qpOASES **User's Manual**

Version 2.0 (June 2009)

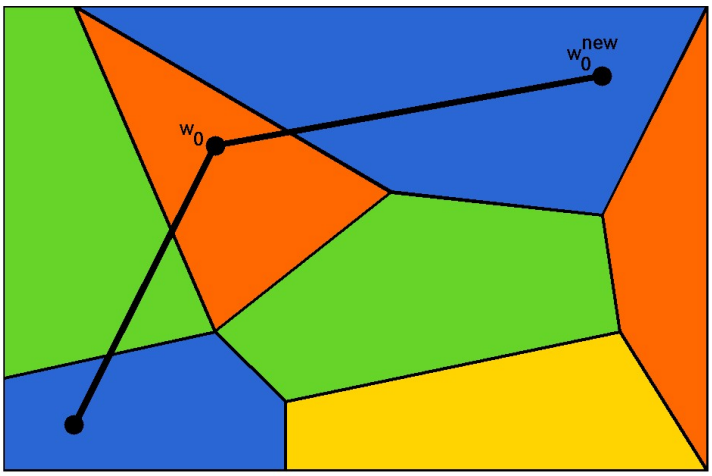

Hans Joachim Ferreau et al.[1]
Optimization in Engineering Center (OPTEC) and
Department of Electrical Engineering, K. U. Leuven
support@qpOASES.org

[1] qpOASES developers in alphabetical order: Hans Joachim Ferreau, Eckhard Arnold, Holger Diedam, Boris Houska, Aude Perrin, Thomas Wiese

Contents

Chapter 1

Introduction

Model predictive control (MPC) is an advanced control strategy which allows to determine inputs of an arbitrary process that optimise the forecasted process behaviour. These inputs, or control actions, are calculated repeatedly using a mathematical *process model* for the prediction. In doing so, the fast and reliable solution of convex *quadratic programming* problems in real-time becomes a crucial ingredient of nearly all algorithms for both linear and nonlinear model predictive control. The success of linear MPC—where just one *quadratic program (QP)* needs to be solved at each sampling instant—can even be attributed to the fact that highly efficient and reliable methods for QP solution have existed for decades, and that their computation times are much smaller than the required sampling times in typical applications. On the other hand, in nonlinear MPC algorithms, quadratic programs often arise as subproblems during the iterative nonlinear solution procedure, so that not only one, but several QPs need to be solved at each sampling instant.

qpOASES is an open-source implementation of the recently proposed online active set strategy (see [3], [4], [2]), which was inspired by important observations from the field of parametric quadratic programming. It builds on the expectation that the optimal active set does not change much from one quadratic program to the next and has several theoretical features that make it particularly suited for model predictive control applications. The software package qpOASES implements these ideas for solving QPs of the following form

$$\begin{aligned} \min_{x} \quad & \tfrac{1}{2} x^T H x + x^T g(w_0) \\ \text{s.t.} \quad & \mathrm{lb}A(w_0) \leq Ax \leq \mathrm{ub}A(w_0)\,, \\ & \mathrm{lb}(w_0) \leq \ x \ \leq \mathrm{ub}(w_0)\,, \end{aligned}$$

where the Hessian matrix is symmetric and positive (semi-)definite and the gradient vector as well as the bound and constraint vectors depend affinely on the parameter w_0. It has already been successfully used for closed-loop control of a *real-world Diesel engine* [5].

This manual is organised as follows: first, the installation of the qpOASES software package is explained in Chapter 2. Afterwards, a concise description of its main functionality is given in Chapter 3, which should enable you to use qpOASES within a couple of minutes. Chapter 4 describes special variants for QPs with varying matrices, simply bounded QPs as well as QPs with semi-definite Hessian matrices; advanced functionality like obtaining

status information or using the concept of a so-called initialised homotopy are presented in Chapter 5. Various interfaces to third-party software are presented in Chapter 6. Finally, Chapter 7 (which is mainly intended for software developers) provides some insight into the internal programming structure of qpOASES and options for further tuning of the algorithm.

Further information and a list of frequently asked questions can be found on

http://www.qpOASES.org/ .

If you have got questions, remarks or comments on qpOASES, please contact the main author:

Hans Joachim Ferreau
Katholieke Universiteit Leuven
Department of Electrical Engineering (ESAT)
Kasteelpark Arenberg 10, bus 2446
B-3001 Leuven-Heverlee, Belgium

Phone: +32 16 32 03 63
E-mail: joachim.ferreau@esat.kuleuven.be / support@qpOASES.org

Also bug reports and source code extensions are most welcome!

Acknowledgements

I would like to express my deep gratitude to all people who helped me in developing qpOASES.

First of all I thank my Diplom thesis supervisors Professor Dr. Dr. h. c. Hans Georg Bock, head of the "Simulation and Optimization Group" of the Interdisciplinary Center for Scientific Computing (IWR) at University of Heidelberg, and Professor Dr. Moritz Diehl, Department of Electrical Engineering (ESAT) and principal investigator of the Optimization in Engineering Center (OPTEC) at K. U. Leuven, for intensive personal support and excellent mathematical advice. The online active set strategy builds on their ideas and they also encouraged me to make the qpOASES source code publicly available.

Moreover, I owe many thanks to all my former colleagues in Heidelberg and my new ones in Leuven for inspiring discussions and pleasant conversations that were more or less related to the qpOASES software package.

Finally, I would like to thank Peter Ortner, Peter Langthaler and Luigi del Re, Institute for Design and Control of Mechatronical Systems at JKU Linz, for making possible the first run of qpOASES on real (Diesel engine) controller hardware.

The initial version of the software has been partly developed within the framework of the REGINS-PREDIMOT European project whose financial support is acknowledged. Further development of the code has been supported by Research Council KUL: CoE EF/05/006 Optimization in Engineering Center (OPTEC). The main author currently holds a PhD fellowship of the Research Foundation – Flanders (FWO) whose financial support and permission to work on this open-source software project is gratefully acknowledged.

Hans Joachim Ferreau

Chapter 2

Installation

The software package qpOASES is written in an object-oriented manner in C++ and comes along with fully commented source code files. Besides some standards libraries[1] *no further external software packages are required*.

For installing qpOASES under LINUX, perform the following steps:

1. *Download the current version of* qpOASES from

   http://www.qpOASES.org/

   by saving the file qpOASES-2.0.tar.gz on your local machine.

2. *Unpack the archive:*

   ```
   gunzip qpOASES-2.0.tar.gz
   tar xf qpOASES-2.0.tar
   ```

   A new directory qpOASES-2.0 will be created; from now on we refer to (the full path of) this directory by <install-dir>. It contains five subfolders, namely

   - SRC (qpOASES source files),
   - INCLUDE (qpOASES header files),
   - EXAMPLES (example files for setting up your own QP problems),
   - INTERFACES (interfaces to third-party software),
   - DOC (this manual and a DOXYGEN configuration file).

3. qpOASES is distributed under the terms of the GNU Lesser General Public License 2.1, which you can find in the file <install-dir>/LICENSE.txt or Appendix A of this manual. *Please read this licence file carefully before you proceed with the installation*, as you implicitly agree with this licence by using qpOASES!

[1] math.h, stdio.h, stdlib.h, string.h (as well as sys/time.h, sys/stat.h or windows.h for runtime measurements)

4. If you want to use qpOASES via the provided third-party interfaces only, you can skip the following steps and proceed as described in Chapter 6. Otherwise continue with the

   Compilation of the qpOASES *library* libqpOASES.a[2]:

   ```
   cd <install-dir>/SRC
   make
   ```

   This library libqpOASES.a provides the complete core functionality of the qpOASES software package. It can be used by, e.g., linking it against a main function from the EXAMPLES folder.

5. *Compilation of a set of simple test examples:*

   ```
   cd <install-dir>/EXAMPLES
   make
   ```

   Among others, an executable called example1 should have been created; run it in order to test your installation. If it terminates after successfully solving two QPs, qpOASES has been successfully installed!

6. *Optional, create source code documentation*[3]*:*

   ```
   cd <install-dir>/DOC
   doxygen doxygen.config
   ```

   Afterwards, you can open the file <install-dir>/DOC/HTML/index.html with your favorite browser in order to view qpOASES's source code documentation.

Remarks:

- It is also possible to install qpOASES on a WINDOWS or MAC OS machine as it does not require LINUX-specific commands.
- If compilation fails due to the fact that the snprintf() function is not supported, you might uncomment line 54 within <install-dir>/INCLUDE/MessageHandling.hpp and try to compile again.

[2] The make command also creates a library called libqpOASESextras.a whose meaning is described in Section 5.6.

[3] All source code files are commented in a way suitable for the documentation system DOXYGEN [6].

Chapter 3

Getting Started

This chapter explains to you within a few minutes how to solve a quadratic programming (QP) problem, or a whole sequence of them, by means of qpOASES. At the end a tutorial example is presented that might serve as a template for your own QPs.

3.1 Outline

Core of qpOASES is the QProblem class which is able to store, process and solve convex quadratic programs using the online active set strategy; it makes use of several auxiliary classes (see Chapter 7). Except for special situations, the QProblem class is intended to be the only *user interface* to qpOASES's functionality.

For solving a series of convex quadratic programs with fixed Hessian and constraint matrix, the following steps are necessary:

1. create an instance of the QProblem class,
2. initialise your QProblem object and solve the first QP (specified by its QP matrices and vectors),
3. solve each following QP by passing its vectors to your QProblem object.

Now, we will explain these three steps in more detail. Various variants and special cases are treated in later chapters for the ease of presentation.

3.2 Main Steps

Creating an Instance of the QProblem Class

Creating an QProblem object is done by means of the following constructor

```
QProblem( int nV, int nC );
```

which takes the number of variables nV and the number of constraints nC of the quadratic program sequence to be solved. At the moment it is not possible to solve QP sequences with varying problem dimensions. However, you might disable/enable constraints within a QP sequence (see Section 5.5).

