

ADMB and TMB Developers' Workshop Report 2017

August 14-18, 2017

International Council for the Exploration of the Sea (ICES)

H.C. Andersens Boulevard 44-46, Copenhagen, Denmark

Contents

Description

Main Objectives

Schedule

Presentation Highlights

Breakout Group Topics and Achievements

ADMB Foundation Meeting

Participant List



Description

This report summarizes the events of the 7th ADMB Developers' Workshop, a meeting held at the International Council for the Exploration of the Sea (ICES) in Copenhagen, Denmark. The meeting was attended by many ADMB and TMB core developers who were joined by invited international experts.

The meeting was held in the typical informal style, allowing open group discussions and demonstrations of new features, possible improvements, and current issues. Meeting participants spent considerable time working in subgroups, each of which tackled a related set of priority issues for the ADMB project.

Main Objectives

- Get an overview of ongoing ADMB and TMB development, decide where work is needed
- Remove a newly discovered bug in likelihood profiling in ADMB
- Understand when we can be confident in the Laplace approximation and checks
- Improve documentation including websites, videos, and examples
- Test new MCMC functionality in ADMB and TMB, produce a set of examples
- Improve core functions in ADMB

Schedule

Monday

- 10am - Welcome and introductions
- 11am - Organize breakout groups and logistics for the week
- Noon - Lunch
- 1pm - Presentation: Johnnoel, "Recent updates to ADMB and overview of test suite"
- 2pm - Presentation: Helen Ogden, "Checking the quality of Laplace-approximate inference"
- 3:30pm - Breakout groups begin working
- 5pm - End

Tuesday

- 9am - Discuss goals for the week
- 10am - Presentation: Kasper Kristensen, "Recent updates to TMB"
- 11:30am - Presentation: Kasper Kristensen, "How to add user contributed C++ files"
- Noon - Lunch
- 1pm - Breakout groups work on objectives
- 4pm - Group boat tour of Copenhagen

Wednesday

- 9am - Presentation: Cole Monnahan, "How to use ADNUTS in ADMB and TMB"
- 11am - Presentation: Arni Magnusson, "ADMB and TMB codebase: Composition and history"
- Noon - Lunch
- 1pm - Breakout groups work on objectives
- 5:30pm - End

Thursday

- 9am - Breakout groups work on objectives
- Noon - Lunch
- 1pm - Breakout groups work on objectives
- 4pm - Foundation meeting
- 6:30pm - End

Friday

- 9am to Lunch - Review and edit Developers' Workshop Report 2017
- Noon - Lunch
- 1pm - Continue to review and edit Developers' Workshop Report 2017
- Breakout groups
- 5pm - Presentation: Arni Magnusson "ADMB and TMB downloads, then and now"
- Closing of the meeting

Presentation Highlights

Johnoel Ancheta, "Recent updates to ADMB"

- Testing coverage is at 56% as calculated by CodeCov
- Using Buildbot server to build and test on multiple operating system and compiler combinations
- Next goal is to investigate upcoming C++ 17 specifications for possible multicore support

Helen Ogden, "Checking the quality of Laplace-approximate inference"

- Laplace approximation tends to underestimate variance parameters
- 4 methods for approximating the likelihood:
 - Laplace (fastest)
 - AGQ - adaptive Gaussian quadrature (can't handle integral > 1 dimension)
 - SR - sequential reduction (new publication from Helen)
 - IS - importance sampling (asymptotically unbiased)
- They all use the normal approximation as a baseline and the non-Laplace ones add complexity and therefore take longer
- A "score" exists to approximate the KL distance between the true and approximate likelihoods

- Read the uploaded [presentation](#)

Kasper Kristensen, “Recent changes to TMB”

- The same version of Matrix package must be used to build and run TMB, otherwise gradients can be completely wrong - a warning is displayed if Matrix version does not match with TMB
- New functionality to simulate from a model from within the template, the documentation for this could be improved
- checkConsistency() is a new function to check the Laplace approximation
- Updates to the code are automatically checked by “Travis”
- Tweedie and Conway-Maxwell-Poisson distributions are available
- Windows users on government machines with limited rights need to install from CRAN

