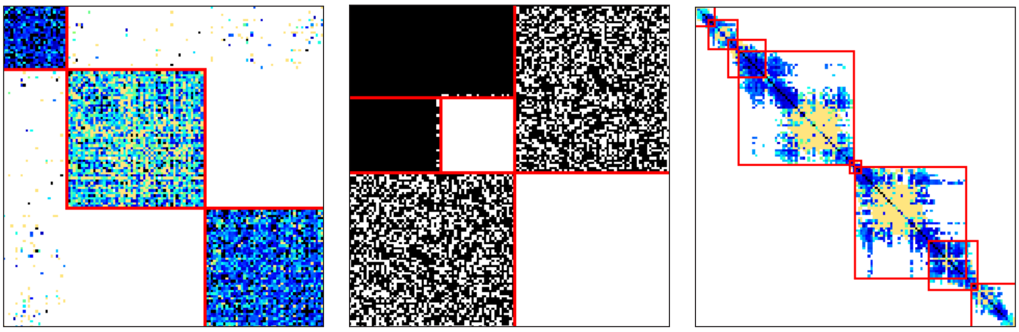

(a) (b) (c)

Figure 1.1: Examples of approximately block matrices from the University of Florida Sparse Matrix Collection [2], specifically in the DIMACS10 dataset [13], a) an approximate block matrix with scattered nonzero elements, b) a block matrix with unaligned blocks, and in the Oberwolfach dataset [11], c) an approximate block matrix with overlapping blocks. (The block boundaries are only suggested – not a part of the original matrices.)

- Remove "in this paper" and rather use "in this chapter", "part" or "section", mostly taken care of by `\autoref`.
- Take care of words like part, section and chapter.
- Use "`\citet{Author2016something}`" rather than "`Author et al. \cite{Author2016something}`" where possible, so as to get the authors' names right (works with natbib, numbers and plainnat).
- See about the `raggedbottom` vs. `flushbottom`. The latter looks nicer but it is more elaborate to achieve without gaps between the paragraphs (see the comments at the beginning of `main.tex`).
- Decide whether matrices are capital bold or not.
- Make sure all the comments in algorithms use the same style (not italics, sentences).
- Get rid of duplicate own papers if your bibliography files are messy.

This chapter should later become the Introduction chapter of your thesis.

1.1 FOCUS OF THE THESIS

You should remember to cite bunch of papers [14], even in groups [4, 5, 12, 10, 14, 15]. This is not surprising though. Let's look at figures.

The figures and tables are best referred using the `\autoref` command, such as Figure 1.1a or Figure 1.1c. Several figures in one should be typeset using `\subfloat`.

Quisque ullamcorper placerat ipsum. Cras nibh. Morbi vel justo vitae lacus tincidunt ultrices. Lorem ipsum dolor sit amet, consectetuer adipiscing elit. In hac habitasse platea dictumst. Integer tempus convallis augue. Etiam facilisis. Nunc elementum fermentum wisi. Aenean placerat. Ut imperdiet, enim sed gravida sollicitudin, felis odio placerat quam, ac pulvinar elit purus eget enim. Nunc vitae tortor. Proin tempus nibh sit amet nisl. Vivamus quis tortor vitae risus porta vehicula.

1.2 contributions

Fusce mauris. Vestibulum luctus nibh at lectus. Sed bibendum, nulla a faucibus semper, leo velit ultricies tellus, ac venenatis arcu wisi vel nisl. Vestibulum diam. Aliquam pellentesque, augue quis sagittis posuere, turpis lacus congue quam, in hendrerit risus eros eget felis. Maecenas eget erat in sapien mattis porttitor. Vestibulum porttitor. Nulla facilisi. Sed a turpis eu lacus commodo facilisis. Morbi fringilla, wisi in dignissim interdum, justo lectus sagittis dui, et vehicula libero dui cursus dui. Mauris tempor ligula sed lacus. Duis cursus enim ut augue. Cras ac magna. Cras nulla. Nulla egestas. Curabitur a leo. Quisque egestas wisi eget nunc. Nam feugiat lacus vel est. Curabitur consectetuer.