Summary of the first step:
You can create an instance `example` of the `QProblem` class with the following command:
`QProblem example( nV,nC );`

Initialisation and Solution of First QP

The second step requires to initialise all internal data structures of the `QProblem` object and the solution of the first QP. Both can be accomodated with a single call to the following function:

```
returnValue init( const double* const H,
                  const double* const g,
                  const double* const A,
                  const double* const lb,
                  const double* const ub,
                  const double* const lbA,
                  const double* const ubA,
                  int&                nWSR,
                  double* const       cputime
                  );
```

which takes the usually positive definite Hessian matrix $H \in \mathbb{R}^{\mathrm{nV} \times \mathrm{nV}}$, the gradient vector $g \in \mathbb{R}^{\mathrm{nV}}$, the constraint matrix $A \in \mathbb{R}^{\mathrm{nC} \times \mathrm{nV}}$ the lower and upper bound vectors $\mathrm{lb}, \mathrm{ub} \in \mathbb{R}^{\mathrm{nV}}$ and the lower and upper constraints' bound vectors $\mathrm{lb}A, \mathrm{ub}A \in \mathbb{R}^{\mathrm{nC}}$ of the initial quadratic program. Equality constraints are imposed be setting the corresponding entries of lower and upper (constraints') bounds vectors to the same value.
All these data must be stored in arrays of type `double` (matrices stored row-wise, i.e. C style, in an one-dimensional array) with appropriate dimensions. If there are, for example, no upper bounds in your QP formulation, you can pass a null pointer instead of vector lb[1].
All `init` functions make deep copies of all arguments, thus afterwards you have to free their memory yourself.
The function `init` initialises all internal data structures, e.g. matrix factorisations, and *solves* the first quadratic program using the initial homotopy idea of the online active set strategy. The integer argument `nWSR` specifies the maximum number of working set recalculations to be performed during the initial homotopy (on output it contains the number of working set recalculations actually performed!). If `cputime` is not the null pointer, it contains the *maximum allowed* CPU time in seconds for the whole initialisation (and the actually required one on output). See Section 5.3 for further details.

The function `init` returns a status code (of type `returnValue`) which indicates whether the initialisation was successful; possible values are:

- `SUCCESSFUL_RETURN`: initialisation successful (including solution of first QP),
- `RET_MAX_NWSR_REACHED`: initial QP could not be solved within the given number of working set recalculations,

[1] If your QP does not comprise constraints (apart from bounds), you should make use of a special variant for simply bounded QPs (cf. Chapter 4).

- RET_INIT_FAILED (or a more detailed error code): initialisation failed.

If init indicates a SUCCESSFUL_RETURN, several functions enable you to obtain information about the solution of the first QP. The most important ones are:

- `returnValue getPrimalSolution( double* const xOpt ) const`

  that writes the optimal primal solution vector (dimension: nV) into the array xOpt, which has to be allocated (and freed) by the user;

- `returnValue getDualSolution( double* const yOpt ) const`

  that writes the optimal dual solution vector[2] (dimension: $\mathrm{nV} + \mathrm{nC}$) into the array yOpt, which has to be allocated (and freed) by the user;

- `double getObjVal( ) const`

  that returns the optimal objective function value.

Summary of the second step:
Having created an QProblem object example, it can be initialised together with solving the first QP with the following command: `example.init( H,g,A,lb,ub,lbA,ubA,nWSR,cputime );`

Solution of the Following QPs

If not only a single quadratic program but a whole sequence of QPs shall be solved—as it is the usual situation for an MPC problem—the next QP can be solved using the function:

```
returnValue hotstart( const double* const g_new,
                      const double* const lb_new,
                      const double* const ub_new,
                      const double* const lbA_new,
                      const double* const ubA_new,
                      int&                nWSR,
                      double* const       cputime
                      );
```

The next QP is specified by passing its gradient vector g_new, its lower and upper bound vectors lb_new and ub_new as well as its lower and upper constraints' bound vectors lbA_new and ubA_new (QP matrices are assumed to be constant). It is solved by means of the online active set strategy using at most nWSR working set recalculations or at most cputime seconds of CPU time (if not null). On output nWSR and cputime contain the number of

[2] We use the following definition of the Lagrange function to define the dual multipliers:

$$Hx^{\mathrm{opt}} + g(w_0) - A^T y^{\mathrm{opt}} = 0 \quad \Longleftrightarrow \quad \mathtt{H \cdot x + g - y[0...nV\text{-}1] - A^T y[nV...nV+nC\text{-}1]} = 0$$

The dual solution vector contains exactly one entry per lower/upper bound as well as exactly one entry per lower/upper constraints' bound. Positive entries correspond to active lower (constraints') bounds, negative entries to active upper (constraints') bounds and a zero entry means that both corresponding (constraints') bounds are inactive.

working set recalculations that were actually performed and the actually required CPU time for solving the next QP, respectively.

Like most qpOASES functions, hotstart returns a status code; possible values are:

- SUCCESSFUL_RETURN: QP has been solved,
- RET_MAX_NWSR_REACHED: QP could not be solved within the given number of working set recalculations,
- RET_HOTSTART_FAILED (or a more detailed error code): QP solution failed.

Summary of the third step:
Having created and initialised a QProblem object example, the next QP can be solved as follows: `example.hotstart( g_new,lb_new,ub_new,lbA_new,ubA_new,nWSR,cputime );`

3.3 A Tutorial Example

A complete example for solving two very simple quadratic programs using qpOASES is given in the file <install-dir>/EXAMPLES/example1.cpp:

```
#include <QProblem.hpp>

int main( )
{
    using namespace qpOASES;

    /* Setup data of first QP. */
    double H[2*2] = { 1.0, 0.0, 0.0, 0.5 };
    double A[1*2] = { 1.0, 1.0 };
    double g[2] = { 1.5, 1.0 };
    double lb[2] = { 0.5, -2.0 };
    double ub[2] = { 5.0, 2.0 };
    double lbA[1] = { -1.0 };
    double ubA[1] = { 2.0 };

    /* Setup data of second QP. */
    double g_new[2] = { 1.0, 1.5 };
    double lb_new[2] = { 0.0, -1.0 };
    double ub_new[2] = { 5.0, -0.5 };
    double lbA_new[1] = { -2.0 };
    double ubA_new[1] = { 1.0 };

    /* Setting up QProblem object. */
    QProblem example( 2,1 );

    /* Solve first QP. */
    int nWSR = 10;
    example.init( H,g,A,lb,ub,lbA,ubA, nWSR,0 );
```

```
    /* Solve second QP. */
    nWSR = 10;
    example.hotstart( g_new,lb_new,ub_new,lbA_new,ubA_new, nWSR,0 );

    return 0;
}
```

In order to access the functionality of the qpOASES software package via the `QProblem` class, the header file `QProblem.hpp` is included.
The `main` function starts with defining the data of two very small-scale QPs. Afterwards, a `QProblem` object is created which is then initialised together with solving the first QP. Finally, the `hotstart` function is used to solve the second QP.
You might wonder about the command `using namespace qpOASES;` at the very top of the `main` function. It is used because all classes, global functions and variables of the qpOASES software package are *collected in a common namespace* that is called qpOASES, too.

3.4 Setting Up Your Own Example

The easiest way for setting up your own example, say `yourexample`, is to use an existing one as a template. In doing so, perform the following steps:

1. *Copy the existing example:*

   ```
   cd <install-dir>/EXAMPLES
   cp example1.cpp yourexample.cpp
   ```

2. *Edit the examples* `Makefile`:

   Open the file `<install-dir>/EXAMPLES/Makefile` and add a new target

   ```
   yourexample:  yourexample.o
       ${CPP} -o $@ ${CPPFLAGS} $@.o -L${LIBS_PATH} -l${QPOASES_LIB}
   ```

   (Do not forget to add its name to the `all` target.)

3. *Implement your own example:*

   Modify your file `<install-dir>/EXAMPLES/yourexample.cpp` and run `make`. An executable called `yourexample` should be at your service.

Chapter 4

Different Solution Variants for Special QP Types

qpOASES is a structure-exploiting active-set QP solver. This chapter details how to most efficiently solve your QPs with qpOASES by choosing a solution variant that matches best your specific problem type. For this purpose, three different QProblem-like classes with overloaded constructors are available.

4.1 Solving QPs in Standard Form

Usually qpOASES expects QPs to be formulated in the following *standard form*:

$$\begin{aligned} \min_{x} \quad & \tfrac{1}{2} x^T H x + x^T g(w_0) \\ \text{s.t.} \quad & \mathrm{lb}A(w_0) \leq Ax \leq \mathrm{ub}A(w_0)\,, \\ & \mathrm{lb}(w_0) \leq \ x \ \leq \mathrm{ub}(w_0)\,, \end{aligned}$$

with a positive definite Hessian matrix H. If your QP is given in exactly this form, you should simply make use of the standard QProblem class as described in Chapter 3. Otherwise, solving your QP is also possible or can be done more efficiently if

- also your QP matrices H and/or A are varying from one QP to the next by using the SQProblem class (see Section 4.2);
- your QP formulation does *not* comprise constraints involving a matrix A by using the QProblemB class (see Section 4.3);
- your Hessian matrix H is not positive definite but only positive semi-definite by using a dedicated constructor (see Section 4.4);
- your Hessian matrix H is zero, i.e. your QP is actually a linear program, by using a dedicated constructor (see Section 4.5);
- your Hessian matrix H happens to be the identity matrix by using a dedicated constructor (see also Section 4.5).

4.2 Solving QPs with Varying Matrices

Although the online active set strategy was originally designed for QP sequences with fixed Hessian and constraint matrices, it can be easily extended to the case where also these matrices vary from QP to the next (see [5] for a mathematical description of this idea).
In order to use this extension, two modifications are necessary:

1. Create an instance of the `SQProblem` class (instead of one of type `QProblem`) by
   - including the header file `SQProblem.hpp` (instead of `QProblem.hpp`) and
   - use the constructor of the `SQProblem` class and a suitable `init` function, both take *exactly the same* arguments as those of the `QProblem` class.
2. Call the modified function

   ```
   returnValue hotstart( const double* const H_new,
                         const double* const g_new,
                         const double* const A_new,
                         const double* const lb_new,
                         const double* const ub_new,
                         const double* const lbA_new,
                         const double* const ubA_new,
                         int&                nWSR,
                         double* const       cputime
                         );
   ```

   for transition from one QP to the next; it also takes the new Hessian `H_new` as well as the new constraint matrix `A_new` as arguments.

A complete example for using the `SQProblem` class can be found within the file `<install-dir>/EXAMPLES/example1a.cpp`.

