



ELSEVIER

Available online at www.sciencedirect.com

Journal of Magnetic Resonance xxx (2004) xxx–xxx

JMR
Journal of
Magnetic Resonancewww.elsevier.com/locate/jmr

ODIN—Object-oriented development interface for NMR

Thies H. Jochimsen* and Michael von Mengershausen

Max-Planck-Institute of Cognitive Neuroscience, Stephanstr. 1a, D 04103 Leipzig, Germany

Received 4 December 2003; revised 4 May 2004

Abstract

A cross-platform development environment for nuclear magnetic resonance (NMR) experiments is presented. It allows rapid prototyping of new pulse sequences and provides a common programming interface for different system types. With this object-oriented interface implemented in C++, the programmer is capable of writing applications to control an experiment that can be executed on different measurement devices, even from different manufacturers, without the need to modify the source code. Due to the clear design of the software, new pulse sequences can be created, tested, and executed within a short time. To post-process the acquired data, an interface to well-known numerical libraries is part of the framework. This allows a transparent integration of the data processing instructions into the measurement module. The software focuses mainly on NMR imaging, but can also be used with limitations for spectroscopic experiments. To demonstrate the capabilities of the framework, results of the same experiment, carried out on two NMR imaging systems from different manufacturers are shown and compared with the results of a simulation.

© 2004 Published by Elsevier Inc.

PACS: 87.59.Pw

Keywords: NMR; Software; Sequence programming; Platform-independent; Pulse design

1. Introduction

Nuclear magnetic resonance (NMR) is a versatile tool to investigate physical properties of materials and living tissue. The flexibility of the NMR technique can be attributed to the fact that a wide range of experiments is designed by solely altering the software that controls the hardware during the measurement. With a given set of hardware components, various properties of the sample can be examined with different software-based experimental setups (i.e., pulse sequences). An important task of the NMR scientist who develops new NMR applications is therefore that of a software engineer. Provided a sophisticated programming interface for sequence design is available, advances in the field of computer science can accelerate the process of creating NMR applications.

Contemporary concepts like object-oriented design, polymorphism, and generic programming are used nowadays in software engineering to create modular,

extensible, and easy-to-use software instead of procedural programming (an excellent overview of these programming paradigms and their implementation in C++ can be found in [1]). By contrast, NMR pulse sequences are usually programmed using the procedural approach. That is, the scientist provides a program that contains a list of sequential instructions to trigger hardware-events together with some calculations to achieve the required properties of the sequence (e.g., resolution, orientation, and contrast). This results in a non-modular, monolithic implementation of the sequence which seriously limits the reuse of certain parts in another sequence, except for duplicating the source code. A modern approach would describe the sequence as a composition of reusable, self-consistent objects that can be combined freely to develop new experimental setups.

Recently, a software architecture has been presented [2] which makes use of this approach by a double-layered design whereby the user interacts with an application framework written in Java [3] which is mapped to corresponding C++ functionality on the hardware

* Corresponding author. Fax: +49-341-99-40-221.

E-mail address: thies@jochimsen.de (T.H. Jochimsen).

60 controller and signal processing computer. The pro-
61 gramming interface is provided not only for sequence
62 programming but also for developing work flows which
63 incorporate different measurement techniques for clinical
64 application. However, this framework is limited to
65 the devices of one manufacturer and its double-layered
66 design may impose a considerable overhead when adding
67 new functionality, for example custom real-time
68 feedback.

69 In contrast, ODIN, which is subsequently introduced,
70 concentrates on platform-independent sequence design,
71 and data processing with a single open-source code basis
72 in C++. The hardware-dependent components that
73 drive the different scanners are encapsulated into low-
74 level objects (pulses, gradients, and data-acquisition)
75 from which complex, platform-independent parts of the
76 sequence are constructed. The same source code is used
77 at all stages of sequence development, from simulation
78 on a stand-alone platform to play-out on a real-time
79 system. ODIN uses the native functionality of the
80 graphical user interface on each platform, allowing a
81 seamless integration of ODIN sequences. Although
82 ODIN is a relatively young software project, its se-
83 quence programming interface has been shown advan-
84 tageous in developing sophisticated functional magnetic
85 resonance imaging (fMRI) applications [4–6], in simu-
86 lations [7], and in the application of its module for pulse
87 design [8].