Cole Monnahan, “How to use the No U-Turn MCMC Sampler in ADMB and TMB”

- Cole gave a demonstration
- Speed comparisons were presented, but further verification is required

Arni Magnusson, “ADMB and TMB codebase: Composition and history”

- ADMB codebase is 284 thousand lines of code, TMB is 204 thousand lines of code
- ADMB codebase is growing at 5% per year, TMB at 9% per year
- TMB consists of 96% C++ and 4% R code
- Arni produced a [document](#) detailing the codebase information and the scripts used to build the tables

Arni Magnusson, “ADMB and TMB downloads, then and now”

- Downloads per month are around 280 for ADMB, up to 4500 for TMB
- ADMB downloads per month have increased over the last few years, TMB downloads are growing fast
- A recent surge in TMB downloads is closely related to increased downloads of glmmTMB
- Arni produced a [document](#) detailing the downloads information

Breakout Group Topics and Achievements

Bug in ADMB Likelihood Profiling - A likelihood profile bug was fixed and resolved by Anders, Chris and Johnnoel. At a recent stock assessment meeting a profile likelihood was produced, which looked strange. The profile likelihood did not look anything like the standard Hessian based normal-approximation. The profile likelihood increased and decreased very sharply around the estimated point (Fig. 1). This was first taken to indicate a problem with the model, but a reviewer mentioned that he had seen this once before, but only for some versions of ADMB. Suspicion grew that this could be caused by a problem with some versions of ADMB. We tested examples that we knew had worked in the past, but now they were failing in the same way.




Figure 1: Likelihood profile (red line) compared to the Gaussian approximation (black line) in a simple case $\log(a)$ in a Beverton-Holt model. The left panel shows the corrected version and the right panel when the bug is present. The code and data for this example is available in appendix A.

Once we knew it had to be an error in ADMB it was a straightforward (but fairly time-consuming) process to identify the code commit which caused the error. The code commit dated back to March 2014 and introduced a wrapper around the gradient calculation function. The wrapper worked for most things, but not for the profile likelihood calculation. The commit was rolled back and the profile likelihood was working again. To make sure that the fix to the profile likelihood did not break any other examples all the other built in examples were validated. To ensure that

the profile likelihood is not broken again in the future, a test of the correctness of the profile likelihood was added, which will be tested after all future commits.

The ADMB users email list was notified of the presence of this bug in versions 11.2 to 11.6. They were advised to install from source or use an earlier version until the next version of ADMB is released, if they need to do profiling.

Adding pbeta from R to ADMB via TINY_AD - Anders, Chris and Johnnoel demonstrated how to include an accurate pbeta function into ADMB. In continuation of the work done in the project Kristensen & Nielsen (2016), where the goal was to extend ADMB and TMB jointly, we aimed to add the pbeta function from R into ADMB via the TINY_AD. The current pbeta function in ADMB is originating from Numerical Recipes. The code from Numerical Recipes is indented to be fairly short and reasonably easy to read, understand, and copy, but it is not intended to be highly accurate. R has put a lot of effort into getting its special functions highly accurate, and the large user base of R generates a lot of testing of the functions in R.

The code from R is however not short and easy to read. For the pbeta function the code from Numerical Recipes is around 60 lines of code, but the code from R is almost 3000 lines (complicated and with plenty of goto statements). It is this long because it needs to be accurate in a lot of special cases. It would not be feasible to manually write adjoint code for such functions. TINY_AD is able to work with such code with a minimum of modifications (see Kristensen & Nielsen 2016), and can supply forward mode derivatives up to any order.

An ADMB program with the following GLOBALS_SECTION was constructed to test this capability.

```
GLOBALS_SECTION
#include <df1b2fun.h>
#include "beta/pbeta.hpp"

template<class Float>
Float ppbeta(Float x, Float a, Float b){
    return toms708::pbeta(x, a, b, 1, 0);
}

TINYFUN3(ppbeta, x, a, b);
VECTORIZ3_ttt(ppbeta);
```

Here the third line includes the code from R for the pbeta function. The version we use here is the code that has already been modified for TMB, so this is an example of using a common code base for TMB and ADMB, both in terms of TINY_AD and the specific code for pbeta.

