



GRUPE D'ÉTUDES ET DE RECHERCHE  
EN ANALYSE DES DÉCISIONS

**Les Cahiers du GERAD**

**CITATION ORIGINALE / ORIGINAL CITATION**

**GERAD** HEC Montréal  
3000, ch. de la Côte-Sainte-Catherine  
Montréal (Québec) Canada H3T 2A7

**Tél. : 514 340-6053**  
Télec. : 514 340-5665  
info@gerad.ca  
www.gerad.ca

**CUTEst: A Constrained and Unconstrained  
Testing Environment with Safe Threads**

N. Gould, D. Orban,  
Ph.L. Toint

G-2013-27

April 2013



# CUTEst: A Constrained and Unconstrained Testing Environment with Safe Threads

## **Nick Gould**

*Scientific Computing Department  
Rutherford Appleton Laboratory  
Chilton, Oxfordshire, OX11 0QX, England*

`nick.gould@stfc.ac.uk`

## **Dominique Orban**

*GERAD & Department of Mathematics and Industrial Engineering  
Polytechnique Montréal  
Montréal (Québec) Canada, H3C 3A7*

`dominique.orban@gerad.ca`

## **Philippe L. Toint**

*Department of Mathematics  
University of Namur  
Namur, Belgium*

`philippe.toint@unamur.be`

April 2013

*Les Cahiers du GERAD*

G-2013-27

Copyright © 2013 GERAD

**Abstract:** We describe the most recent evolution of our constrained and unconstrained testing environment and its accompanying SIF decoder. Code-named `SIFDecode` and `CUTEst`, these updated versions feature dynamic memory allocation, a modern thread-safe Fortran modular design, a new Matlab interface and a revised installation procedure integrated with `GALAHAD`.

**Key Words:** CUTE, CUTEr, CUTEst, optimization, modeling, benchmarking.

**Résumé:** Nous décrivons la plus récente version de notre environnement de test avec et sans contraintes ainsi que de son décodeur SIF. Nommés `SIFDecode` et `CUTEst`, ces nouvelles versions comportent l'allocation dynamique de mémoire, une conception modulaire à fil sécurisé en Fortran 2003, une interface Matlab mise à jour ainsi qu'une procédure d'installation intégrée avec celle de `GALAHAD`.

**Mots clés :** CUTE, CUTEr, CUTEst, optimisation, langage de modélisation, étalonnage.

---

**Acknowledgments:** We are extremely grateful to Roger Fletcher, Philip Gill, Michal Kočvara, Michael Powell, Klaus Schittkowski and Elizabeth Wong for making their latest codes available to us so that we could build and test interfaces.

Research of the second author was partially supported by NSERC Discovery Grant 299010-04.

## 1 Introduction

The Constrained and Unconstrained Environment (CUTE) [3] and its associated set of optimization examples have been widely adopted and used since its introduction in 1993. The test set has grown over time and now numbers approximately 1150 examples, many of them of variable dimension. In addition, since both the MPS linear programming format [6] and its quadratic programming extensions [16; 17; 13] are compatible with CUTE's Standard Input Format (SIF) [5, Ch.7], CUTE and its successor CUTEr [12] provide access to other test sets such as those from Netlib [9] and Maros and Mészáros [16].

To set the scene, recall that a SIF file provides a portable description of group-partially separable optimization problem [4]. Such a file is translated by the package `SifDec` to a number of Fortran subroutines that compute values and derivatives of constituent element and group functions, together with data that explains how the functions are glued together. Armed with these components, CUTE(r) reassembles them to allow users to compute values and derivatives of both the objective function and/or its constraints as required, as well as providing static information such as bounds on variables and constraints.

The core (Fortran 77) routines behind CUTE(r) have not changed significantly since their original release. A main limitation of standard Fortran 77 is that it offers no mechanism for dynamic memory management, something programmers of other languages (particularly C) take for granted. This defect had a significant implication for CUTE and CUTEr, namely that a one-size-fits-all set of array dimensions are set at compile time, and CUTEr and its relatives return to the user if this choice is insufficient, with a recommendation for recompilation with values that might be appropriate. Many CUTEr users have learned to detest this inflexibility, and it is certainly the main source of complaint we receive.