1.3 structure of the thesis

The next chapter serves a brief introduction into the NLS problems and their applications. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Donec odio elit, dictum in, hendrerit sit amet, egestas sed, leo. Praesent feugiat sapien aliquet odio. Integer vitae justo. Aliquam vestibulum fringilla lorem. Sed neque lectus, consectetuer at, consectetuer sed, eleifend ac, lectus. Nulla facilisi. Pellentesque eget lectus. Proin eu metus. Sed porttitor. In hac habitasse platea dictumst. Suspendisse eu lectus. Ut mi mi, lacinia sit amet, placerat et, mollis vitae, dui. Sed ante tellus, tristique ut, iaculis eu, malesuada ac, dui. Mauris nibh leo, facilisis non, adipiscing quis, ultrices a, dui.

Suspendisse vel felis. Ut lorem lorem, interdum eu, tincidunt sit amet, laoreet vitae, arcu. Aenean faucibus pede eu ante. Praesent enim elit, rutrum at, molestie non, nonummy vel, nisl. Ut lectus eros, malesuada sit amet, fermentum eu, sodales cursus, magna. Donec eu purus. Quisque vehicula, urna sed ultricies auctor, pede lorem egestas dui, et convallis elit erat sed nulla. Donec luctus. Curabitur et nunc. Aliquam dolor odio, commodo pretium, ultricies non, pharetra in, velit. Integer arcu est, nonummy in, fermentum faucibus, egestas vel, odio.

1.4 math notations

This text makes use of various more or less standard mathematical notations. Make sure to typeset equations using `\begin{equation}`. Alternatively, use `\begin{align}` to insert more equations in a single block, aligned at the assignments (or at any suitable position, really):

$$\begin{aligned} \Lambda &= \begin{pmatrix} \Lambda_{11} & \Lambda_{12} \\ \Lambda_{21} & \Lambda_{22} \end{pmatrix} , \\ R &= \mathrm{chol}(\Lambda) . \end{aligned}$$

Note that I like to include punctuation in the equations as above but that's not a must. Make sure to *avoid at all costs* the `eqnarray` environment, as it introduces a number of irregularities! Make sure to `\mbox` inline maths so it does not break across more lines. Note that I don't do that here as I'm lazy. Would have done that if I knew from the start.

This section briefly revises the used conventions. Vectors are denoted by small bold letters, e.g. v is a scalar but $\mathbf{v}$ is a vector. Matrices are denoted by capital Latin or Greek letters, e.g. A and Λ are matrices. Matrices (and as a special case also vectors) which are logically partitioned to blocks are denoted by *bold* capital Latin or Greek letters, e.g. Σ is a matrix but $\mathbf{\Sigma}$ is a matrix where the elements are matrices.

To assemble a (column) vector, one writes $\mathbf{c} = [1; 2; 3]$ while to assemble a row vector, one writes $\mathbf{r} = [1, 2, 3]$. By an extension, writing $A = [1, 2; 3, 4]$ is the same as:

$$A = \begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix} .$$

Part I

BACKGROUND

Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetuer id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.

2

LEAST SQUARES METHODS AND THEIR APPLICATIONS

Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetuer id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.

2.1 NONLINEAR LEAST SQUARES

Let's have some fancy math. Note that below, we label the equations using `\label{eq:something}` and refer to them using `\eqref{eq:something}`, e.g. (2.1). That automatically adds the appropriate brackets. The joint probability distribution can be written as:

$$P(\boldsymbol{\theta}, \mathbf{z}) \propto P(\theta_0) \prod_{z}^{n} P\left(z_k \mid \theta_{i_k}, \theta_{j_k}\right), \tag{2.1}$$

where $P(\theta_0)$ is the prior and z_k are the constraints between the variables θ_{i_k} and θ_{j_k}. The goal is to obtain the Maximum Likelihood Estimation (MLE) of a set of variables in θ, given the available observations in $\mathbf{z}$:

$$\theta^* = \underset{\theta}{\operatorname{argmax}}\ P(\theta \mid \mathbf{z}) = \underset{\theta}{\operatorname{argmin}}\ \left(-\log(P(\theta \mid \mathbf{z}))\right). \tag{2.2}$$