4.3 Solving Simply Bounded QPs

We call a quadratic program "simply bounded" whenever it does not comprise constraints but only bounds:

$$\begin{aligned} \min_{x} \quad & \tfrac{1}{2} x^T H x + x^T g(w_0) \\ \text{s.t.} \quad & \mathrm{lb}(w_0) \leq \ x \ \leq \mathrm{ub}(w_0) \,. \end{aligned}$$

This special form can be exploited within the solution algorithm for speeding up the computation, typically by a factor of three to five. Therefore, the qpOASES software package implements the special class `QProblemB` for solving simply bounded QPs.
In order to make use of this feature do the following:

1. You have to create a `QProblemB` object using the following constructor

   ```
   QProblemB( int nV );
   ```

2. Afterwards you can initialise the QProblemB object together with solving the first simply bounded QP by calling, for example,

```
returnValue init( const double* const H,
                  const double* const g,
                  const double* const lb,
                  const double* const ub,
                  int&                nWSR,
                  double* const       cputime
                  );
```

   The only difference from the QProblem class is the fact that the arguments specifying the constraints—i.e. A, lbA, ubA, and nC—are missing.

3. For solving the next problem within your QP sequence, the following variant of the hotstart function is available:

```
returnValue hotstart( const double* const g_new,
                      const double* const lb_new,
                      const double* const ub_new,
                      int&                nWSR,
                      double* const       cputime
                      );
```

   Again, it takes exactly the same arguments as the corresponding QProblem member function except for the two arguments lbA_new, ubA_new.

A complete example for using the QProblemB class can be found within the file <install-dir>/EXAMPLES/example1b.cpp.

4.4 Solving QPs with Positive Semi-Definite Hessian Matrix

qpOASES provides a built-in regularisation procedure for dealing with semi-definite Hessian matrices. This procedure first adds a small multiple of the identity matrix[1] to the Hessian and solves the corresponding regularised QP. Afterwards, a few post-iterations[2] are performed that improve solution accuracy signigificantly over a plain regularisation at virtually now extra computational cost.

If your QP involves a Hessian matrix that is only positive semi-definite, this regularisation scheme *is used automatically* (i.e. without any change in the constructor or other function calls) as semi-definiteness can be easily detected. However, if your Hessian is only positive semi-definite, this causes a certain computational overhead[3] that can be avoided by a dedicated constructor call, e.g.,

[1] given by the global constant EPS_FOR_REGULARISATION
[2] given by the global constant MAX_NUMBER_OF_REGULARISATION_STEPS
[3] an additional Cholesky decomposition

```
QProblem( int nV, int nC, HessianType hessianType );
```

Therein, hessianType can take one of the following values:

- HST_POSDEF: Hessian matrix is positive definite,
- HST_SEMIDEF: Hessian matrix is positive semi-definite,
- HST_ZERO: Hessian matrix is zero matrix (see next section),
- HST_IDENTITY: Hessian matrix is identity matrix (see next section).

If hessianType is set to HST_SEMIDEF or HST_ZERO, the built-in regularisation scheme is switched on at no additional computational cost. Corresponding overloaded constructors also exist for the SQProblem and QProblemB class, respectively.

When using the built-in regularisation procedure, the default values for the regularisation parameter and the maximum number of post-iterations are taken from the global constants EPS_FOR_REGULARISATION and MAX_NUMBER_OF_REGULARISATION_STEPS, respectively. However, it is possible to change these values at runtime by calling the following member functions:

- `returnValue setEpsForRegularisation( double _eps );`
- `returnValue setMaxNumOfRegularisationSteps( int _nRegSteps );`

Both functions expect positive arguments, though the latter function also accepts the argument value 0 which effectively disables the built-in regularisation procedure.

4.5 Solving QPs with Trivial Hessian Matrix

Whenever a Hessian matrix is passed to qpOASES, i.e. when calling a init function or performing a hotstart while using the SQProblem class, it is internally checked whether the Hessian is trivial. It is considered trivial if and only if it is the zero or identity matrix, corresponding to HST_ZERO or HST_IDENTITY as mentioned in the previous section. If the Hessian is trivial, several simplifications of the internal linear algebra operations apply, cutting computational load by about a factor of two.

If your Hessian is trivial, you might explicitly provide this information to qpOASES via a dedicated constructor call, e.g.,

```
QProblem( int nV, int nC, HessianType hessianType );
```

(corresponding overloaded constructors also exist for the SQProblem and QProblemB class, respectively). If you set hessianType to HST_ZERO or HST_IDENTITY, no internal memory for storing the Hessian matrix is allocated. Moreover, when doing so you are allowed to pass a null pointer as argument within all function calls involving the Hessian matrix, e.g.,

```
// assumes that a QProblem object "qp" exists
qp.init( 0,g,A,lb,ub,lbA,ubA,nWSR,cputime );
```

A null pointer is then interpreted as zero or identity matrix, respectively. Whenever you pass a non-null argument, a full Hessian matrix is expected and its type is automatically determined internally.

Solving Linear Programming (LP) Problems

The regularisation scheme mentioned in Section 4.4 in principle also allows to solve linear programming (LP) problems by means of qpOASES. However, qpOASES *is not a dedicated (parametric) LP solver*, thus using it for solving LPs might be highly inefficient due to the dense linear algebra and also might fail in certain circumstances! Therefore, this additional feature should be only used for small-scale LPs (comprising, say, a hundred variables) and in situations where computational time is not the main concern.

A complete example for solving two small-scale LPs with qpOASES can be found within the file <install-dir>/EXAMPLES/exampleLP.cpp.

Solving QPs whose Hessian is the Identity Matrix

Via a coordinate transformation, every strictly convex QP can be transformed into an equivalent one whose Hessian is the identity matrix. Also ℓ_2-norm minimisation problems naturally pose QPs whose Hessian is the identity matrix. Thus, it is possible to provide such QP sequence to qpOASES by specifying the Hessian type to be HST_IDENTITY within the above-mentioned constructor call; all other function calls remain unaltered.

Chapter 5

Advanced Functionality

5.1 Obtaining Status Information

There are many functions for obtaining status information on the current iterate. Firstly, you can obtain the primal and dual iterate as well as the corresponding objective function value by using, respectively:

- `returnValue getPrimalSolution( double* const xOpt ) const`,
- `returnValue getDualSolution( double* const yOpt ) const`,
- `double getObjVal( ) const`.

If you wonder why these are the same functions as for obtaining the optimal solution after a QP has been solved (cf. Section 3.2), you should recall that qpOASES uses a homotopy for solving the current QP that produces a sequence of iterates that are *optimal for intermediate QPs* along the homotopy path.
The first two functions expect an allocated `double` array and store the optimal solution vector if and only if the (intermediate) QP has been solved; otherwise the error code `RET_QP_NOT_SOLVED` is returned. The function `getObjVal( )` calculates and returns the optimal objective function value or returns `INFTY` if the (intermediate) QP has not been solved.

Secondly, you can ask for the total number of variables and constraints and for the cardinality of certain subsets (at current iterate!) of them:

- `int getNV( ) const`: returns number of variables,
- `int getNFR( ) const`: returns number of free variables,
- `int getNFX( ) const`: returns number of fixed variables,
- `int getNC( ) const`: returns number of constraints,
- `int getNEC( ) const`: returns number of (implicitly defined) equality constraints,
- `int getNAC( ) const`: returns number of active constraints,

- `int getNIAC( ) const` : returns number of inactive constraints.

  Moreover,

- `int getNZ( ) const` : returns dimension of the null space of active constraints.

Finally, you can ask for the overall status of the QP (object):

- `BooleanType isInitialised( ) const` : returns `BT_TRUE` if and only if the QP object has been initialised,
- `BooleanType isSolved( ) const` : returns `BT_TRUE` if and only if QP has been solved,
- `BooleanType isInfeasible( ) const` : returns `BT_TRUE` if and only if QP was found to be infeasible.

5.2 Initialised Homotopy

For solving a QP, `qpOASES` always starts at the optimal solution of the previous QP and performs a homotopy to the optimal solution of the QP to be solved. At the very beginning of a sequence (when `init` is called) an auxiliary QP is constructed internally whose optimal solution is known. This optimal solution serves as a starting point for the homotopy to the optimal solution of the (actual) initial QP. By default, this auxiliary QP has the origin as solution and its active set is empty (or comprising implicitly fixed variables and equality constraints only).

The notion *initialised homotopy* refers to the possibility to incorporate an initial guess for the optimal solution or the active set at the solution into the construction of the auxiliary QP. This is done by calling a special variant of the `init` function:

```
returnValue init( const double* const      H,
                  const double* const      g,
                  const double* const      A,
                  const double* const      lb,
                  const double* const      ub,
                  const double* const      lbA,
                  const double* const      ubA,
                  int&                     nWSR,
                  double* const            cputime,
                  const double* const      xOpt,
                  const double* const      yOpt,
                  const Bounds* const      guessedBounds,
                  const Constraints* const guessedConstraints
                  );
```

Besides the arguments of the usual `init` function, it (optionally) takes guesses for the primal solution vector `xOpt`, the dual solution vector `yOpt` or the status (active/inactive) of bounds and constraints at the solution (see below). Null pointers can be passed for all of these

arguments. The construction of the auxiliary QP now depends on the arguments passed (for convenience we summarise guessedBounds and guessedConstraints to guess which is null if and only if both parts are null) as follows:

1. xOpt == 0, yOpt == 0, guess == 0: start at primal/dual origin with empty active set (usual auxiliary QP setup);

2. xOpt != 0, yOpt == 0, guess == 0: start at primal/dual origin and determine active set by "clipping"[1];

3. xOpt == 0, yOpt != 0, guess == 0: start with primal variables equal to zero, dual variables equal to given vector and determine active set from signs of dual variables;

4. xOpt == 0, yOpt == 0, guess != 0: start at primal/dual origin and with given active set;

5. xOpt != 0, yOpt != 0, guess == 0: start with given vectors for primal and dual variables and determine active set from signs of dual variables;

6. xOpt != 0, yOpt == 0, guess != 0: start with primal variables equal to given vector, dual variables equal to zero and with given active set;

7. xOpt != 0, yOpt != 0, guess != 0: start with given vectors for primal and dual variables and with given active set (assume them to be consistent!).

The remaining eighth combination is not allowed for consistency reasons.