88 In this paper, the first section gives an introduction
89 into the ODIN sequence programming interface and its
90 underlying concepts. The design of radio frequency (RF)
91 pulses will be described in more detail as this is one of
92 the major strengths of ODIN. The next two sections
93 contain additional information about the internal rep-
94 resentation of the sequence within the ODIN library and
95 the mechanisms that are used to execute the experiment
96 in different hardware environments. After that, strate-
97 gies to visualize and simulate the sequence are presented,
98 and the data processing framework of ODIN is dis-
99 cussed. Finally, experimental results obtained with
100 ODIN on different platforms are shown and compared
101 with the results of a simulation.

102 2. Platform-independent sequence design

103 An NMR experiment is basically a sequence of pe-
104 riods where the sample is exposed to different magnetic
105 field configurations, such as RF pulses and magnetic
106 field gradients, or periods where data are acquired.
107 From these basic sequence elements, complex experi-
108 ments can be composed which measure spectroscopic
109 properties, relaxation, and transportation processes of
110 the spins within the sample. Magnetic field gradients
111 extend these experiments to spatially resolved data sets,
112 i.e., images of these parameters. In addition, repetitive

113 measurements yield time series of physiological pro-
114 cesses within living tissue, for example, neuronal activity
115 in the human brain.

116 The NMR sequence can be described in terms of the
117 physical properties of their elements and the arrange-
118 ment of these sequence elements as a function of time. A
119 simple NMR sequence is shown in Fig. 1. This level of
120 description is independent of the measurement device.
121 ODIN provides a programming interface in terms of a
122 C++ class hierarchy which reflects the physical aspects
123 of a sequence. A sequence program which is written
124 using this framework can be executed on different NMR
125 hardware. The system-specific actions are performed by
126 a library that transfers the sequence-specific requests to
127 the actual measurement hardware as depicted in Fig. 2.
128 The benefit of separating the physical logic of the ex-
129 periment from the peculiarities of the current hardware
130 is the portability of the sequence program. It can be
131 reused with other hardware, even from another manu-
132 facturer.

2.1. Sequence programming interface 133

134 In the following, the term *basic sequence objects* refers
135 to elements of the sequence that cannot be divided into
136 smaller elements from the physical point of view. Ex-
137 amples of such “sequence atoms” are periods of RF
138 irradiation, the application of temporary field gradients
139 or intervals of data acquisition. Each basic sequence
140 object is represented by a C++ class which handles its
141 physical properties, for example the duration. These
142 objects are constructed during the initialization of the
143 sequence according to the instructions given by the se-
144 quence programmer. From this collection, the sequence
145 is constructed by grouping the sequence objects into
146 container objects. To simplify the notion of composing
147 new container objects, the operators + and / are over-
148 loaded, i.e., they are redefined with sequence objects as
149 operands, and can be used to specify whether two se-
150 quence objects a and b should be played out sequentially
151 (a+b) or in parallel (a/b). As an example, the source
152 code for the simple sequence visualized in Fig. 1 is
153 printed in Fig. 3.

154 Besides this technique of building sequences from
155 scratch by grouping basic sequence objects together, the
156 ODIN library offers many predefined high-level se-
157 quence objects as C++ classes. For example, the object
158 `acq` in Figs. 1 and 3 is an acquisition window with the
159 simultaneous application of a gradient field that is used
160 in many imaging sequences for spatial frequency en-
161 coding. These more complex objects are constructed
162 from basic sequence objects within the library, using the
163 same mechanism of building container objects as the
164 sequence programmer would. In addition, the class of
165 these composite objects provides an interface that is
166 adjusted to its high-level concept. For instance, the ob-




Fig. 1. A simple gradient-echo sequence. The inner part contains a slice-selective RF pulse, gradients G_x , G_y , and G_z for spatial encoding, and a period during which the signal is received. This part is repeated N times for linear stepping of the gradient strength of G_y . The sequence objects of these elements are indicated below. The operators $+$ and $/$ between these objects combines them to form the sequence.

167 ject `acq` has a member function that returns the point in
 168 time of the center of the acquisition window with proper
 169 consideration of the delayed onset due to the ramp of
 170 the simultaneous gradient.

171 2.2. Pulse design

172 A crucial part of the sequence is the application of
 173 RF pulses to generate a detectable signal from a limited
 174 spatial or spectral range of spins within the sample. The
 175 ODIN framework contains a flexible module for the
 176 generation and simulation of RF pulses. A wide range of
 177 pulses is supported by a plug-in style mechanism. The
 178 desired excitation profile, gradient shape, and frequency
 179 filter can be selected and modified separately to match
 180 the pulse optimally to the specific application. It can be
 181 easily extended by supplying the module with new plug-
 182 ins which generate k -space trajectories or calculate the
 183 RF waveform as a function of time or k -space coordi-
 184 nate. The following pulse types are already supported by
 185 existing plug-ins of the ODIN library:

- Slice-selective pulses (Sinc, Gauss), optionally with a VERSE [9] trajectory for reduced power excitation. 186
- Adiabatic pulses (Sech [10], WURST [11]). 187
- Spectrally and spatially selective pulses [12] for slice-selection with a predefined spectral profile (e.g., for fat suppression). 188
- Two dimensional (2D) pulses [13] with various excitation shapes and different spiral trajectories. 189
- Composite pulses [14] which are created by concatenating one of the above pulses with different transmitter phases and flip angles. 190

In addition, these pulses can be filtered either in k -space or in the time domain using a filter plug-in. The benefit from separating the pulse shape and the trajectory into different plug-ins can be illustrated by considering the generation of 2D pulses: each of the excitation profiles (point, box, disk, and user-defined list of points) can be used in combination with any of the 2D trajectories in order to generate a pulse profile that is well adjusted to the requirements. For example, an excitation profile that consists of a chain of adjacent points together with a

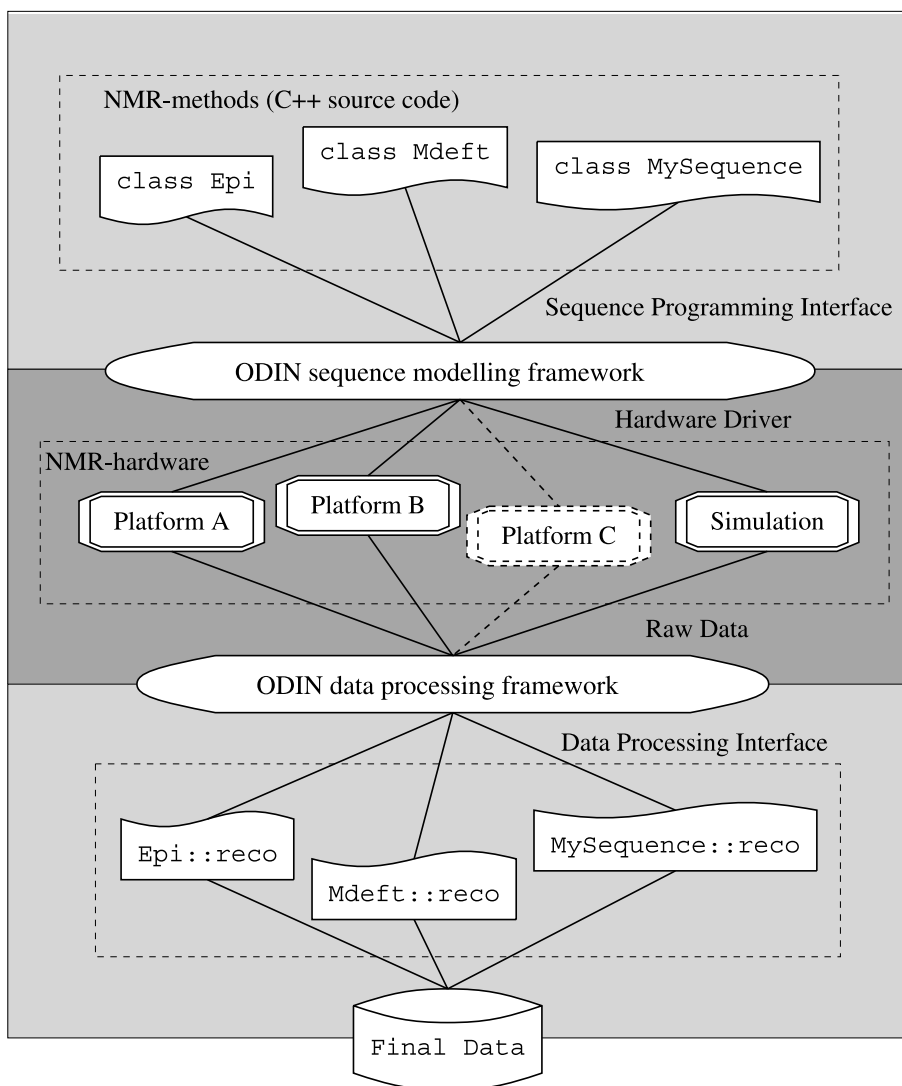


Fig. 2. Flowchart of an NMR experiment performed with the ODIN framework. The sequence programmer implements a C++ class that represents the experimental method and uses the platform-independent sequence programming interface. An object of this class is then used by the ODIN library to execute the sequence on the different platforms by means of hardware-specific instructions within the library. The acquired raw data is then post-processed by a member function `reco` of the same class that was used for the measurement. Finally, the processed data (images, spectra) are written to disk.