Of course, modern Fortran (90 and later) provide dynamic memory-allocation, and it has mostly been the scale of the task of rewriting CUTEr to do this that has stopped us; CUTEr (and its dependent `SifDec`) numbered roughly 50,000 lines of code. Now, we have done so. The new, Fortran 2003 packages CUTEst and SIFDecode both request array storage as needed. The data for any array that is not currently large enough is written to temporary store (or, if there is insufficient room in memory, to disk), the array de-allocated and re-allocated with some “elbow-room” and the existing data copied back. A more general concern about the lack of freely-available reliable modern Fortran compilers has also vanished with the arrival of both `g95` and `gFortran`; there are of course many excellent commercial Fortran compilers available.

An additional limitation of CUTE/CUTEr was that it relied on Fortran common blocks to share data between tools. Some of this data, such as that required to describe the structure of a problem, was fixed after calling “setup” routines, while the remainder was constantly rewritten as problem function (and derivative) evaluations took place. This use of common meant that the packages were not thread safe, and could not be used in a multi-threaded (parallel) environment (for example to test branch-and-bound methods for integer programming or global optimization). This deficiency has also been addressed in CUTEst, and indeed the data is now split so that a single copy of the fixed data is available to all threads, while dynamic data is stored on a per-thread basis.

We have also taken the opportunity to revise the way the packages are organised and installed. We have adopted the scheme we currently use within GALAHAD [11], and this leads to some isolation of common components that are distributed separately. Since GALAHAD was itself a major user of CUTEr, we have also updated components of GALAHAD to use CUTEst packages. In addition, we have upgraded all of CUTEr's interfaces to external optimization packages—written in a variety of languages—that are still distributed (there have been a few casualties since 2003), and have also provided interfaces of a number of new ones. We have also considerably simplified the interface to Matlab.

This short paper is arranged as follows. We first describe how CUTEst is now organised, and how it may be used to test external packages. We then give a few details of the new features provided. A few words about the SIF decoder are followed by a description of the installation procedure.

## 2 New package organisation

CUTEst is provided as a series of directories and files, all lying beneath a root directory that we shall refer to as `$CUTEst`. The directory structure is illustrated in Figure 1.