Besides initialising the homotopy at startup of the QP sequence, it is also possible to incorporate an initial guess for the active set when calling the hotstart function:

```
returnValue hotstart( const double* const      g_new,
                      const double* const      lb_new,
                      const double* const      ub_new,
                      const double* const      lbA_new,
                      const double* const      ubA_new,
                      int&                     nWSR,
                      double* const            cputime,
                      const Bounds* const      guessedBounds,
                      const Constraints* const guessedConstraints
                      );
```

In this case only the active set can be specified, primal and dual solution vectors are always taken from the previous QP solution. This hotstart variant updates the active set according to the user's guess and performs a usual homotopy afterwards.

[1] i.e. add all bounds and constraints to active set that are violated for given primal solution vector

Specifying an Initial Guess for the Active Set

For specifying an initial guess for the active set, you have to setup a Bounds and/or Constraints object. This can either be done from scratch or by modifying an exisiting one. For the first variant you might use the following code fragment:

```
// assumes that a QP object "qp" exists
int nV = qp.getNV( );
int nC = qp.getNC( );

Bounds guessedBounds( nV );
guessedBounds.setupAllLower( );

Constraints guessedConstraints( nC );
guessedConstraints.setupAllInactive( );
```

First, a Bounds object comprising a working set of nV bounds is constructed and afterwards all bounds are set to be active at their lower limit. Second, a Constraints object is constructed analogously and all constraints are set to be inactive. For a Bounds object you can call one of the following functions:

- returnValue setupAllFree( ): all variables are free, i.e. bounds are inactive,
- returnValue setupAllLower( ): all variables are fixed at their lower limits,
- returnValue setupAllUpper( ): all variables are fixed at their upper limits.

For a Constraints object you can call one of the following functions:

- returnValue setupAllInactive( ): all constraints are inactive,
- returnValue setupAllLower( ): all constraints are active at their lower limits,
- returnValue setupAllUpper( ): all constraints are active at their upper limits.

Moreover, you might setup the status of each bound/constraint one by one by calling:

- returnValue setupBound(int number, SubjectToStatus status) or
- returnValue setupConstraint(int number, SubjectToStatus status),

repectively, where number specifies the number of the repective bound/constraint (starting at zero!) and status is one of the following types:

- ST_INACTIVE: bound/constraint is inactive,
- ST_LOWER: bound/constraint is active at its lower limit,
- ST_UPPER: bound/constraint is active at its upper limit.

Please note that you can call *either* exactly one setupAll* variant *or* exactly one of setupBound/setupConstraint for each single bound/constraint!

Instead of setting up a Bounds/Constraints object from scratch, you might want to *modify an existing one*. For achieving this, you will most commonly first obtain a copy of the active set of the current QP by calling:

```
// assumes that a QP object "qp" exists
Bounds guessedBounds;
qp.getBounds( guessedBounds );

Constraints guessedConstraints;
qp.getConstraints( guessedConstraints );
```

Afterwards you might use one of the following functions to manipulate a `Bounds` object:

- `returnValue moveFixedToFree( int number )`: moves the `number`-th bound from the working set of fixed variables to that of free ones,
- `returnValue moveFreeToFixed( int number, SubjectToStatus status )`: moves the `number`-th bound from the working set of free variables to that of fixed ones (where `status` must be either `ST_LOWER` or `ST_UPPER`).

For a `Constraints` object you can call one of the following functions:

- `returnValue moveActiveToInactive( int number )`: moves the `number`-th constraint from the working set of active constraints to that of inactive ones,
- `returnValue moveInactiveToActive( int number, SubjectToStatus status )`: moves the `number`-th constraint from the working set of inactive constraints to that of active ones (where `status` must be either `ST_LOWER` or `ST_UPPER`).

Moreover, in the model predictive control context it is very common that the active set is *shifted* between two consecutive sampling instants. Therefore, for both `Bounds` and `Constraints` you can also call one of the following functions:

- `returnValue shift( int offset )`: shifts forward the working set of bounds/constraints by a given offset (which has to be an integer divisor of the total number of bounds/constraints), i.e. the status information of the first `offset` bounds/constraints is thrown away and the one of the last `offset` ones is duplicated;
- `returnValue rotate( int offset )`: rotates forward the working set of bounds/constraints by a given offset.

We refer to the DOXYGEN documentation (cf. installation step six described in Chapter 2) for more details.

5.3 Specifying a CPU Time Limit for QP Solution

For all `init` and `hotstart` function calls the input argument `nWSR` is mandatory. Additionally, it is possible to specify a maximum amount of CPU time to be spent on the respective QP solution. For doing so, a non-null pointer to a `double` containing the maximum allowed CPU time in seconds needs to be specified. If both, a maximum number of working set recalculations `nWSR` and a maximum allowed CPU time `cputime` is given, the solution procedure stops as soon as *one of these limits* is reached, whatever may occur first.

The CPU time limitation is based on a *heuristic* that estimates the required CPU time for the next working set change; if there is not enough time left, the solution procedure stops. This heuristic is based on the CPU time measurements of the previous working set changes, thus the actual total CPU time might be slightly higher that the allowed one due to time measurement inaccuracies. However, it is guaranteed that *at most one* working set change too much is performed.

Note that the CPU timit limit only can take effect if a system clock is available via the global `getCPUtime` function (implemented within the file `SRC/Utils.cpp`).

5.4 Speeding-Up Solution for QPs Comprising Many Constraints

Heuristic for Approximating the Constraint Product

In case the QP comprises much more constraints than optimisation variables, the step length determination requires a major part of the overall computational load per QP iteration. That is because the (costly) matrix-vector product Ax has to be formed for determining if an inactive constraint is going to become active at the next iterate.
qpOASES has implemented a strategy that only *approximates* this matrix-vector product when inactive constraints are so far off their limits that they cannot become active during the next step. This strategy still ensures exact QP solution and usually leads to considerable computational savings. However, in worst-case it can even prolong computation time, thus it needs to be explicitly enabled by defining the compiler flag `__MANY_CONSTRAINTS__`. Note, that this strategy relies on the fact that *each constraint has ℓ_1-norm not greater than 1*! Thus, before setting this compiler flag, you might need to re-scale your constraints (otherwise QP solution can fail!).

Specifying a Function for Evaluating the Constraints

Another possibility to speed-up QP solution in case of many constraints is available whenever the calculation of the matrix-product of the constraint matrix `A` with the current primal iterate `x` can be simplified. In that case, the user can provide a dedicated function that can evaluate the product of any constraint at a given primal iterate. Once such a function is specified and passed to an QP object, qpOASES will use this user-provided function for calculating the constraint products instead of doing a standard (but possibly naive) matrix-vector multiplication.

For using this functionality, you have to perform the following steps:

1. Derive a customized class from the abstract base class `ConstraintProduct` as declared within `<install-dir>/INCLUDE/ConstraintProduct.hpp`. Within this class, you have to implement the function operator which has the following form:

```
virtual int operator()( int constrIndex,
                        const double* const x,
                        double* const constrValue
                        ) const;
```

It takes the index of the constraint to be evaluated (between 0 and nC) and an array containing the current primal iterate (of size nV) as input arguments and writes the corresponding product into constrValue. The function operator needs to return 0 on success and might return an error code otherwise.

2. Make this derived class available within your example, instantiate an object of this class and pass it to the QP object by calling

```
// assumes that a QP object "qp" exists
MyConstraintProduct myCP( );
qp.setConstraintProduct( &myCP );
```

A full tutorial example illustrating this feature of qpOASES can be found within the file <install-dir>/EXAMPLES/example4.cpp.

5.5 Further Useful Functionality

Reading Data From Files

Both the init and the hotstart functions are overloaded with variants that are able to read the required data directly from a plain ASCII file, e.g.:

- ```
  returnValue init( const char* const H_file,
                    const char* const g_file,
                    const char* const A_file,
                    const char* const lb_file,
                    const char* const ub_file,
                    const char* const lbA_file,
                    const char* const ubA_file,
                    int&              nWSR,
                    double* const     cputime
                    );
  ```

- ```
  returnValue hotstart( const char* const g_file,
                        const char* const lb_file,
                        const char* const ub_file,
                        const char* const lbA_file,
                        const char* const ubA_file,
                        int&              nWSR,
                        double* const     cputime
                        );
  ```

Instead of a double array, they expect a string with the name of the ASCII file containing the respective data. Data files must be stored row-wise; all entries within one row should be space- or tabulator-separated.
These variants also exists for the case when an initial guess for the active set is provided (as described in Section 5.2).

Output Settings

You can adjust the text output of qpOASES using the following functions:

- `PrintLevel getPrintLevel( ) const`,
- `void setPrintLevel( PrintLevel _printlevel )`.

The function `getPrintLevel` returns one of the following print levels:

- `PL_NONE`: no output at all,
- `PL_LOW`: print error messages only,
- `PL_MEDIUM`: print error messages, warnings, some info messages as well as a concise iteration summary (default value),
- `PL_HIGH`: print all messages that occur while iterating.

By means of the function `setPrintLevel` you can specify one of the above-mentioned print levels whenever desired.

Resetting a `QProblem` Object

Sometimes it can be useful to reset an existing `QProblem` object. This is particularly helpful if you want to restart while solving a QP sequence (e.g. after an internal error has occured) without creating a new object. This feature is provided by the following function:

```
returnValue reset( );
```

It resets all internal data structures and matrix factorisations and thus leaves the `QProblem` object in exactly the same state as it would be after a constructor call. Therefore, you need to call an `init` function for solving the first QP after an execution of `reset`.

Disabling/Enabling Constraints

For the special situation where some of the constraints leave (and possibly re-enter) the QP formulation, there exist the following two functions:

- `returnValue disableConstraint( const Indexlist* const numbers );`
- `returnValue enableConstraint( const Indexlist* const numbers );`

Both take a list (of internal data type `Indexlist`[2]) containing the indices of the constraints to be disabled/enabled. A disabled constraint will be automatically removed from the QP formulation (incorporated into the usual homotopy framework); by enabling it afterwards, it will be added to the QP formulation again.

Remark: This functionality has not been sufficiently tested yet.

[2] Please consult the DOXYGEN documentation on how to setup an `Indexlist` object.

Printing QP Properties

At any time you might print a concise list of properties of the QP object by calling:

```
returnValue printProperties( );
```

Besides other information, it displays number and type of bounds and constraints, respectively, the type of the Hessian matrix as well as the status of the QP object.

5.6 Add-Ons for qpOASES

When compiling the source code of qpOASES, a second library libqpOASESextras.a is created. Its functionality comprises all the functionality of the standard libqpOASES.a and additionally provides several add-ons which are described in the following subsections. Header and implementation files of these add-ons are located within a sub-folder EXTRAS of INCLUDE and SRC, respectively.