For every measurement $z_k = h_k\left(\theta_{i_k}, \theta_{j_k}\right) - \nu_k$, the noise ν_k is assumed to be normally distributed, with zero mean and covariance Σ_k:

$$P\left(z_k \mid \theta_{i_k}, \theta_{j_k}\right) \propto \exp\left(-\frac{1}{2}\left\|h_k\left(\theta_{i_k}, \theta_{j_k}\right) \ominus z_k\right\|_{\Sigma_k}^2\right), \tag{2.3}$$

where $h_k\left(\theta_{i_k}, \theta_{j_k}\right)$ is the nonlinear measurement function, z_k are the measurements, $\ominus$ is the vectorial inverse composition operator. Note that binary measurements are assumed here but measurements of any degree can be combined at will. Setting $\boldsymbol{\Sigma} = I$ yields ordinary nonlinear least squares, otherwise weighted

Nonlinear Least Squares (NLS) are obtained. Finding the MLE from (2.2) is done by solving the following NLS problem:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \left(\frac{1}{2} \sum_{k=1}^{m} \left\| h_k \left(\theta_{i_k}, \theta_{j_k} \right) \ominus z_k \right\|_{\Sigma_k}^2 \right) . \tag{2.4}$$

Gathering all residuals in $\mathbf{r}(\theta) = [r_1, \ \ldots, \ r_m]^\top$ where $r_k = h_k \left(\theta_{i_k}, \theta_{j_k} \right) \ominus z_k$ and gathering the measurement noise in $\boldsymbol{\Sigma} = \operatorname{diag} \left([\Sigma_1, \ \ldots, \ \Sigma_m] \right)$, the sum in (2.4) can be written in the vectorial form and expressed in terms of a L2-norm:

$$\left\| \mathbf{r} (\theta) \right\|_{\boldsymbol{\Sigma}}^2 = \mathbf{r}^\top (\theta) \ \boldsymbol{\Sigma}^{-1} \mathbf{r} (\theta) = \left\| \boldsymbol{\Sigma}^{-\top/2} \mathbf{r} (\theta) \right\|^2 . \tag{2.5}$$

Table 2.1: A few of the commonly used robust functions. Note that a, b and c are constant parameters of the individual functions (i.e. not the same variable).

	loss function $\rho(u)$	score function $\psi(u) = \frac{\partial \rho(u)}{\partial u}$
Ordinary LS	$\frac{1}{2}u^2$	u
Huber [9]	$\begin{cases} \frac{1}{2}u^2 & \text{if } \lvert u \rvert \le a \\ \frac{1}{2}a(2\lvert u \rvert - a) & \textit{otherwise} \end{cases}$	$\begin{cases} 1 & \text{if } \lvert u \rvert \le a \\ a \operatorname{sign}(u) & \textit{otherwise} \end{cases}$
Cauchy [8]	$\frac{a^2}{2} \log\left(1 + \left(\frac{u}{a}\right)^2\right)$	$\frac{u}{1+\left(\frac{u}{a}\right)^2}$
Tukey [1]	$\begin{cases} \frac{a^2}{6}\left(1 - \left(1 - \left(\frac{u}{a}\right)^2\right)^3\right) & \text{if } \lvert u \rvert \le a \\ \frac{a^2}{6} & \textit{otherwise} \end{cases}$	$\begin{cases} u\left(1 - \left(\frac{u}{a}\right)^2\right)^2 & \text{if } \lvert u \rvert \le a \\ 0 & \textit{otherwise} \end{cases}$
Hampel [7]	$\begin{cases} \frac{1}{2}u^2 & \text{if } \lvert u \rvert < a \\ a\lvert u \rvert - \frac{1}{2}a^2 & \text{if } a \le \lvert u \rvert < b \\ a\frac{c\lvert u \rvert - \frac{1}{2}u^2}{c-b} - \frac{7}{6}a^2 & \text{if } b \le \lvert u \rvert < c \\ a(b+c-a) & \textit{otherwise} \end{cases}$	$\begin{cases} u & \text{if } \lvert u \rvert < a \\ a \operatorname{sign}(u) & \text{if } a \le \lvert u \rvert < b \\ a\frac{c \operatorname{sign}(u) - u}{c-b} & \text{if } b \le \lvert u \rvert < c \\ 0 & \textit{otherwise} \end{cases}$