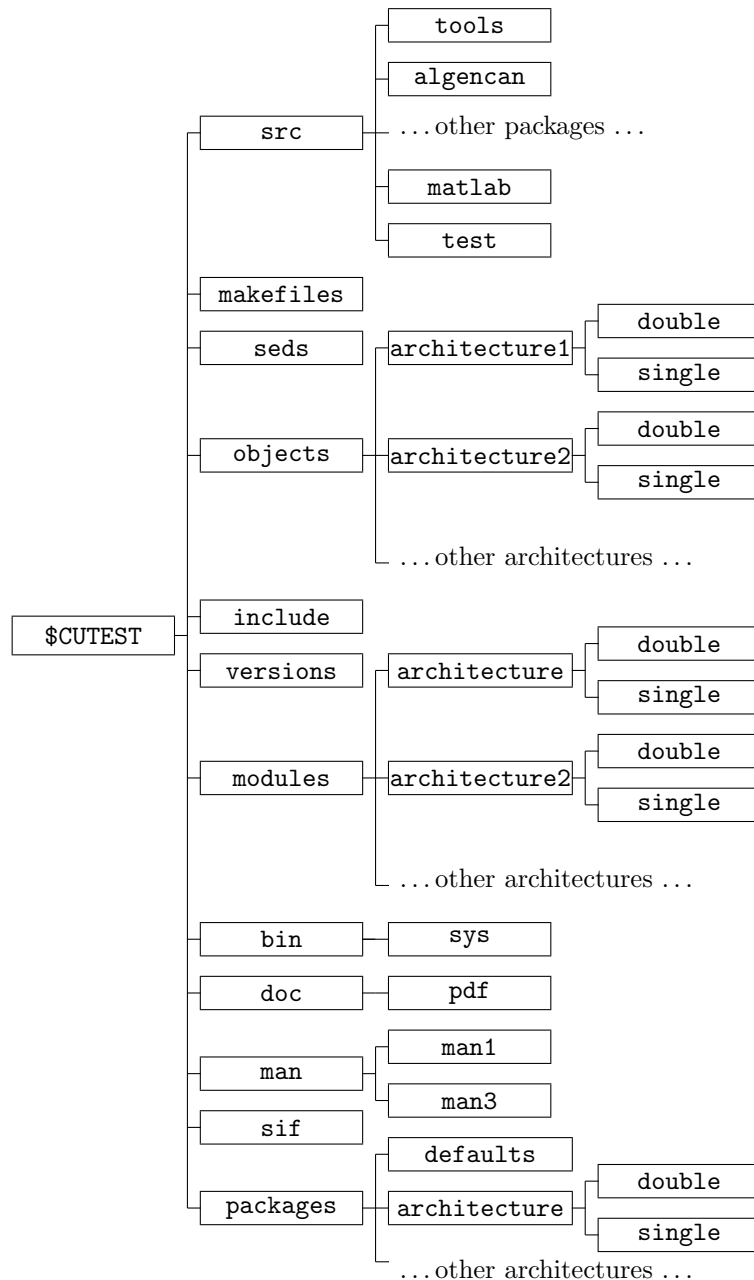


Figure 1: Structure of the CUTEst directories

Before installation the sub-directories `objects`, `modules`, `makefiles`, `versions` and `bin/sys` will all be empty. The script `install_cutest` prompts the user for the answers to a series of questions aimed at determining what machine type, operating system and compiler (from a large list of predefined possibilities encoded in a separate package, `ArchDefs`, that is distributed with CUTEst, and shared with SIFDecode and GALAHAD) to build CUTEst—we call this combination of a machine, operating system and compiler an

*architecture*. Each architecture is assigned a simple (mnemonic) architecture *code* name, say **architecture**—for example a version for the NAG Fortran 95 compiler on a PC running Linux is coded `pc.lnx.n95`, while another for the IBM Fortran 95 compiler on an PowerPC system running AIX is `ppc.aix.x95`. Having determined the architecture, the installation script builds sub-directories of **objects**, **modules** and **packages** named **architecture**, as well as further sub-directories **double** and **single** of these to hold architecture-dependent compiled libraries, module file information and external package linking information if required. In addition, architecture-dependent makefile information and environment variables for execution scripts are placed in files named **architecture** in the **makefiles** and **bin/sys** sub-directories, and a file recording how the code is related to the architecture is put in **versions**.

The source codes for each CUTEst package interface occurs in a separate sub-directory of the **src** directory. The main CUTEst evaluation tools all lie in the sub-directory **tools**, while a set of comprehensive test programs are available in **test**. The remaining sub-directories contain interface programs between the tools and each of the external packages supported. Each sub-directory contains the interface, a default options specification file, an example program that tests the interface without requiring the external package, a README that explains what a user needs to do to make the external package work with CUTEst, and a makefile. Since the order of compilation of Fortran modules is important, and as we have seen there is a strong interdependency between the CUTEst packages, the makefiles have to be carefully crafted. For this reason, we have chosen not to use variants of tools such as **imake** to build and maintain the makefiles. A set of configuration files that provides default link details between CUTEst and each external package is provided in the directory **packages/defaults**, and users may copy and modify these to architecture and dimension specific sub-directories of **packages** to override the default choices. Man pages for CUTEst as a whole and each individual tool are provided sub-directories of the **man** directory, and translations are also available in PDF format in the **pdf** sub-directory of the documentation directory **doc**.

Once the correct directory structure is in place, the installation script builds a random-access library of the required precision by visiting each of the sub-directories of **src** and calling the Unix utility **make**. CUTEst package interfaces are all written in double precision, but if a user prefers to use single precision, the makefiles call suitable Unix **sed** scripts (stored in **seds**) to transform the source prior to compilation. A user may choose to install all of CUTEst with or without Matlab support or just the tools themselves; the tools for unconstrained and constrained optimization may also be installed separately. Recompilation following updates is easily performed by issuing the command **make** from the **src** directory, while **make test** from the same directory runs comprehensive tests of all compiled components.

### 3 Interfaces to the CUTEst test set

To run one of the supported packages on an example stored in **EXAMPLE.SIF**, say, a user needs to issue the command

```
runcutest -A architecture -p package -D EXAMPLE[.SIF]
```

where **architecture** is the architecture code discussed in §2, **package** defines the package to be used—the manual page for **runcutest** gives a list of current possibilities—and the suffix **[.SIF]** is optional. This command translates the SIF file into Fortran subroutines and related data using the decoder provided in **SIFDecode**, and then calls the required optimization package to solve the problem. A default architecture may be defined by setting the environment variable **\$MYARCH**, and if so the **-A** flag may be avoided. Once a problem has been decoded, it may be re-used (perhaps with different options) using the auxiliary command

```
runcutest -A architecture -p package
```

For Matlab use, the command

```
cutest2matlab EXAMPLE[.SIF]
```

may be used instead; since Matlab is very specific about the Fortran compilers it supports, requests for a CUTEst-Matlab installation will adjust compiler options accordingly.