5.6.1 Solution Analysis

For a posteriori analysis of a QP solution the SolutionAnalysis class is provided as an add-on to qpOASES. Currently it implements the following two functions:

- Determination of the maximum violation of the KKT optimality conditions:

  ```
  returnValue getMaxKKTviolation( QProblem* qp,
                                  double& maxKKTviolation
                                  ) const;
  ```

  This function takes a pointer to a QProblem object which is assumed to have readily solved an (intermediate) QP and writes the maximum violation of the KKT optimality conditions into the argument maxKKTviolation. If the QProblem object has not solved the current QP, the status code RET_UNABLE_TO_ANALYSE_QPROBLEM is returned.

- Computation of the variance-covariance matrix of the QP output for uncertain inputs:

  ```
  returnValue getVarianceCovariance( QProblem* qp,
                                     double* g_b_bA_VAR,
                                     double* Primal_Dual_VAR
                                     ) const;
  ```

  It also takes a QProblem object which is assumed to have readily solved an (intermediate) QP as well as the variance-covariance of the gradient, the bounds and the constraints' bounds, respectively (matrix dimension: $2\mathtt{nV}+\mathtt{nC} * 2\mathtt{nV}+\mathtt{nC}$). The variance-covariance matrix of the primal and dual variables is written into the argument Primal_Dual_VAR (matrix dimension: $2\mathtt{nV}+\mathtt{nC} * 2\mathtt{nV}+\mathtt{nC}$), which needs to be allocated by the user.

For using the SolutionAnalysis class you need to include its header SolutionAnalysis.hpp into your source file, a complete example can be found in the file <install-dir>/EXAMPLES/example2.cpp.

5.6.2 Solving Test Problems from the Online QP Benchmark Collection

A second qpOASES add-on is intended to facilitate the solution of test problems from the Online QP Benchmark Collection [1]. Data for a whole QP sequence with constant matrices along with its optimal primal/dual solution vectors and the optimal objective function value is stored in plain ASCII files. For conveniently reading these files, three functions are provided (see <install-dir>/INCLUDE/EXTRAS/OQPinterface.hpp for a detailed documentation):

- readOQPdimensions for reading the dimensions of the QP sequence,
- readOQPdata for reading data and solution information of the QP sequence,
- solveOQPbenchmark for solving a given benchmark QP sequence.

Moreover, the following function summarises the functionality of the three above-mentioned ones:

```
returnValue runOQPbenchmark( const char* path,
                             int& nWSR,
                             double& maxCPUtime,
                             double& maxPrimalDeviation,
                             double& maxDualDeviation,
                             double& maxObjDeviation
                             );
```

It takes the path to the directory where the benchmark problem is stored as well as the maximum number of working set recalculations per QP as input arguments. On output nWSR contains the maximum number of working set recalculations that have been actually performed, maxCPUtime contains the maximum CPU time that have been required for solving each of the QPs and maxPrimalDeviation, maxDualDeviation, maxObjDeviation contain the maximum primal, dual and objective function value deviation (ℓ_∞-norm) from the given optimal QP solutions, respectively.

For using this add-on you need to include the header file OQPinterface.hpp into your source code, a complete example can be found in the file <install-dir>/EXAMPLES/example3.cpp. In order to run this example, you need to download example no. 01 from the Online QP Benchmark Collection website [1] first and extract its archive into the sub-folder <install-dir>/EXAMPLES/chain80w/.

Chapter 6

Interfaces for Third-Party Software

If you want to use qpOASES via one of the following third-party interfaces, make sure that you have performed the installation steps 1 through 3 from Chapter 2. Afterwards, proceed with the installation of the desired interface as described in this chapter.

6.1 Interface for Matlab

Installation

It is possible to use qpOASES directly within the MATLAB environment. This is facilitated by compiling it into a so-called MEX function, which can be done as follows:

1. Start MATLAB and run `mex -setup` for choosing a C++ compiler (e.g. gcc).

2. Execute the following commands:

   ```
   cd <install-dir>/INTERFACES/MATLAB
   make
   ```

   The latter command runs the MATLAB script `make.m` which does the compilation. Executables `qpOASES.<ext>`, `qpOASES_sequence.<ext>`, `qpOASES_sequenceSB.<ext>` and `qpOASES_sequenceVM.<ext>` should be created, where `<ext>` (e.g. `mexglx`) depends on your operating system.

Remarks:

- The compilation was tested under LINUX using MATLAB 7.3 and the gcc compiler. Modifications of the `make.m` script might be necessary depending on your operating system, your MATLAB version and your compiler. For compiling the MATLAB interface for WINDOWS operating systems, the BORLAND BCC 5.5 compiler (available at `http://www.codegear.com`) has proven to work.

- If compilation fails due to the fact that the `snprintf()` function is not supported, you might uncomment line 54 within `<install-dir>/INCLUDE/MessageHandling.hpp` and try to compile again.

Interface for Solving a Single QP

After a successful installation, you can call qpOASES as conventional QP solver from the MATLAB environment (using a cold start every time):

```
[ obj,x,y,status,nWSRout ] = qpOASES( H,g,A,lb,ub,lbA,ubA,{nWSR,{x0}} )
```

This command combines the creation of a QProblem object and a calls to the function init (see Chapter 3): the *input arguments*[1] specify the Hessian matrix, the gradient vector, the constraint matrix, the lower and upper bound vectors, the lower and upper constraints' vectors, respectively. Again, the Hessian has to be symmetric and positive definite and all vectors must be stored as column vectors. Moreover, the maximum number of working set recalculations and an initial guess for the primal solution (cf. Section 5.2) can be specified optionally (either both or only x0 can be left away).
If the input argument nWSR is not specified, the default value $5 * (\mathtt{nV} + \mathtt{nC})$ is chosen. If no initial guess is given, the usual homotopy starting at the origin is performed. Furthermore, it is possible to leave one or more of the input arguments lb, ub, lbA, ubA empty if your QP formulation does not comprise the corresponding bounds.
The *output arguments* contain the optimal objective function value, the optimal primal solution vector, the optimal dual solution vector, a status flag, and the number of working set recalculations actually performed, respectively. The status flag can take the following values:

- 0: QP was solved,
- 1: QP could not be solved within the given number of working set recalculations,
- -1: QP solution failed.

If you do not need all output information, you can leave all but the first one away, e.g.

```
[ obj,x ] = qpOASES( H,g,A,lb,ub,[],ubA )
```

Remark: The function qpOASES also allows you to solve a pre-computed sequence of QPs with fixed matrices: you just have to pass a whole sequence of input vectors. Each vector must be stored column-wise in a matrix, i.e. the ith QP is given by the ith columns of the QP "vectors" g, lb, ub, lbA, ubA, and all these five matrices must have the same number of columns. As both the Hessian and the constraint matrix remain constant, they are passed as in the case of a single QP. If a whole sequence of QPs is to be solved, also the outputs are given column-wise, i.e. obj is a row vector, x is a matrix with optimal primal solution vectors stored column-wise inside, and so on.

The interface allows you to directly use the QProblemB class for simply bounded QPs (cf. Section 4.3) by simply leaving the arguments A, lbA, ubA away:

```
[ obj,x,y,status,nWSRout ] = qpOASES( H,g,lb,ub,{nWSR,{x0}} )
```

Again, a default value for the number of working set recalculations is used (here $5 * \mathtt{nV}$) if nWSR is not specified; and you can leave lb or ub empty if they do not occur within your QP formulation.

[1] all matrices have to be passed in dense format

Interface for Solving a QP Sequence

As the online active set strategy is intended to solve a whole sequence of parameterised QPs, there exist a special MATLAB function for hotstarting each QP from the solution of the previous one:

```
[ obj,x,y,status,nWSRout ] = qpOASES_sequence( 'i',H,g,A,lb,ub,lbA,ubA,{nWSR,{x0}} )
[ obj,x,y,status,nWSRout ] = qpOASES_sequence( 'h',g,lb,ub,lbA,ubA,{nWSR} )
                             qpOASES_sequence( 'c' )
```

As in the C++ implementation (cf. Chapter 3), the first QP of the sequence is solved together with the initialisation all internal data structures. For this purpose, the function qpOASES_sequence (called with first input argument `'i'`) takes all QP data and optionally the maximum number of working set recalcutations for solving the initial QP and an initial primal solution guess as further input arguments. It provides the usual output information (see above) and you can leave all but the first output argument away, again.
Afterwards, each subsequent QP can be solved by performing a so-called "hot start" using the function qpOASES_sequence, again (this time called with first input argument `'h'`). It takes the QP vectors of the new QP as well as the maximum number of working set recalcutations as further input arguments, and provides the usual output information.
Having solved the last QP of the sequence, you are encouraged to free the internal memory by calling `qpOASES_sequence( 'c' )`.

For solving QPs of special types as described in Chapter 4, special variants of the above function are provided: first, you can run the commands

```
[ obj,x,y,status,nWSRout ] = qpOASES_sequenceSB( 'i',H,g,lb,ub,{nWSR,{x0}} )
[ obj,x,y,status,nWSRout ] = qpOASES_sequenceSB( 'h',g,lb,ub,{nWSR} )
                             qpOASES_sequenceSB( 'c' )
```

for solving *s*imply *b*ounded QPs (input arguments correspoding to constraints are simply left away); second, call

```
[ obj,x,y,status,nWSRout ] = qpOASES_sequenceVM( 'i',H,g,A,lb,ub,lbA,ubA,{nWSR,{x0}} )
[ obj,x,y,status,nWSRout ] = qpOASES_sequenceVM( 'h',H,g,A,lb,ub,lbA,ubA,{nWSR} )
                             qpOASES_sequenceVM( 'c' )
```

for solving QPs with *v*arying *m*atrices, where qpOASES_sequenceVM also takes the new matrices of the next QP of the sequence. Again, the internal memory is freed by calling `qpOASES_sequenceSB( 'c' )` and `qpOASES_sequenceVM( 'c' )`, respectively. This memory is kept independently for all three QP types.