Next a small wrapper function is constructed called ppbeta, as ADMB already has a function named pbeta.

Next the macro TINYFUN3, which was developed and described in Kristensen & Nielsen (2016), is used to write the derivatives found by TINY_AD into the structure needed for ADMB. Versions are constructed for constant types, for fixed effects models, and for random effects models. Finally, a macro is called to vectorize the function.

The function was validated to give exactly the same results as the pbeta function in R.

The entire example is added to the ADMB source tree under `admb/tests/admb/tiny`

First tests revealed that the long and accurate version was a bit slower than the version from Numerical Recipes, but more tests are needed. In any case accuracy is important for these special functions, so it is recommended to make the accurate function the default in ADMB.

Reference:

Kristensen, K. and Nielsen, A. 2016. Extending TMB and ADMB functionality jointly (available [here](#)).

Profile likelihood of derived quantities in TMB: Initial steps - Anders experimented with calculating profile likelihood of a derived quantities in TMB.

TMB already supports profile likelihood calculation of model parameters and linear functions of model parameters via the function 'tmbprofile'. It does not directly support likelihood profile calculation of a quantity which is a nonlinear function of model parameters (or random effects). The following describes an experimental way of getting the profile likelihood of a derived variable.

The derived quantity Q is a function of one or more model parameters (or random effects, which is a case not tested here). When calculating a point on the profile we must restrict the model parameters such that Q is equal to a given point q_1 and optimize subject to that restriction. This is then repeated for a number of points q_1, q_2, \dots, q_n spanning the support of the distribution.

A simple solution to restrict the optimizer to find a solution where $Q=q$ is to add a penalty term $K(Q-q)^2$ to the negative log-likelihood. The constant K must be big enough to ensure that the solution has Q close to q , but not so big that numerical problems are occurring.

As an example a derived quantity is defined as the sum of a and b in a Beverton-Holt model, where the model parameters are $\log(a)$ and $\log(b)$. The code to add the penalty to the cpp part of the TMB program could be:

```
Type ab=exp(loga)+exp(logb);
ADREPORT(ab);
DATA_SCALAR(abfix);
Type abpen=0.0;
```

```

if(!isNA(abfix)){
  abpen = Type(1.0e6)*(abfix-ab)*(abfix-ab);
}
REPORT(abpen);
return nll+abpen;

```

Here the first line computes the derived quantity 'ab', the second line reports the estimate and Hessian-based standard deviation. The third line reads in the value to fix it at.

The next 4 lines (4-7) calculates the penalty. If NA (not assigned) is read as the value to fix it at, then the penalty is set to zero.

Finally, the penalty is reported and the negative log-likelihood is returned as the negative log-likelihood of the model plus the penalty.

On the R-side of the TMB program, the value to fix at is first set to NA to allow the model run as normal (without the penalty). Then after the model has been optimized the Hessian-based sdreport of the quantity can be used to set the relevant fixed points to evaluate the profile likelihood in. Here 200 equidistant points between the estimate minus 7 standard deviations and the estimate plus 7 standard deviations were chosen.

```

est<-sdr$value["ab"]
std<-sdr$sd[names(sdr$value)=="ab"]
xx <- seq(est-7*std,est+7*std, length=200)

```

To make the optimization more stable we set up a model object which initializes all model parameters at the estimate found from the un-penalized run. Then we define a function to return the profile likelihood at a given fixed value. This function set the value to the fixed value. Rebuilds the computational tape. does the now penalized optimization and returns the value of the negative log-likelihood at the minimum (removing the penalty part). Finally, the function is evaluated in all points defined above and the values returned are converted to likelihood values and standardized. The code is the following:

```

obj <- MakeADFun(data, opt$par, DLL=basename)
fun<-function(x){
  obj$env$data$abfix <- x
  obj$retape()
  opt <- nlminb(obj$par, obj$fn, obj$gr)
  opt$objective-obj$report()$abpen
}
yy <- exp(-Vectorize(fun)(xx))
A <- sum(diff(xx) * (yy[-length(yy)]+yy[-1]))/2
prof<- list(x=xx,y=yy/A)

```





Figure 2: The estimated densities for $a+b$ in a Beverton-Holt model with $\log(a)$ and $\log(b)$ as model parameters. The methods are: Hessian-based from TMB (black solid line), Hessian based from ADMB (dashed blue line), profile likelihood from TMB (solid red line), and profile likelihood from ADMB (dashed green line).