A few SIF examples are given in the `sif` directory, while the `runcutest` and `cutest2matlab` commands are in the `bin` sub-directory, and have man-page descriptions in the `man/man1` sub-directory.

## 4 Improvements

### 4.1 New features

As we mentioned in the introduction, CUTEst and SIFDecode both use the Fortran 2003 allocate/deallocate features to create and modify array storage. Each CUTEst tool uses a module CUTEEST that provides access to two derived types `CUTEEST_data_type` and `CUTEEST_work_type` used to store workspace arrays. The former collects data that describes problem structure and is set prior to any problem function evaluation, and unchanged thereafter, while the latter is used to hold data that may change at every evaluation. Data in these arrays is made available to the tools through a scalar `CUTEEST_data_global` of type `CUTEEST_data_type` and an allocatable array `CUTEEST_work_global` of type `CUTEEST_work_type`.

All evaluation tools are available as both simple (unthreaded) and threaded versions. For the latter, which may be distinguished by the suffix `_threaded`, `CUTEEST_work_global` is allocated to be large enough to hold all the threads that will be used by the setup subroutines, and each evaluation call specifies access to `CUTEEST_work_global(i)` for the particular thread `i` required.

We give a list of all CUTEst tools and their functionality in Appendix A. Most are rewritten version of their CUTEr counterparts, with argument-list changes to remove redundant size parameters and occasional order changes to handle inconsistencies. An additional `status` argument has been added to each to report any fatal memory errors (such as array allocation/deallocation failures) and inability to evaluate functions at specified values; CUTEr dealt with such eventualities by terminating execution. We have added termination tools to allow users to deallocate all storage created by the setup procedure when they have no longer need for it. In addition, constraints may now be ordered so that equalities precede or follow inequalities, or so that linear constraints precede or follow nonlinear ones, and variables so that those that only appear linearly in the problem precede or follow those that appear nonlinearly. New tools have also been introduced to describe the sparsity patterns of the Hessian of the objective and Lagrangian functions. Finally, a new tool has been added to compute the sparse gradient of the objective function for constrained problems; this corrects an oversight since we already provide similar functionality for the gradients of individual constraints.

### 4.2 New Matlab calls

The Matlab interface has been substantially revised and simplified. It is now entirely written in C instead of Fortran to facilitate interaction with Matlab's own API and usage of its index types. The interface merges the constrained and unconstrained tools together so users may use familiar and consistent function calls such as `cutest_obj()` to obtain the objective function value and possibly its gradient regardless of the presence of constraints. In this regard, the interface exploits Matlab's ability to determine how many output arguments are required by the user. This allows both `f = cutest_obj(x)` and `[f,g] = cutest_obj(x)`. Another example is the `cuter_cons()` function, which allows to evaluate all or individual constraints and/or constraint gradients.

A problem is "loaded" into Matlab by calling a simple function with no arguments: `prob = cutest_setup()`. The single output argument of the setup function is a Matlab structure containing problem data such as the number of variables, number of constraints, number of nonzeros in the Jacobian and Hessian, initial guess, bounds, initial multipliers, and arrays indicating which constraints are linear and which are equality constraints. Those fields may be accessed using the familiar dot notation, e.g., `prob.n`, `prob.x`, `prob.cl`, etc.

Help is included with all CUTEst tools, available to Matlab users by way of the familiar `help` call. In particular, `help cutest` gives an overview of the tools available.

Finally, the current interface makes it easier to decode problems and build the corresponding MEX files from inside Matlab. For instance, the commands `probnam = 'LUBRIFC'; unix(['cutest2matlab ', probnam])` generate the MEX file corresponding to problem LUBRIFC in the current directory.

See Appendix B for a complete description of the Matlab CUTEst tools and their functionality.

### 4.3 New interfaces

In addition to the still-current packages supported by CUTEr, CUTEst provides new interfaces to ALGENCAN [2], BOBYQA [19], Direct Search [7], filterSD [8], NEWUOA [18], NLPQLP [21], NOMAD [15], PENNLP [14], QL [20], SPG [1] and SQIC [10], as well as various new packages within GALAHAD [11].

Interfaces to the obsolete packages `hsl_ve12`, `osl`, `va15`, `ve09` and `ve14` supported in CUTEr have been withdrawn.

### 4.4 New test examples

Almost 200 new examples have been added to the test-problem collection since the release of CUTEr. These include large collections of problems arising from linear complementarity, and of real-life quadratic programming problems.

All test problems are now under version control in the same way as the source code and users may update their local repository easily when new problems are added or changes are made to existing problems.