2.1.1 *Dealing with Outlier Measurements*

We can also have tables. We also use `\autoref` to refer to tables, e.g. Table 2.1.

3 SPARSE MATRIX REPRESENTATIONS

Let's look at listings and algorithms.

3.1 COORDINATE FORMAT

The coordinate format, often abbreviated `COO`, is a very simple sparse matrix format; it stores each nonzero entry as a triplet of *row*, *column* and the associated *value*, with no ordering imposed by a rule or at least a convention. It is suitable for assembly of the sparse matrices and it is simple to add or erase more values at any time. For the matrix in Figure 3.1a, the `COO` representation is in Listing 3.1.

Listing 3.1: The matrix from Figure 3.1a stored in the `COO` format.

```
m = 6; n = 4; nnz = 12;
row = {0, 0, 1, 2, 2, 3, 3, 3, 4, 4, 5, 5};
col = {0, 3, 0, 2, 3, 0, 1, 3, 0, 3, 1, 3};
val = {6, 4, 7, 9, 4, 2, 5, 3, 2, 1, 1, 2};
```

To better illustrate the properties of this format, let us consider a simple matrix vector product of the form $\mathbf{y} = A\mathbf{x} + \mathbf{y}$, listed in Algorithm 3.1. Looking at the source code of this algorithm, two things quickly become apparent. We can refer to the individual lines of the algorithm by using `\label{algline:something}` and then `\ref{algline:something}`. For example, the loop starts at line 2 and all the hard work is done at line 5, The clickable hyperlinks will not work, unless the label of the algorithm itself is on the top (otherwise they will wrongly point to

$$\begin{bmatrix} 6 & & & 4 \\ 7 & & & \\ & & 9 & 4 \\ 2 & 5 & & 3 \\ 2 & & & 1 \\ & 1 & & 2 \end{bmatrix}$$

(a)

$$\begin{bmatrix} 6 & & 5 & \\ & 1 & 7 & \\ & & 9 & 4 \\ & & & 3 \end{bmatrix}$$

(b)

Figure 3.1: Example sparse matrices.

Algorithm 3.1: A basic matrix - vector multiplication algorithm.

function COO_GEMV(A, $\mathbf{x}$, $\mathbf{y}$)

  ▷ calculates $\mathbf{y} = A\mathbf{x} + \mathbf{y}$ where $\mathbf{x}$, $\mathbf{y}$ are dense and A is sparse in `COO` format

  for $i = 0$ **to** $A.nnz - 1$ **do**

    $r = A.row[i]$

    $c = A.col[i]$

    $\mathbf{y}[r] = \mathbf{y}[r] + A.val[i] \cdot \mathbf{x}[c]$

  end for

end function

the previous algorithm). Note that this is probably the *only* reasonable place to use `\ref`, otherwise you should use `\autoref` or `\eqref`.

3.2 MORE LISTINGS

You can put C++ code inside listings, easily:

Listing 3.2: Example of a C++ code (the `std::mismatch` function).

```
template<class InputIt1, class InputIt2>
std::pair<InputIt1, InputIt2>
    mismatch(InputIt1 first1, InputIt1 last1, InputIt2 first2)
{
    while (first1 != last1 && *first1 == *first2)
        ++first1, ++first2; // awkward comma
    return std::make_pair(first1, first2);
}
```

Note that you can also include math in the listings, by specifying `mathescape=true` in the listing option, Listing 3.3 uses `$\star$` to display the $\star$ character. Note that these listings (and in fact some tables and figures as well) have a different title in the list of listings and a different – and usually longer – title displayed above them. This is done using `caption={[Short title]Long title}`.