The approach produced a plausible profile likelihood curve (Fig. 2). It is not identical to the similar plot produced by ADMB, but the definition of the profile likelihood in ADMB is a bit unconventional (ADMB manual p. 1-45).

For a single parameter, the approach outlined above gave identical profiles to the profiles from the 'tmbprofile' function, but more tests are needed in cases where the true profile likelihood function is known. An automatic way to set the penalty constant should be devised. When building this into TMB it could be wrapped in convenient macros and added to the existing tmbprofile function.

Video - The introductory TMB video produced by Hans, Arni, and Mollie at the previous developers' meeting was popular with new users. This year, the same team produced an additional video covering a multivariate normal distribution. The 5 minute video covers the multivariate normal density function, using the density namespace, transforming parameters, simple matrix operations, and typecasting. Code for the example (along with an a related example) is available from the ADMB Foundation's GitHub page in a new repository called "tmb-examples". The video and the accompanying example is useful to teach and learn TMB.

Documentation - Many changes were made to the numerous forms of documentation based on feedback from the group. Hans, Kasper, and Casper are making changes to the TMB Wiki and the [TMB manual](#). A set of introductory level TMB examples were collected on the ADMB Foundation web site in the directory [tmb-examples](#). This directory will also contain links to the videos produced by the foundation.

RStudio TMB integration - Gavin, Andrea, and Kasper have developed functionality that makes it easier to write and debug TMB programs in RStudio. This version of the TMB program is currently only available in a branch of the code on Github named [setupRStudio](#).

RStudio is a very popular IDE for the R programming environment. Many TMB users will be using RStudio. RStudio has many built-in features for supporting users, such as debugging tools, static code analysis for error-checking, and auto-complete code, including 'snippets' that allow shortcut addition of commonly used code. While the R and normal C++ components of TMB are already integrated through RStudio's support for these languages, it would be useful if TMB-specific objects and workflow were integrated into R Studio. Such functionality would improve the usability of TMB for RStudio users.

Andrea, Kasper, and Gavin worked to develop integration of TMB into RStudio, through the use a R function 'TMB:::setupRStudio()'. This function modifies the user's '.Rprofile' settings to enable direct sourcing of the TMB .cpp file, which provides access to R Studio's built-in error checking – giving concise error messages, and linking of errors to the appropriate places in the code. Secondly, a set of R and .cpp code snippets are added to RStudio's library of snippets so the user can access commonly used TMB functions (and entire script/code templates) by only typing a few characters.

This functionality is made available in the TMB package, currently in branch 'setupRStudio', and can be installed via:

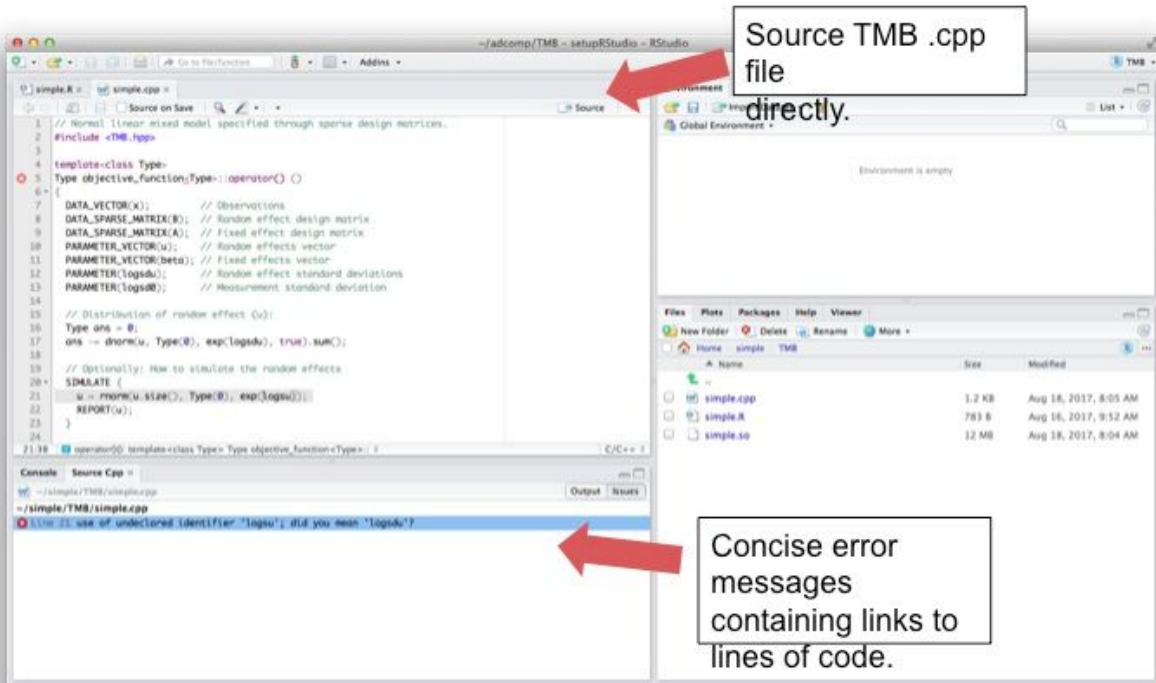
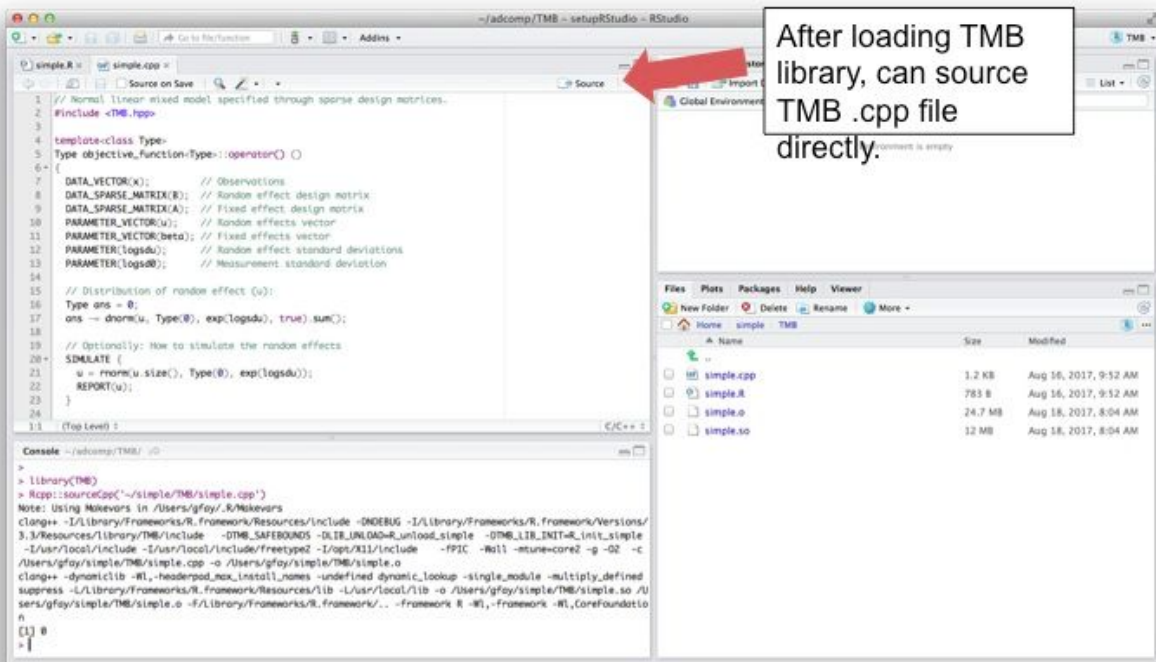
```
devtools::install_github("kaskr/adcomp/TMB", ref="setupRStudio")
```