Examples

The files `example1.mat`, `example1a.mat` and `example1b.mat` contain, respectively, very basic examples for solving a sequence comprising two QPs with fixed matrices, varying matrices, and with simple bounds only. For solving the first one do the following:

1. Start MATLAB and execute the following commands:

   ```
   cd <install-dir>/INTERFACES/MATLAB
   load example1.mat
   ```

2. Solve the first QP by typing

   ```
   [ obj,x,y,status,nWSRout ] =
           qpOASES_sequence( 'i',H,g,A,lb,ub,lbA,ubA,10 )
   ```

3. Solve the second QP by typing

   ```
   [ obj,x,y,status,nWSRout ] =
           qpOASES_sequence( 'h',g_new,lb_new,ub_new,lbA_new,ubA_new,10 )
   ```

4. Free the internal memory by calling `qpOASES_sequence( 'c' )`.

6.2 Interface for Simulink

Installation

You can use qpOASES directly within the SIMULINK environment, too. This requires to compile it into a so-called S function, which can be done as follows:

1. Start MATLAB and run `mex -setup` for choosing a C++ compiler (e.g. gcc).

2. Execute the following commands:

   ```
   cd <install-dir>/INTERFACES/SIMULINK
   make
   ```

   The latter command runs the MATLAB script `make.m` which does the compilation. Three executables called `qpOASES_QProblemB.<ext>`, `qpOASES_QProblem.<ext>` and `qpOASES_SQProblem.<ext>` should be created, where `<ext>` (e.g. `mexglx`) depends on your operating system.

Remarks:

- The compilation was tested under LINUX using MATLAB 7.3 and the gcc compiler. Modifications of the `make.m` script might be necessary depending on your operating system, your MATLAB version and your compiler. For compiling the SIMULINK interface for WINDOWS operating systems, the BORLAND BCC 5.5 compiler (available at `http://www.codegear.com`) has proven to work.

- If compilation fails due to the fact that the `snprintf()` function is not supported, you might uncomment line 54 within `<install-dir>/INCLUDE/MessageHandling.hpp` and try to compile again.

Interface

There exist three different S function interfaces corresponding to the three different types of QP sequences to be solved (see also Chapter 4):

1. qpOASES_QProblemB.<ext> for solving simply bounded QPs,
2. qpOASES_QProblem.<ext> for solving QPs with fixed matrices,
3. qpOASES_SQProblem.<ext> for solving QPs with varying matrices.

For each of these interfaces a simple example is provided within the folder <install-dir>/INTERFACES/SIMULINK. We only give details for the one for QPs with fixed matrices, as the other ones work analoguously.

In order to run the example, start MATLAB and execute the corresponding script file as follows:

```
cd <install-dir>/INTERFACES/SIMULINK
load_example_QProblem
```

The sample QP data is loaded into the workspace and the file qpOASES_QProblem.mdl (depicted in Figure 6.1) is opened.

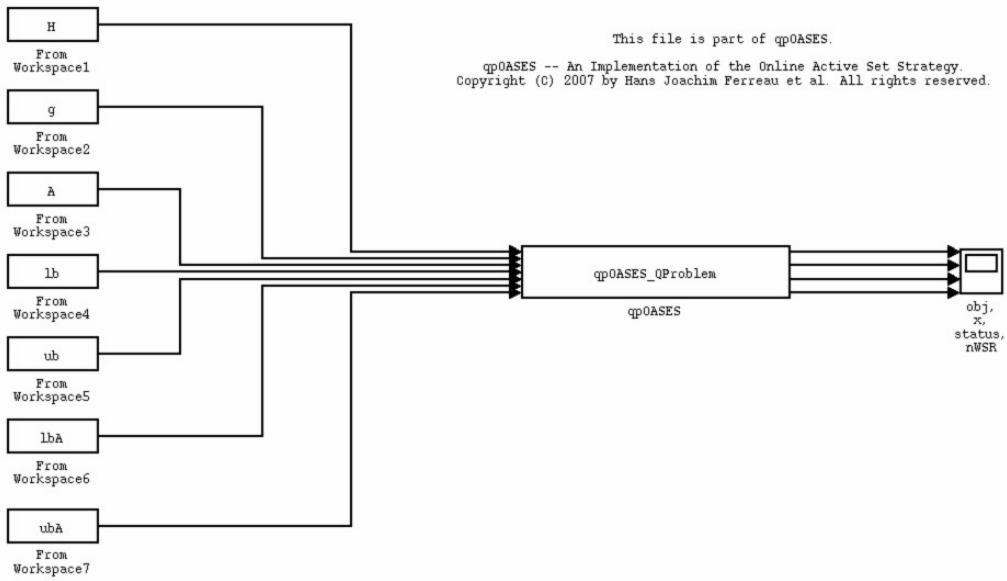

Figure 6.1: qpOASES working as SIMULINK S function.

The qpOASES S function has *seven inputs*:

- the (fixed) QP matrices H and A as well as
- the QP vectors g, lb, ub, lbA, ubA, which can be updated at each sampling instant.

The dimensions of the inputs are detected automatically, but they have to be consistent (e.g. the dimension of H needs to be the squared size of g).

Moreover, you have to *define three additional values* near top of the file qpOASES_QProblem.cpp *before compilation of the S function*:

- #define SAMPLINGTIME <value>: the sample time of the SIMULINK block,
- #define NCONTROLINPUTS <value>: the number of control inputs of your system (the leading NCONTROLINPUTS components of the optimal primal solution vector are returned as optimal output by the S function),
- #define NWSR <value>: the maximum number of working set recalculations to be performed per QP.

For running the example you can use the specified default values; but do not forget to adjust them to the requirements of your own problem.

At each sampling instant the qpOASES S function provides the following *four outputs*:

- obj: the optimal objective function value;
- x: the leading NCONTROLINPUTS components of optimal primal solution vector;
- status: a status flag which can take the values
  - 0: QP was solved,
  - 1: QP could not be solved within the given number of working set recalculations,
  - -1: QP solution failed;
- nWSR: the number of working set recalculations actually performed.

An Example

Having executed the script load_example_QProblem as described above, you can simply start the SIMULINK simulation given by the file example_QProblem.mdl. The simulation runs for 0.5 s with a sample time of 0.1 s. At the first two sampling instants the QPs as specified in the file example1.mat of the MATLAB interface are solved; at the remaining sampling instants the last QP is solved repeatedly (requiring zero iterations as the hotstart feature of the online active set strategy is used).

6.3 Interface for scilab

Installation

For using qpOASES within SCILAB, you have to perform the following steps:

1. Compile the SCILAB interface by executing the following commands:

   ```
   cd <install-dir>/INTERFACES/SCILAB
   make
   ```

2. Start SCILAB and link the interface to the SCILAB environment:

   ```
   exec qpOASESinterface.sce;
   ```

Interface for Solving a Single QP

If you simply want to use qpOASES as conventional QP solver (using a cold start every time), you can call it as follows:

```
[ obj,x,y,status,nWSRout ] = qpOASES( H,g,A,lb,ub,lbA,ubA,nWSR )
```

The *input arguments* specify the Hessian matrix, the gradient vector, the constraint matrix, the lower and upper bound vectors, the lower and upper constraints' vectors, and the maximum number of working set recalculations, respectively. As usual, the Hessian must be symmetric and positive definite and all vectors must be stored as column vectors.
The *output arguments* contain the optimal objective function value, the optimal primal solution vector, the optimal dual solution vector, a status flag, and the number of working set recalculations actually performed, respectively. The status flag can take the following values:

- `0`: QP was solved,
- `1`: QP could not be solved within the given number of working set recalculations,
- `-1`: QP solution failed.

If you do not need all output information, you can leave all but the first one away.

Remark: A special variant for simply bounded QPs is not yet interfaced.

Interface for Solving a QP Sequence

As the online active set strategy is intended to solve a whole sequence of parameterised QPs, there exist special routines for doing so:

```
[ obj,x,y,status,nWSRout ] = qpOASES_init( H,g,A,lb,ub,lbA,ubA,nWSR )
[ obj,x,y,status,nWSRout ] = qpOASES_hotstart( g,lb,ub,lbA,ubA,nWSR )
                             qpOASES_cleanup
```

As in the C++ implementation (cf. Chapter 3), the first QP of the sequence is solved together with the initialisation all internal data structures. For this purpose, the function `qpOASES_init` takes all QP data and the maximum number of working set recalcutations for solving the initial QP as input arguments, and provides the usual output information (see above).
Afterwards, each subsequent QP is can be solved by performing a so-called "hot start" using the function `qpOASES_hotstart`. It takes the QP vectors of the new QP as well as

the maximum number of working set recalcutations as input arguments, and provides the usual output information, again.
Having solved the last QP of the sequence, you are encouraged to free the internal memory by calling qpOASES_cleanup.

For solving QPs of special types as described in Chapter 4, special variants of the above functions are provided. First, the functions

```
[ obj,x,y,status,nWSRout ] = qpOASES_initSB( H,g,lb,ub,nWSR )
[ obj,x,y,status,nWSRout ] = qpOASES_hotstartSB( g,lb,ub,nWSR )
                             qpOASES_cleanupSB
```

for *s*imply *b*ounded QPs (input arguments correspoding to constraints are simply left away). Second, the functions

```
[ obj,x,y,status,nWSRout ] = qpOASES_initVM( H,g,A,lb,ub,lbA,ubA,nWSR )
[ obj,x,y,status,nWSRout ] = qpOASES_hotstartVM( H,g,A,lb,ub,lbA,ubA,nWSR )
                             qpOASES_cleanupVM
```

for QPs with *v*arying *m*atrices, where qpOASES_hotstartVM also takes the new matrices of the next QP of the sequence.

Again, the internal memory is freed by calling qpOASES_cleanupSB and qpOASES_cleanupVM, respectively. This memory is kept independently for all three QP types.