## 5 A revised SIF decoder

As we mentioned in the introduction, the lack of dynamic memory allocation affects the SIF decoding package SifDec just as severely as it does CUTEr. A new stand-alone Fortran package SIFDecode has been written to address this issue. All of the functionality of the subroutines previously in SifDec have been subsumed into a single Fortran 2003 module SIFDECODE. Since required array sizes are not known beforehand, default initial values are increased as required as the package makes a single pass through the SIF file under consideration; default initial values may be changed to make the processing more efficient, but this is not crucial.

The distributed package is organised in the same way as CUTEst (see §2), although now the `src` directory simply contains two source sub-directories, `decode` that holds the decoder and its main program, and `select` containing the test-problem database interrogation tools from CUTE [3, §2.3].

Once the package has been installed, the decoder is called by issuing the command

```
sifdecoder -A architecture EXAMPLE[.SIF]
```

where `architecture` and `EXAMPLE.SIF` are as before; the `-A` option may be omitted when using the default architecture.

## 6 New installation procedures

The installation procedure has been updated to recognise that most users will need to install both SIFDecode and CUTEst, and may also wish to integrate these with GALAHAD. A single script, `install_optsuite`, prompts the user to describe what features are needed and which architecture is desired. An opportunity to modify default compilation flags is provided, after which the script will automatically download and install the software.

## 7 Obtaining the packages

All of the required and optional packages ArchDefs, SIFDecode, CUTEst and GALAHAD are available from the CCPForge project, funded by The Joint Information Systems Committee (<http://www.jisc.ac.uk>) and maintained by the Scientific Computing Department of the Science and Technology Facilities Council (<http://www.stfc.ac.uk/SCD/default.aspx>) under the departmental SLA with EPSRC. See

<http://ccpforge.cse.rl.ac.uk/gf/project/cutest/wiki/>

for download details.

Both SIFDecode and CUTEst are distributed and made available under the terms of the GNU Lesser General Public License. See

<http://www.gnu.org/licenses/lgpl-3.0.txt>

for details.

## 8 Conclusions and perspectives

We believe CUTEst is a considerable improvement over past versions because of its improved modular and thread-safe design exploiting recent additions to the Fortran standard, dynamic allocation, simplified and unified tool calling sequences, improved Matlab interface and, last but not least, the more user-friendly installation process. Despite the age of the standard input format and the advent of more modern modeling languages, CUTE and CUTEr remained widely-used tools in the optimization community and beyond, as illustrated by the large number of user comments and request that we receive and the associated problem collection remains a staple of optimization software testing and benchmarking.

Since all packages and test problems are maintained and distributed via a source code revision system, bug fixes, improvements and additions are easily available.

## A Available tools

Separate evaluation tools are provided for unconstrained and constrained problems. Both unthreaded and threaded versions are available when this is relevant. See the appropriate man page for full details.

### A.1 Unconstrained problems

`cutest_udimen` (both threaded and unthreaded) determine the number of variables.

`cutest_usetup` (unthreaded) and `cutest_usetup_threaded` (threaded) setup internal data structures and determine variable bounds.

`cutest_unames` (both threaded and unthreaded) determine the names of the problem and the variables.

`cutest_uvartype` (both threaded and unthreaded) determine whether the variables are continuous or discrete.

`cutest_udimsh` (both threaded and unthreaded) determine the number of nonzeros in the sparse Hessian.

`cutest_udimse` (both threaded and unthreaded) determine the number of nonzeros in the finite-element Hessian.

`cutest_ufn` (unthreaded) and `cutest_ufn_threaded` (threaded) evaluate the objective function value.

`cutest_ugr` (unthreaded) and `cutest_ugr_threaded` (threaded) evaluate the gradient of the objective function.

`cutest_uofg` (unthreaded) and `cutest_uofg_threaded` (threaded) evaluate both the values and gradients of the objective function.

`cutest_udh` (unthreaded) and `cutest_udh_threaded` (threaded) evaluate the Hessian of the objective function as a dense matrix.

`cutest_ugrdh` (unthreaded) and `cutest_ugrdh_threaded` (threaded) evaluate the objective gradient and dense Hessian.

`cutest_ushp` (both threaded and unthreaded) evaluate the sparsity pattern of the Hessian of the objective function.

`cutest_ush` (unthreaded) and `cutest_ush_threaded` (threaded) evaluate the Hessian of the objective function as a sparse matrix.