The user can then run

```
TMB:::setupRStudio()
```

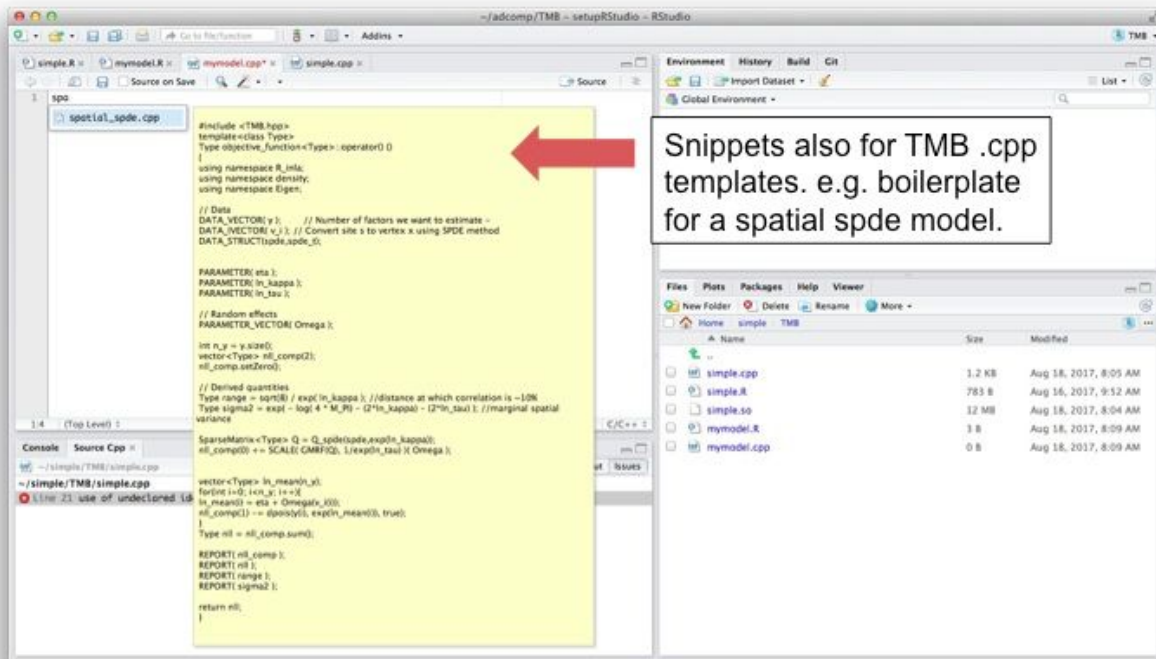
and will be prompted to restart RStudio to complete the integration.

Additional snippets can be added to the package. Of interest is to also make the TMB objects available to R Studio's static code analyzer, to have instant error checking, as is currently enabled for R and ordinary C++ (e.g. error markup in the RStudio gutter). It is envisaged that once this functionality is added, the R Studio integration can be merged into the main TMB package branch.



The screenshot shows the RStudio interface with a script editor on the left containing R code for TMB. A red arrow points to a snippet menu that is open, showing options like `tmb.load`, `tmb.nla`, `tmb.af`, `tmb.sp`, `tmb.profile`, and `tmb`. A text box with a red arrow pointing to the snippet menu contains the text: "R snippets for easy code completion for common TMB functions / operations." The Environment pane on the right shows "Global Environment" and "Environment is empty". The Files pane on the bottom right shows files: `simple.cpp` (1.2 KB), `simple.R` (783 B), and `simple.so` (12 MB).

The screenshot shows the RStudio interface with a script editor on the left containing R code for TMB workflow templates. A red arrow points to a snippet menu that is open, showing options like `tmb.rcscript`, `tmb.mesh`, `tmb.spde`, `tmb.unload`, `tmb.load`, `tmb.nla`, and `tmb.af`. A text box with a red arrow pointing to the snippet menu contains the text: "R snippets also contain complete TMB R script workflow templates." The Environment pane on the right shows "Global Environment" and "Environment is empty". The Files pane on the bottom right shows files: `simple.cpp` (1.2 KB), `simple.R` (783 B), `simple.so` (12 MB), and `mymodel.R` (0 B).