Examples

The files example1.dat, example1a.dat and example1b.dat contain, respectively, very basic examples for solving a sequence comprising two QPs with fixed matrices, varying matrices, and with simple bounds only. For solving the first one do the following:

1. Start SCILAB and execute the following commands:

   ```
   cd <install-dir>/INTERFACES/SCILAB
   load('example1.dat')
   ```

2. Solve the first QP by typing

   ```
   [ obj,x,y,status,nWSRout ] =
             qpOASES_init( H,g,A,lb,ub,lbA,ubA,10 )
   ```

3. Solve the second QP by typing

   ```
   [ obj,x,y,status,nWSRout ] =
             qpOASES_hotstart( g_new,lb_new,ub_new,lbA_new,ubA_new,10 )
   ```

4. Free the internal memory by calling qpOASES_cleanup.

6.4 Running qpOASES on dSPACE

qpOASES can be easily run on a dSPACE board via its Simulink interface, provided that a C++ compiler is available. This has been tested for dSPACE boards version 5.3 or higher together with the dSPACE C++ Integration Kit 1.0.2 or higher. The following additional notes hopefully facilitate the setup:

1. Setup your dSPACE system
2. Install the dSPACE C++ Integration Kit
3. Install qpOASES (its Simulink interface, to be more precisely)
4. Compile qpOASES with compiler flag `__DSPACE__`. This can be done, e.g., by uncommenting line 42 within the file `INCLUDE/MessageHandling.hpp`.
5. Setup your Simulink project
6. Open MK(make) file of your project (eventually you have to compile it once before) and add the following lines at the head of this file:

   ```
   # enable c++ support
   USER_BUILD_CPP_APPL = ON
   ```

7. Also complete the following lines:

   ```
   USER_SRCS = qpOASES_SQProblem.cpp qpOASES_QProblem.cpp
   qpOASES_QProblemB.cpp SQProblem.cpp Qproblem.cpp QproblemB.cpp
   SubjectTo.cpp Bounds.cpp Constraints.cpp Cyclingmanager.cpp
   Indexlist.cpp MessageHandling.cpp Utils.cpp
   ```
   (i.e. all source files of qpOASES and its Simulink interface)

   ```
   USER_SRCS_DIR = ./SRC
   ```
   (i.e. directory of qpOASES source files)

   ```
   USER_INCLUDES_PATH = ./INCLUDE ./SRC
   ```
   (i.e. directories of qpOASES header and source files)

8. Compile your project
9. Run the compiled project on dSPACE

6.5 Using qpOASES within the ACADO Toolkit

ACADO Toolkit is a software framework for automatic control and dynamic optimisation available at

`http://www.acadotoolkit.org`.

It is an open-source (LGPL) environment for setting up a great variety of dynamic optimization problems for use in control, in particular (nonlinear) model predictive control. ACADO Toolkit uses qpOASES as default QP solver, for linear MPC as well as for the QP sequences resulting from SQP-type methods.

6.6 Using qpOASES within MUSCOD-II

MUSCOD-II is a proprietary software package for numerical solution of optimal control problems involving differential-algebraic equations, developed by the members of the "Simulation and Optimization Group" of the Interdisciplinary Center for Scientific Computing (IWR) at University of Heidelberg. The current version of MUSCOD-II also contains an interface for using qpOASES as underlying QP solver.

Chapter 7

Developer Information and Compiling Options

This chapter provides a very brief introduction to the qpOASES software design. If you are interested in using qpOASES within your own software project or in developing extensions for it yourself, we recommend to consult its DOXYGEN documentation (cf. installation step six described in Chapter 2) for detailed information. Moreover, you are encouraged to pose questions or remarks to `support@qpOASES.org`.

7.1 Class Hierarchy

So far, we mentioned three different classes: `QProblem`, `QProblemB` and `SQProblem`. These are the only classes which provide user interfaces for accessing qpOASES's functionality. However, they are not the only classes of the qpOASES software package but are embedded in a more complex hierarchy (see Figures 7.1 and 7.2).

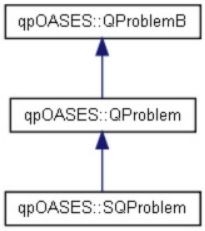

Figure 7.1: `QProblem` class hierarchy (illustrated with DOXYGEN [6]).

The class `QProblemB` is at the bottom of the hierarchy and provides all functionality necessary for solving a simply bounded quadratic program (cf. Section 4.3). The `QProblem` class is derived from it and implements all necessary additional functionality for solving a QPs comprising general constraints. The class `SQProblem`, in turn, inherits all features of the `QProblem` class and provides further functionality for handling QPs with varying matrices (cf. Section 4.2).

All the three classes `QProblemB`, `QProblem` and `SQProblem` make use of further auxiliary classes: they possess members of type `Bounds` or `Constraints` (which are derived from a common type `SubjectTo`) in order to store bounds or constraints of a QP. Both the `Bounds` and the `Constraints` class manages lists (of type `Indexlist`) of free and fixed variables and active and inactive (and disabled) constraints, respectively. Moreover, the `QProblem` class also keeps an `CyclingManager` object for detecting and dealing with possible cycling (only rudimentary implemented, as it seems to occur quite rarely!).

Finally, all the above mentioned classes use a class called `MessageHandling` for providing errors messages, warnings or other information to the user and for handling return values of their member functions in a unified framework. The current implementation uses a single *global* instance of the `MessageHandling` class; the global function

```
MessageHandling* getGlobalMessageHandler( );
```

returns a pointer to it.

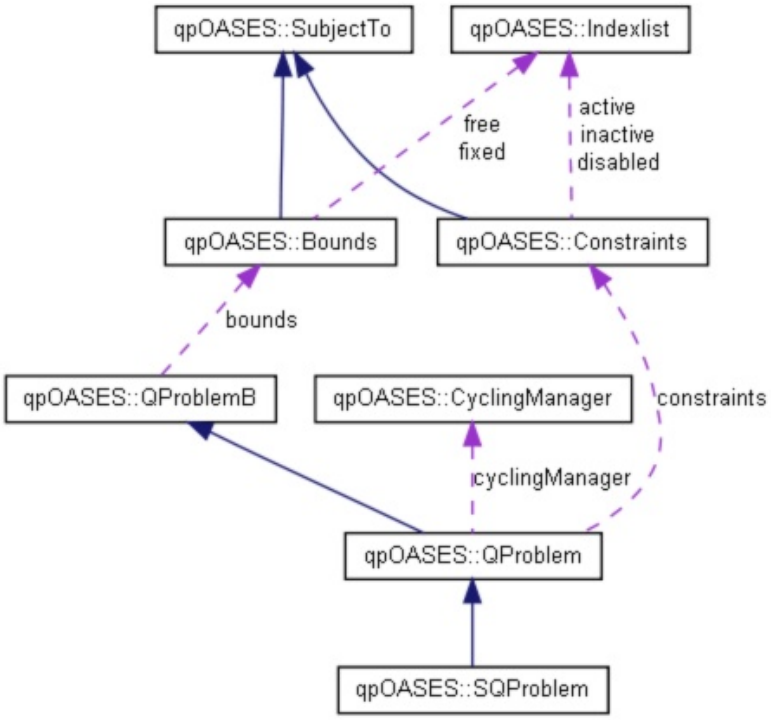

Figure 7.2: qpOASES class hierarchy (illustrated with DOXYGEN [6]).

7.2 Global Constants

Some useful global constants are defined in file <install-dir>/INCLUDE/Constants.hpp. Their default values seem to work reasonably, but you might change them if necessary:

- EPS: numerical value of machine precision,
- ZERO: numerical value of zero (for situations in which it would be unreasonable to compare with 0.0),
- INFTY: numerical value of infinity (e.g. for non-existing bounds),
- BOUNDTOL: lower/upper (constraints') bound tolerance (an inequality constraint whose lower and upper bound differ by less than BOUNDTOL is regarded to be an equality constraint),
- BOUNDRELAXATION: offset for relaxing bounds at beginning of an initial homotopy.
- ENABLINGFACTOR / ENABLINGOFFSET: when enabling a constraint, it is relaxed to (1 +/- ENABLINGFACTOR) * lbA/ubA +/- ENABLINGOFFSET,
- EPS_FOR_REGULARISATION: scaling factor of identity matrix used for Hessian regularisation,
- MAX_NUMBER_OF_REGULARISATION_STEPS: maximum number of successive regularisation steps.

7.3 Compiler Flags

When compiling qpOASES, you can define the following compiler flags:

- LINUX : activates LINUX-specific functionality (e.g. time measurement),
- WIN32 : activates WINDOWS-specific functionality (e.g. time measurement),
- __DEBUG__ : activates more detailed output messages during the QP solution,
- __MATLAB__ : activates MATLAB-specific functionality (in particular, the use of mexPrintf instead of printf),
- __cplusplus : necessary for building C++ S functions for SIMULINK,
- __DSPACE__ : define this compiler flag in order to disable the qpOASES namespace (and switching off all text messages) for ensuring backward compatibility with dSPACE compilers,

- __XPCTARGET__: define this compiler flag in order to disable all text messages for ensuring compatibility for xPC Target compilers,

- __MANY_CONSTRAINTS__: enables a usually faster way for determining the current step length for QPs comprising many constraints (see Section 5.5),

- __USE_THREE_MULTS_GIVENS__: switches to a different way of calculating Givens rotations that requires only three multiplications,

- __ALWAYS_INITIALISE_WITH_ALL_EQUALITIES__: forces to always include all implicitly fixed bounds and all equality constraints into the initial working set when setting up an auxiliary QP.

Bibliography

[1] Online QP Benchmark Collection. http://homes.esat.kuleuven.be/∼optec/software/onlineQP/.

[2] M.J. Best. *Applied Mathematics and Parallel Computing*, chapter An Algorithm for the Solution of the Parametric Quadratic Programming Problem, pages 57–76. Physica-Verlag, Heidelberg, 1996.

[3] H.J. Ferreau. An Online Active Set Strategy for Fast Solution of Parametric Quadratic Programs with Applications to Predictive Engine Control. Master's thesis, University of Heidelberg, 2006.

[4] H.J. Ferreau, H.G. Bock, and M. Diehl. An online active set strategy to overcome the limitations of explicit MPC. *International Journal of Robust and Nonlinear Control*, 18(8):816–830, 2008.

[5] H.J. Ferreau, P. Ortner, P. Langthaler, L. del Re, and M. Diehl. Predictive Control of a Real-World Diesel Engine using an Extended Online Active Set Strategy. *Annual Reviews in Control*, 31(2):293–301, 2007.

[6] D. van Heesch. Doxygen homepage. http://www.doxygen.org.