Listing 3.3: The matrix from Figure 3.1a stored in the `ELL` format.

```
m = 6; n = 4; ndiag = 3;
coef = {6, 7, 9, 2, 2, 1,
        4, 0, 4, 5, 1, 2,
        0, 0, 0, 3, 0, 0}; // Note the filled-in zeros in shorter rows.
col  = {0, 0, 2, 0, 0, 1,  // The ⋆ entries do not correspond to any
        3, ⋆, 3, 1, 3, 3,  // non-zero value and can point to arbitrary
        ⋆, ⋆, ⋆, 3, ⋆, ⋆}; // column (e.g. the last one - 3).
```

3.3 chapter summary

Fusce mauris. Vestibulum luctus nibh at lectus. Sed bibendum, nulla a faucibus semper, leo velit ultricies tellus, ac venenatis arcu wisi vel nisl. Vestibulum diam. Aliquam pellentesque, augue quis sagittis posuere, turpis lacus congue quam, in hendrerit risus eros eget felis. Maecenas eget erat in sapien mattis porttitor. Vestibulum porttitor. Nulla facilisi. Sed a turpis eu lacus commodo facilisis. Morbi fringilla, wisi in dignissim interdum, justo lectus sagittis dui, et vehicula libero dui cursus dui. Mauris tempor ligula sed lacus. Duis cursus enim ut augue. Cras ac magna. Cras nulla. Nulla egestas. Curabitur a leo. Quisque egestas wisi eget nunc. Nam feugiat lacus vel est. Curabitur consectetuer.

Suspendisse vel felis. Ut lorem lorem, interdum eu, tincidunt sit amet, laoreet vitae, arcu. Aenean faucibus pede eu ante. Praesent enim elit, rutrum at, molestie non, nonummy vel, nisl. Ut lectus eros, malesuada sit amet, fermentum eu, sodales cursus, magna. Donec eu purus. Quisque vehicula, urna sed ultricies auctor, pede lorem egestas dui, et convallis elit erat sed nulla. Donec luctus. Curabitur et nunc. Aliquam dolor odio, commodo pretium, ultricies non, pharetra in, velit. Integer arcu est, nonummy in, fermentum faucibus, egestas vel, odio.

The common design goals in sparse block storage are:

dense blocks: the blocks are stored as dense matrices so that they can be easily written by the Jacobian function rather than scattered into an elementwise format.

immutable addresses: in an incremental setting, the Jacobian matrix is augmented with new blocks as new observations come and it can double as a cache if the addresses of the existing blocks do not change.

integral representation: the nonlinear solvers often employ direct methods and matrix factorizations which are not compatible with split matrix schemes.

efficient arithmetics: the current solvers, much to their disadvantage, only use sparse block matrices as a convenient platform for generation of elementwise sparse matrices to be passed to the linear solver.

Part II

SECOND PART

Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetuer id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.

4 SLAM ++ BLOCK MATRIX DESIGN

Many applications ranging from physics, computer graphics, computer vision to robotics rely on efficiently solving large nonlinear systems of equations, as illustrated in the previous chapter. In the case of using a Gauss-Newton-like algorithm, the solution can be approximated by iteratively solving a series of linearized problems. In some applications, the size of the system can be considerably large. The most computationally demanding part is to assemble and solve the linearized system at each iteration. This chapter shows solutions that exploit both, the block structure and the sparsity of the corresponding matrices and offers very efficient methods to manipulate, store and perform arithmetic operations.

Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetuer id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.

Part III

APPENDIX

Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Ut purus elit, vestibulum ut, placerat ac, adipiscing vitae, felis. Curabitur dictum gravida mauris. Nam arcu libero, nonummy eget, consectetuer id, vulputate a, magna. Donec vehicula augue eu neque. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Mauris ut leo. Cras viverra metus rhoncus sem. Nulla et lectus vestibulum urna fringilla ultrices. Phasellus eu tellus sit amet tortor gravida placerat. Integer sapien est, iaculis in, pretium quis, viverra ac, nunc. Praesent eget sem vel leo ultrices bibendum. Aenean faucibus. Morbi dolor nulla, malesuada eu, pulvinar at, mollis ac, nulla. Curabitur auctor semper nulla. Donec varius orci eget risus. Duis nibh mi, congue eu, accumsan eleifend, sagittis quis, diam. Duis eget orci sit amet orci dignissim rutrum.