`cutest_ugrsh` (unthreaded) and `cutest_ugrsh_threaded` (threaded) evaluate the objective gradient and sparse Hessian.

`cutest_ueh` (unthreaded) and `cutest_ueh_threaded` (threaded) evaluate the Hessian of the objective function as a finite-element matrix.

`cutest_ugreh` (unthreaded) and `cutest_ugreh_threaded` (threaded) evaluate the objective gradient and finite-element Hessian.

`cutest_uhprod` (unthreaded) and `cutest_uhprod_threaded` (threaded) evaluate the product of the Hessian of the objective function with a vector.

`cutest_ubandh` (unthreaded) and `cutest_ubandh_threaded` (threaded) obtain the part of the Hessian of the objective that lies within a specified band.

`cutest_ureport` (unthreaded) and `cutest_ureport_threaded` (threaded) discover how many evaluations have occurred and how long this has taken.

`cutest_uterminate` (both unthreaded and threaded) remove internal data structures when they are no longer needed.

## A.2 Constrained problems

`cutest_cdimen` (both threaded and unthreaded) determine the number of variables and constraints.

`cutest_csetup` (unthreaded) and `cutest_csetup_threaded` (threaded) setup internal data structures and determine variable and constraint bounds.

`cutest_cnames` (both threaded and unthreaded) determine the names of the problem, the variables and the constraints.

`cutest_connames` (both threaded and unthreaded) determine the names of the constraints.

`cutest_cvartype` (both threaded and unthreaded) determine whether the variables are continuous or discrete.

`cutest_cdimsj` (both threaded and unthreaded) determine the number of nonzeros in sparse constraint Jacobian.

`cutest_cdimsh` (both threaded and unthreaded) determine the number of nonzeros in the sparse Hessian.

`cutest_cdimse` (both threaded and unthreaded) determine the number of nonzeros in the finite-element Hessian.

`cutest_cfn` (unthreaded) and `cutest_cfn_threaded` (threaded) evaluate the objective function and constraint values.

`cutest_cgr` (unthreaded) and `cutest_cgr_threaded` (threaded) evaluate the gradients of the objective function and constraints.

`cutest_cofg` (unthreaded) and `cutest_cofg_threaded` (threaded) evaluate both the value and gradient of the objective function.

`cutest_cofsg` (unthreaded) and `cutest_cofsg_threaded` (threaded) evaluate both the value and sparse gradient of the objective function.

`cutest_csgr` (unthreaded) and `cutest_csgr_threaded` (threaded) evaluate the sparse gradients of the objective function and constraints.

`cutest_ccfg` (unthreaded) and `cutest_ccfg_threaded` (threaded) evaluate the values and gradients of the constraints.

`cutest_ccfsg` (unthreaded) and `cutest_ccfsg_threaded` (threaded) evaluate the values and sparse gradients of the constraints.

`cutest_ccifg` (unthreaded) and `cutest_ccifg_threaded` (threaded) evaluate the value and gradient of an individual constraint.

`cutest_ccifsg` (unthreaded) and `cutest_ccifsg_threaded` (threaded) evaluate the value and sparse gradient of an individual constraint.

`cutest_cdh` (unthreaded) and `cutest_cdh_threaded` (threaded) evaluate the Hessian of the Lagrangian function as a dense matrix.

`cutest_cidh` (unthreaded) and `cutest_cidh_threaded` (threaded) evaluate the Hessian of the objective function or an individual constraint as a dense matrix.

`cutest_cgrdh` (unthreaded) and `cutest_cgrdh_threaded` (threaded) evaluate the constraint Jacobian and Hessian of the Lagrangian function as dense matrices.

`cutest_cshp` (both threaded and unthreaded) evaluate the sparsity pattern of the Hessian of the Lagrangian function.

`cutest_csh` (unthreaded) and `cutest_csh_threaded` (threaded) evaluate the Hessian of the Lagrangian function as a sparse matrix.

`cutest_cshc` (unthreaded) and `cutest_cshc_threaded` (threaded) evaluate the Hessian of the Lagrangian function not including the objective as a sparse matrix.

`cutest_cish` (unthreaded) and `cutest_cish_threaded` (threaded) evaluate the Hessian of the objective function or an individual constraint as a sparse matrix.

`cutest_csgrsh` (unthreaded) and `cutest_csgrsh_threaded` (threaded) evaluate the constraint Jacobian and Hessian of the Lagrangian function as sparse matrices.

`cutest_ceh` (unthreaded) and `cutest_ceh_threaded` (threaded) evaluate the Hessian of the Lagrangian function as a finite-element matrix.