Appendix A

qpOASES Software Licence

qpOASES is distributed under the terms of the GNU *Lesser* General Public License (LGPL) as published by the Free Software Foundation:

```
                    GNU LESSER GENERAL PUBLIC LICENSE
                         Version 2.1, February 1999

 Copyright (C) 1991, 1999 Free Software Foundation, Inc.
 51 Franklin Street, Fifth Floor, Boston, MA  02110-1301  USA
 Everyone is permitted to copy and distribute verbatim copies
 of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL.  It also counts
 as the successor of the GNU Library Public License, version 2, hence
 the version number 2.1.]

                            Preamble

  The licenses for most software are designed to take away your
freedom to share and change it.  By contrast, the GNU General Public
Licenses are intended to guarantee your freedom to share and change
free software--to make sure the software is free for all its users.

  This license, the Lesser General Public License, applies to some
specially designated software packages--typically libraries--of the
Free Software Foundation and other authors who decide to use it.  You
can use it too, but we suggest you first think carefully about whether
this license or the ordinary General Public License is the better
strategy to use in any particular case, based on the explanations below.

  When we speak of free software, we are referring to freedom of use,
not price.  Our General Public Licenses are designed to make sure that
you have the freedom to distribute copies of free software (and charge
for this service if you wish); that you receive source code or can get
it if you want it; that you can change the software and use pieces of
it in new free programs; and that you are informed that you can do
these things.

  To protect your rights, we need to make restrictions that forbid
distributors to deny you these rights or to ask you to surrender these
rights.  These restrictions translate to certain responsibilities for
you if you distribute copies of the library or if you modify it.

  For example, if you distribute copies of the library, whether gratis
or for a fee, you must give the recipients all the rights that we gave
you.  You must make sure that they, too, receive or can get the source
```

```
code.  If you link other code with the library, you must provide
complete object files to the recipients, so that they can relink them
with the library after making changes to the library and recompiling
it.  And you must show them these terms so they know their rights.

  We protect your rights with a two-step method: (1) we copyright the
library, and (2) we offer you this license, which gives you legal
permission to copy, distribute and/or modify the library.

  To protect each distributor, we want to make it very clear that
there is no warranty for the free library.  Also, if the library is
modified by someone else and passed on, the recipients should know
that what they have is not the original version, so that the original
author's reputation will not be affected by problems that might be
introduced by others.

  Finally, software patents pose a constant threat to the existence of
any free program.  We wish to make sure that a company cannot
effectively restrict the users of a free program by obtaining a
restrictive license from a patent holder.  Therefore, we insist that
any patent license obtained for a version of the library must be
consistent with the full freedom of use specified in this license.

  Most GNU software, including some libraries, is covered by the
ordinary GNU General Public License.  This license, the GNU Lesser
General Public License, applies to certain designated libraries, and
is quite different from the ordinary General Public License.  We use
this license for certain libraries in order to permit linking those
libraries into non-free programs.

  When a program is linked with a library, whether statically or using
a shared library, the combination of the two is legally speaking a
combined work, a derivative of the original library.  The ordinary
General Public License therefore permits such linking only if the
entire combination fits its criteria of freedom.  The Lesser General
Public License permits more lax criteria for linking other code with
the library.

  We call this license the "Lesser" General Public License because it
does Less to protect the user's freedom than the ordinary General
Public License.  It also provides other free software developers Less
of an advantage over competing non-free programs.  These disadvantages
are the reason we use the ordinary General Public License for many
libraries.  However, the Lesser license provides advantages in certain
special circumstances.

  For example, on rare occasions, there may be a special need to
encourage the widest possible use of a certain library, so that it becomes
a de-facto standard.  To achieve this, non-free programs must be
allowed to use the library.  A more frequent case is that a free
library does the same job as widely used non-free libraries.  In this
case, there is little to gain by limiting the free library to free
software only, so we use the Lesser General Public License.

  In other cases, permission to use a particular library in non-free
programs enables a greater number of people to use a large body of
free software.  For example, permission to use the GNU C Library in
non-free programs enables many more people to use the whole GNU
operating system, as well as its variant, the GNU/Linux operating
system.

  Although the Lesser General Public License is Less protective of the
users' freedom, it does ensure that the user of a program that is
linked with the Library has the freedom and the wherewithal to run
```

that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) The modified work must itself be a software library.

b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.

```
    c) You must cause the whole of the work to be licensed at no
    charge to all third parties under the terms of this License.

    d) If a facility in the modified Library refers to a function or a
    table of data to be supplied by an application program that uses
    the facility, other than as an argument passed when the facility
    is invoked, then you must make a good faith effort to ensure that,
    in the event an application does not supply such function or
    table, the facility still operates, and performs whatever part of
    its purpose remains meaningful.

    (For example, a function in a library to compute square roots has
    a purpose that is entirely well-defined independent of the
    application.  Therefore, Subsection 2d requires that any
    application-supplied function or table used by this function must
    be optional: if the application does not supply it, the square
    root function must still compute square roots.)

These requirements apply to the modified work as a whole.  If
identifiable sections of that work are not derived from the Library,
and can be reasonably considered independent and separate works in
themselves, then this License, and its terms, do not apply to those
sections when you distribute them as separate works.  But when you
distribute the same sections as part of a whole which is a work based
on the Library, the distribution of the whole must be on the terms of
this License, whose permissions for other licensees extend to the
entire whole, and thus to each and every part regardless of who wrote
it.

Thus, it is not the intent of this section to claim rights or contest
your rights to work written entirely by you; rather, the intent is to
exercise the right to control the distribution of derivative or
collective works based on the Library.

In addition, mere aggregation of another work not based on the Library
with the Library (or with a work based on the Library) on a volume of
a storage or distribution medium does not bring the other work under
the scope of this License.

  3. You may opt to apply the terms of the ordinary GNU General Public
License instead of this License to a given copy of the Library.  To do
this, you must alter all the notices that refer to this License, so
that they refer to the ordinary GNU General Public License, version 2,
instead of to this License.  (If a newer version than version 2 of the
ordinary GNU General Public License has appeared, then you can specify
that version instead if you wish.)  Do not make any other change in
these notices.

  Once this change is made in a given copy, it is irreversible for
that copy, so the ordinary GNU General Public License applies to all
subsequent copies and derivative works made from that copy.

  This option is useful when you wish to copy part of the code of
the Library into a program that is not a library.

  4. You may copy and distribute the Library (or a portion or
derivative of it, under Section 2) in object code or executable form
under the terms of Sections 1 and 2 above provided that you accompany
it with the complete corresponding machine-readable source code, which
must be distributed under the terms of Sections 1 and 2 above on a
medium customarily used for software interchange.

  If distribution of object code is made by offering access to copy
```

from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

b) Use a suitable shared library mechanism for linking with the

Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.

c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.

d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.

e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and

all its terms and conditions for copying, distributing or modifying
the Library or works based on it.

  10. Each time you redistribute the Library (or any work based on the
Library), the recipient automatically receives a license from the
original licensor to copy, distribute, link with or modify the Library
subject to these terms and conditions. You may not impose any further
restrictions on the recipients' exercise of the rights granted herein.
You are not responsible for enforcing compliance by third parties with
this License.

  11. If, as a consequence of a court judgment or allegation of patent
infringement or for any other reason (not limited to patent issues),
conditions are imposed on you (whether by court order, agreement or
otherwise) that contradict the conditions of this License, they do not
excuse you from the conditions of this License. If you cannot
distribute so as to satisfy simultaneously your obligations under this
License and any other pertinent obligations, then as a consequence you
may not distribute the Library at all. For example, if a patent
license would not permit royalty-free redistribution of the Library by
all those who receive copies directly or indirectly through you, then
the only way you could satisfy both it and this License would be to
refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any
particular circumstance, the balance of the section is intended to apply,
and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any
patents or other property right claims or to contest validity of any
such claims; this section has the sole purpose of protecting the
integrity of the free software distribution system which is
implemented by public license practices. Many people have made
generous contributions to the wide range of software distributed
through that system in reliance on consistent application of that
system; it is up to the author/donor to decide if he or she is willing
to distribute software through any other system and a licensee cannot
impose that choice.

This section is intended to make thoroughly clear what is believed to
be a consequence of the rest of this License.

  12. If the distribution and/or use of the Library is restricted in
certain countries either by patents or by copyrighted interfaces, the
original copyright holder who places the Library under this License may add
an explicit geographical distribution limitation excluding those countries,
so that distribution is permitted only in or among countries not thus
excluded. In such case, this License incorporates the limitation as if
written in the body of this License.

  13. The Free Software Foundation may publish revised and/or new
versions of the Lesser General Public License from time to time.
Such new versions will be similar in spirit to the present version,
but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library
specifies a version number of this License which applies to it and
"any later version", you have the option of following the terms and
conditions either of that version or of any later version published by
the Free Software Foundation. If the Library does not specify a
license version number, you may choose any version ever published by
the Free Software Foundation.

  14. If you wish to incorporate parts of the Library into other free

programs whose distribution conditions are incompatible with these,
write to the author to ask for permission. For software which is
copyrighted by the Free Software Foundation, write to the Free
Software Foundation; we sometimes make exceptions for this. Our
decision will be guided by the two goals of preserving the free status
of all derivatives of our free software and of promoting the sharing
and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO
WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW.
EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR
OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY
KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE
LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME
THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN
WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY
AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU
FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR
CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE
LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING
RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A
FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF
SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH
DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest
possible use to the public, we recommend making it free software that
everyone can redistribute and change. You can do so by permitting
redistribution under these terms (or, alternatively, under the terms of the
ordinary General Public License).

To apply these terms, attach the following notices to the library. It is
safest to attach them to the start of each source file to most effectively
convey the exclusion of warranty; and each file should have at least the
"copyright" line and a pointer to where the full notice is found.

    <one line to give the library's name and a brief idea of what it does.>
    Copyright (C) <year> <name of author>

    This library is free software; you can redistribute it and/or
    modify it under the terms of the GNU Lesser General Public
    License as published by the Free Software Foundation; either
    version 2.1 of the License, or (at your option) any later version.

    This library is distributed in the hope that it will be useful,
    but WITHOUT ANY WARRANTY; without even the implied warranty of
    MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU
    Lesser General Public License for more details.

    You should have received a copy of the GNU Lesser General Public
    License along with this library; if not, write to the Free Software
    Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Also add information on how to contact you by electronic and paper mail.

```
You should also get your employer (if you work as a programmer) or your
school, if any, to sign a "copyright disclaimer" for the library, if
necessary.  Here is a sample; alter the names:

  Yoyodyne, Inc., hereby disclaims all copyright interest in the
  library ‘Frob’ (a library for tweaking knobs) written by James Random Hacker.

  <signature of Ty Coon>, 1 April 1990
  Ty Coon, President of Vice

That’s all there is to it!
```