MCMC Examples - A main objective of the working group was to confirm that the MCMC samplers used in ADMB and TMB produce the same results as other established Bayesian methods. To accomplish this goal, Andrea and Mollie replicated STAN models from github.com/stan-dev/example-models/BPW into TMB .cpp code and ran models with adnuts replicating STAN setting for bounds, init values, and mcmc settings. One example (3.3.1 GLM Poisson) did not work using the TMB version of adnuts (wide bounds seemed to cause overflow and 0% acceptance rate). It was not possible to fix this during the workshop, but a reproducible example has been sent to Cole.

People doing stock assessments would like it if ADMB could do MCMC faster. The NUTS algorithm, contributed by foundation member Cole Monahan, has the potential to speed up sampling, by running fewer iterations with better mixing properties. Chris took on the task of confirming the speed on some stock assessment models that he uses at work. Initially, he had difficulty getting the NUTS version of ADMB installed on his computer, but eventually succeeded with Johnnoel's help. Now that Chris has a working version, he will confirm that the MCMC NUTS functionality is ready to be used in stock assessments.

Gavin also was successfully able to install ADMB with NUTS on Mac OS (10.9.5), and ran NUTS in ADMB from the command line. Gavin also installed Cole's R package adnuts, for implementing NUTS in both TMB and ADMB. As warned by Cole, some of the admb functionality in adnuts did not work out of the box, with some functionality appearing to currently

only being implemented on Windows. Gavin was able to successfully modify the appropriate R code to allow the adnuts R library to run on his Mac. Now that Gavin has a working version he will compare MCMC NUTS to RWM for simple fisheries assessment models used in his stock assessment course.

Laplace Approximation and Checking - Helen, Kasper, and Hans discussed various possible methods which could be used to give a more accurate approximation to the likelihood in cases where the Laplace approximation fails. Kasper suggested using variational methods based on a large class of nonlinear transformations of the normal distribution (for which the Laplace approximation is always exact). As an alternative approach, he proposed using a parametric family of transformations, controlled by an extra parameter, and choosing a good value for that parameter using the fact that the derivatives with respect to that parameter should be zero. Helen described the ideas behind the sequential reduction approximation, which can reduce the cost of finding a good approximation if the integrand has certain factorization structure, and we talked about how that could be combined with importance sampling.

To allow these new ideas to be tested, it was agreed to create a directory of simple examples which demonstrate the types of situation in which the Laplace approximation could fail.

Helen and Kasper discussed ways to check the accuracy of the Laplace approximation. Helen ran a simulation study on a simple two-level GLMM to compare the various approaches to this problem in that example. The code and a report on this will be made available on the workshop github page. In this example, a scaled version of the approximate bias computed by `checkConsistency` agrees closely with other measures in most cases. One problem case identified is when the estimate of the variance of random effects is close to zero, in which case the variance of the approximate bias is too large, even with a large number of runs. However, in that case problems in the inference should already be clear from the very large variance estimate for the parameter, with the need to run `checkConsistency`.

Phases in TMB - Gavin and Mollie continued work from the 2016 developers meeting to implement a wrapper for TMB that provides easy implementation of sequential parameter estimation in TMB via phases. While this functionality is already possible in TMB via the 'map' argument, it is not automatic. Phases is something ADMB users are very comfortable with and parameter estimation in phases is a common part of the model fitting workflow for many stock assessors, adding the option to do this in TMB should appeal to some users. A few issues with the code that had been raised since the 2016 meeting were fixed. The function now needs some more testing. `TMBphase` is added as an R package to the set of contributed 'TMB helper' R packages available via https://github.com/kaskr/TMB_contrib_R

ADMB-IDE - Two related objectives were identified as important ADMB-IDE development tasks:

(1) Prepare an ADMB-IDE release for the upcoming version of ADMB 12.0. Arni will continue to work on this after the workshop, synchronizing the newest versions of the IDE components: ADMB, Rtools, Emacs, modes, user settings, and installation scripts.