`cutest_csgreh` (unthreaded) and `cutest_csgreh_threaded` (threaded) evaluate the constraint Jacobian as a sparse matrix and the Hessian of the Lagrangian function as a finite-element matrix.

`cutest_chprod` (unthreaded) and `cutest_chprod_threaded` (threaded) evaluate the product of the Hessian of the Lagrangian function with a vector.

`cutest_chcprod` (unthreaded) and `cutest_chcprod_threaded` (threaded) evaluate the product of the Hessian of the Lagrangian function not including the objective with a vector.

`cutest_cjprod` (unthreaded) and `cutest_cjprod_threaded` (threaded) evaluate the product of the constraint Jacobian or its transpose with a vector.

`cutest_creport` (unthreaded) and `cutest_creport_threaded` (threaded) discover how many evaluations have occurred and how long this has taken.

`cutest_cterminate` (both unthreaded and threaded) remove internal data structures when they are no longer needed.

### A.3 Both unconstrained problems and constrained problems

`cutest_probname` (both threaded and unthreaded) determine the name of the problem.

`cutest_varnames` (both threaded and unthreaded) determine the names of the variables.

A call to `cutest_u/csetup[_threaded]` must precede calls to any other evaluation tool with the exception of `cutest_u/cdimen`. Once `cutest_u/cterminate[_threaded]` has been called, no further calls should be made without first recalling `cutest_u/csetup[_threaded]`.

## B Matlab interfaces

Table 1: Available Matlab tools

Matlab tool	CUTEst tool(s)	Purpose
<code>cutest_dims</code>	<code>cdimen</code>	Obtain problem dimensions
<code>cutest_setup</code>	<code>usetup / csetup</code>	Setup problem data structure
<code>cutest_obj</code>	<code>uofg / cofg</code>	Evaluate objective function value and its gradient if requested
<code>cutest_grad</code>	<code>ugr / cgr</code>	Evaluate objective function gradient
<code>cutest_sobj</code>	<code>cofsg</code>	Evaluate objective function value and its gradient as a sparse vector if requested
<code>cutest_objcons</code>	<code>cfn</code>	Evaluate objective and constraints
<code>cutest_cons</code>	<code>ccifg</code>	Evaluate constraint bodies and their gradients if requested. Evaluate a single constraint value and its gradient if requested
<code>cutest_scons</code>	<code>ccifsg</code>	Evaluate constraint bodies and Jacobian in sparse format. Evaluate a single constraint value and its gradient as a sparse vector
<code>cutest_lagjac</code>	<code>cgr</code>	Evaluate Jacobian and gradient of either objective or Lagrangian
<code>cutest_slagjac</code>	<code>csgr</code>	Evaluate Jacobian in sparse format and gradient of either objective or Lagrangian as a sparse vector
<code>cutest_Jprod</code>	<code>cjprod</code>	Evaluate the matrix-vector product between the Jacobian and a vector
<code>cutest_Jtprod</code>	<code>cjprod</code>	Evaluate the matrix-vector product between the transpose Jacobian and a vector
<code>cutest_hess</code>	<code>udh / cdh</code>	Evaluate the Hessian matrix of the Lagrangian, or of the objective if the problem is unconstrained
<code>cutest_ihess</code>	<code>udh / cidh</code>	Evaluate the Hessian matrix of the i-th problem function (i=0 is the objective function), or of the objective if problem is unconstrained
<code>cutest_hprod</code>	<code>uhprod / chprod</code>	Evaluate the matrix-vector product between the Hessian of the Lagrangian (or the objective if unconstrained) and a vector
<code>cutest_gradhess</code>	<code>ugrdh / cgrdh</code>	Evaluate the gradient of either the objective or the Lagrangian, the Jacobian (or its transpose) and the Hessian of the Lagrangian in dense format
<code>cutest_sphess</code>	<code>ush / csh</code>	Evaluate the Hessian matrix of the Lagrangian, or of the objective if the problem is unconstrained, in sparse format
<code>cutest_isphess</code>	<code>ush / cish</code>	Evaluate the Hessian matrix of the i-th problem function (i=0 is the objective function), or of the objective if problem is unconstrained, in sparse format
<code>cutest_varnames</code>	<code>varnames</code>	Obtain variable names as a list of strings
<code>cutest_connames</code>	<code>cnames</code>	Obtain constraint names as a list of strings
<code>cutest_terminate</code>	<code>utermminate / cterminate</code>	Remove existing internal workspace