A APPENDIX

The list of publications cited in the text follows. Note that special care was taken so that all the publications are justified and correctly aligned. This was not particularly easy to do – the solution used here is a bit hackish and may require a bit of user intervention in `Bibliography.tex` or `Publication.tex` if the space for the bracket on the left is too big or too small (shouldn't really happen though, unless you exceed 1000 references or use a different citation style).

BIBLIOGRAPHY

[1] Albert E Beaton and John W Tukey. The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data. *ASQ/ASA Technometrics*, 16(2):147–185, 1974.

[2] T.A. Davis. The university of florida sparse matrix collection. In *NA Digest*. Citeseer, 1994.

[3] Ulrich Drepper. What every programmer should know about memory. Technical report, Red Hat, Inc., 2007.

[4] I. S. Duff and J. K. Reid. An implementation of Tarjan's algorithm for the block triangularization of a matrix. *ACM Trans. Math. Software*, 4(2):137–147, 1978. ISSN 0098–3500. doi: http://doi.acm.org/10.1145/355780.355785.

[5] Iain S Duff. On algorithms for obtaining a maximum transversal. *ACM Trans. Math. Software*, 7(3):315–330, 1981.

[6] Kazushige Goto and Robert Van De Geijn. High-performance implementation of the level-3 BLAS. *ACM Trans. Math. Software*, 35(1):4, 2008.

[7] Frank R Hampel. Robust estimation: A condensed partial survey. *Probability Theory and Related Fields*, 27(2):87–104, 1973.

[8] Paul W Holland and Roy E Welsch. Robust regression using iteratively reweighted least-squares. *Communications in Statistics – Theory and Methods*, 6(9):813–827, 1977.

[9] Peter J Huber. *Robust statistics*. Springer Heidelberg, 2011. ISBN 978-0-470-12990-6.

[10] Eun-Jin Im and Katherine Yelick. Optimizing sparse matrix computations for register reuse in SPARSITY. In *Computational Science–ICCS 2001*, pages 127–136. Springer Heidelberg, 2001.

[11] Jan Lienemann, Dag Billger, Evgenii B Rudnyi, Andreas Greiner, and Jan G Korvink. Mems compact modeling meets model order reduction: Examples of the application of arnoldi methods to microsystem devices. In *Proceedings of the NSTI Nanotechnology Conference and Trade Show (Nanotech)*, volume 4, 2004.

[12] Alex Pothen and Chin-Ju Fan. Computing the block triangular form of a sparse matrix. *ACM Trans. Math. Software*, 16(4):303–324, 1990.

[13] Ilya Safro, Peter Sanders, and Christian Schulz. Advanced coarsening schemes for graph partitioning. *ACM J. of Experimental Algorithmics*, 19:2–2, 2015.

[14] Richard Vuduc, James W Demmel, Katherine Yelick, Shoaib Kamil, Rajesh Nishtala, Benjamin Lee, et al. Performance optimizations and bounds for sparse matrix-vector multiply. In *Proc. of the ACM/IEEE Conf. on Supercomputing*, pages 26–26. IEEE, 2002.

[15] Richard W Vuduc and Hyun-Jin Moon. Fast sparse matrix-vector multiplication by exploiting variable block structure. In *Proc. of the Intl. Conf. on High Performance Computing and Commun. (HPCC)*, pages 807–816. Springer Heidelberg, 2005.

COLOPHON

This document was typeset using the typographical look-and-feel `classicthesis` developed by André Miede. The style was inspired by Robert Bringhurst's seminal book on typography "*The Elements of Typographic Style*". It is available for LaTeX via CTAN as `classicthesis`.