(2) Continue work on the semi-automatic build procedure of ADMB-IDE, an effort that was started at the Developers' Workshop 2016. The goal is to have a fully operational release of ADMB-IDE (exe and zip) as a product of the Buildbot suite of procedures that are applied to the development version of the ADMB codebase. For this workshop, Arni downloaded the current Buildbot zip, extracted the contents, and concluded that most or all components were there, for the user to install and configure manually. The installer (exe) has been a more popular way to fully install ADMB-IDE, but the installer was currently not available from the Buildbot. Johnoel and Arni will continue to work on this after the workshop, after task (1) above is completed.

Outreach Planning - Kasper, Anders and Mollie discussed plans to teach a course on TMB at the International Statistical Ecology Conference in St. Andrews, Scotland in July 2018. Originally, Hans was registered to teach with Kasper and Anders, but had to cancel attending the conference due to scheduling conflicts. Many foundation members will present models using ADMB and TMB at the conference.

Social Media - The ADMB foundation now has a Twitter account, handle @AdmbProject. Gavin has volunteered to maintain this account, which will focus on promoting new features to ADMB & TMB, highlight new papers using ADMB & TMB, and provide outreach for trainings and workshops. Gavin and Arni discussed Gavin getting access to the ADMB facebook account so he can link postings between the two accounts. The Twitter presence provides an additional opportunity for exposure of ADMB/TMB to new users, as there is a large R community that uses this platform.

ADMB Foundation Meeting

Allan Hicks (ADMB Foundation Treasurer) and Jim Ianelli joined remotely via Google Hangouts.

We discussed the possibility that funding from NOAA may be lost in the future. If NOAA funding were cut, the loss of a paid professional developer for ADMB would quickly deteriorate the maintenance and reliability of the software. One of the main consequences would be that as compilers and operating systems evolve, ADMB would become incompatible with many systems.

We agreed that the main purposes of the Foundation are to (1) fund Johnnoel to continue maintaining the core ADMB code, its compatibility with various operating systems, and its distribution through the Buildbot; and (2) to fund developers' meetings where non-routine tasks such as major changes are tackled.

Allan explained what he has learned from taking over the duties as treasurer. John Sibert and Allan are working on clearly documenting the definitions of the roles of board members and foundation members so that the next board transition will go smoother. This documentation should include a list of deadlines for routine duties such as the annual filing with the state of Hawaii. Having a clear list of duties will alleviate John Sibert from the need to remind the board of the tasks.

Many of the foundation's correspondences are addressed to John Sibert. It may be a good idea to establish another permanent address for the Foundation. Allan is ok with keeping the mailing address at IPHC for now.

Currently a Wells Fargo account in Seattle contains about 6000 USD. Funds in this account come from fees charged for teaching courses such as those at the biannual International Statistical Ecology Conference. Ownership of the account needs to be transferred from the former treasurer Steve Martell to Allan Hicks.

The board invited Andrea Havron and Helen Ogden to become Foundation members. It was agreed that the voluntary duties of Foundation members are to (1) answer user questions on the email lists when possible, (2) present ADMB and TMB models at conferences, courses, and workshops when appropriate, and (3) test the code and offer advancements and bug fixes when possible.

It was announced that foundation member Teresa A'mar will switch from doing fisheries research to doing cancer research. We expect that this will lead to an expansion in the user base of our methods.

Participant List

Name	Attendance	Affiliation	Notes
Anders Nielsen	in person	Technical University of Denmark	TMB developer
Andrea Havron	in person	University of Auckland, NZ	Statistics PhD student doing spatial models
Arni Magnusson	in person	ICES Secretariat, Denmark	ADMB Foundation President
Casper Berg	in person	Technical University of Denmark	TMB developer
Chris Grandin	in person	DFO Pacific Biological Station, Canada	
Colin Millar	part-time	ICES Secretariat, Denmark	
Gavin Fay	in person	University of Massachusetts Dartmouth, USA	
Hans Skaug	in person	University of Bergen, Norway	
Helen Ogden	in person	University of Southampton, UK	Mathematician researching the Laplace approximation
Jhnoel Ancheta	in person	University of Hawaii, USA	ADMB Core Developer
Kasper Kristensen	in person	Technical University of Denmark	TMB developer
Mollie Brooks	in person	Technical University of Denmark	ADMB Foundation Secretary