## References

- [1] E. G. Birgin, J. M. Martinez, and M. Raydan. Algorithm 813: SPG—software for convex-constrained optimization. *ACM Transactions on Mathematical Software*, 27:340–349, 2001.
- [2] E. G. Birgin, R. Castillo, and J. M. Martinez. Numerical comparison of augmented Lagrangian algorithms for nonconvex problems. *Computational Optimization and Applications*, 31(1):31–56, 2005.
- [3] I. Bongartz, A. R. Conn, N. I. M. Gould, and Ph. L. Toint. CUTE: Constrained and unconstrained testing environment. *ACM Transactions on Mathematical Software*, 21(1):123–160, 1995.
- [4] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. An introduction to the structure of large scale nonlinear optimization problems and the LANCELOT project. In R. Glowinski and A. Lichnewsky, editors, *Computing Methods in Applied Sciences and Engineering*, pages 42–51, Philadelphia, USA, 1990. SIAM.
- [5] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. *LANCELOT: a Fortran package for Large-scale Nonlinear Optimization (Release A)*. Springer Series in Computational Mathematics. Springer Verlag, Heidelberg, Berlin, New York, 1992.
- [6] International Business Machine Corporation. Mathematical programming system/360 version 2, linear and separable programming—user’s manual. Technical Report H20-0476-2, IBM Corporation, 1969. MPS Standard.
- [7] E. D. Dolan, A. P. Gurson, P. L. Shepherd, C. M. Siefert, V. J. Torczon, and A. Yates. C++ direct searches. [http://www.cs.wm.edu/~va/software/DirectSearch/direct\\_code/](http://www.cs.wm.edu/~va/software/DirectSearch/direct_code/), 2001.
- [8] R. Fletcher. A sequential linear constraint programming algorithm for NLP. *SIAM Journal on Optimization*, 22(3):772–794, 2012.
- [9] D. M. Gay. Electronic mail distribution of linear programming test problems. Mathematical Programming Society COAL Newsletter, December 1985. See <http://www.netlib.org/lp/data/>.
- [10] P. E. Gill and E. Wong. Methods for convex and general quadratic programming. Technical Report NA 10-1, Dept. of Mathematics, University of California, San Diego, 2013.
- [11] N. I. M. Gould, D. Orban, and Ph. L. Toint. GALAHAD—a library of thread-safe fortran 90 packages for large-scale nonlinear optimization. *ACM Transactions on Mathematical Software*, 29(4):353–372, 2003.
- [12] N. I. M. Gould, D. Orban, and Ph. L. Toint. CUTer (and SifDec), a Constrained and Unconstrained Testing Environment, revisited. *ACM Transactions on Mathematical Software*, 29(4):373–394, 2003.
- [13] IBM Optimization Solutions and Library. *QP Solutions User Guide*. IBM Corporation, 1998.
- [14] M. Kocvara and M. Stingl. PENNON: A code for convex nonlinear and semidefinite programming. *Optimization Methods and Software*, 18(3):317–333, 2003.
- [15] S. Le Digabel. Algorithm 909: NOMAD: Nonlinear optimization with the MADS algorithm. *ACM Transactions on Mathematical Software*, 37(4):1–15, 2011.
- [16] I. Maros and C. Mészáros. A repository of convex quadratic programming problems. *Optimization Methods and Software*, 11-12:671–681, 1999.
- [17] D. B. Ponceleón. *Barrier methods for large-scale quadratic programming*. PhD thesis, Department of Computer Science, Stanford University, Stanford, California, USA, 1990.
- [18] M. J. D. Powell. The NEWUOA software for unconstrained optimization without derivatives. In G. Di Pillo and M. Roma, editors, *Large-Scale Nonlinear Optimization, volume 83 of Nonconvex Optimization and Its Applications*, pages 255–297, Heidelberg, Berlin, New York, 2006. Springer Verlag.
- [19] M. J. D. Powell. The BOBYQA algorithm for bound constrained optimization without derivatives. Technical Report DAMTP NA2009/06, Department of Applied Mathematics and Theoretical Physics, Cambridge University, Cambridge, UK, 2009.
- [20] K. Schittkowski. QL: A Fortran code for convex quadratic programming—User’s guide, Version 2.11. Technical report, University of Bayreuth, Department of Computer Science, 2005.
- [21] K. Schittkowski. NLPQLP: A Fortran implementation of a sequential quadratic programming algorithm with distributed and non-monotone line search. Technical report, University of Bayreuth, Department of Computer Science, 2010.