207 slew-rate optimized trajectory is useful for curved slice
208 imaging [15].

209 Because the pulse module is a regular sequence object,
210 it can be integrated seamlessly into any NMR sequence.
211 For example, the object `pulse` in Figs. 1 and 3 is a slice-
212 selective specialization of this module using the `Sinc`
213 plug-in for the pulse shape. In addition, a graphical user
214 interface (Fig. 4) which acts as a front-end to the pulse
215 module can be used for interactive pulse design and
216 monitoring of the corresponding excitation profile.

217 2.3. Loops and vectors

218 An essential aspect in most NMR experiments is to
219 repeat certain parts of the sequence unchanged or with
220 different settings. Examples are the repetition of a gra-

221 dient-echo with different strength of the phase-encoding 221
222 gradient in conventional Fourier imaging as used in the 222
223 sequence of Fig. 1, or the repetition with different pulse 223
224 frequencies for multi-slice acquisition. 224

To use this technique in a uniform manner, ODIN 225
introduces the concept of vector objects and loop objects. 226
Vector objects are elements of the sequence that are used 227
repeatedly with different settings. The following prede- 228
fined vector classes, derived from a common base class 229
`Seq Vector`, are available to the sequence programmer: 230

- Gradient pulses with different gradient strengths for 231
phase encoding or diffusion weighting. 232
- Sequence objects that drive the transmitter (RF 233
pulses) or receiver (acquisition windows) contain two 234
vector objects for frequency and phase switching to 235
be used for multi-slice experiments or phase cycling. 236

```

class SimpleSequence : public SeqMethod {
private:
  // Sequence objects:
  SeqPulsarSinc pulse;      SeqGradPhaseEnc phase;
  SeqAcqRead    acq;        SeqAcqDeph    deph;
  SeqObjLoop    loop;       SeqDelay      delay;
  SeqObjList    oneline;

public:
  SimpleSequence(const tjstring& label) : SeqMethod(label) {
    set_description("Simple Gradient Echo Sequence");
  }

  void method_pars_init() {
    // This is the place where sequence-specific parameters can be initialized
  }

  void method_seq_init() {
    // This function builds the sequence, it is called every time a parameter has been changed by the user

    // The global objects commonPars, geometryInfo and systemInfo hold parameters
    // that are common to most sequences, the information about the selected
    // geometry and system specific properties, respectively. These parameters
    // can be accessed via the appropriate 'get' and 'set' functions.

    // Excitation pulse:
    pulse=SeqPulsarSinc("pulse",geometryInfo->get_sliceThickness());

    // Geometry:
    // calculate the resolution in the read direction and set the number of
    // phase encoding steps so that a uniform resolution will be obtained
    float resolution=geometryInfo->get_FOV(readChannel) /commonPars->get_MatrixSize(readChannel);
    commonPars->set_MatrixSize(phaseChannel,geometryInfo->get_FOV(phaseDirection) / resolution);

    // Phase encoding:
    phase=SeqGradPhaseEnc("phase",commonPars->get_MatrixSize(phaseChannel),
      geometryInfo->get_FOV(phaseChannel),phaseChannel,0.25*systemInfo->get_max_grad());

    // Frequency encoding:
    acq=SeqAcqRead("acq",commonPars->get_AcqSweepWidth(),
      commonPars->get_MatrixSize(readChannel),geometryInfo->get_FOV(readChannel),readChannel);

    // Dephasing for frequency encoding
    deph=SeqAcqDeph("deph",acq,FID);

    // One gradient echo to sample one line in k-space
    oneline = pulse + phase/deph + acq;

    // Sequence layout:
    set_sequence( loop ( oneline + delay ) [phase] );
  }

  void method_rels() {
    // This is the place where sequence timing is performed

    // ensure correct repetition time by setting the duration of 'delay'
    double linedur=oneline.get_duration();
    if(linedur>commonPars->get_RepetitionTime()) commonPars->set_RepetitionTime(linedur);
    delay.set_duration( (commonPars->get_RepetitionTime()-linedur));
  }

  void method_pars_set() {
    // This function is called once before the measurement is started
  }
};

```

Fig. 3. The source code of a simple gradient-echo sequence, implemented as a C++ class to be used within the ODIN framework.

- 237 • Delay objects with a variable duration, which is changed for each iteration. 242
- 238
- 239 • A list of user-defined rotation matrices that can be attached to gradient-related objects to alter their direction subsequently. 243
- 240
- 241
- A container object that holds a list of other sequence objects which are played out sequentially for each repetition. 244
- Although this set of specialized vector classes is probably not exhaustive, the last class may be used to easily 245
- 246

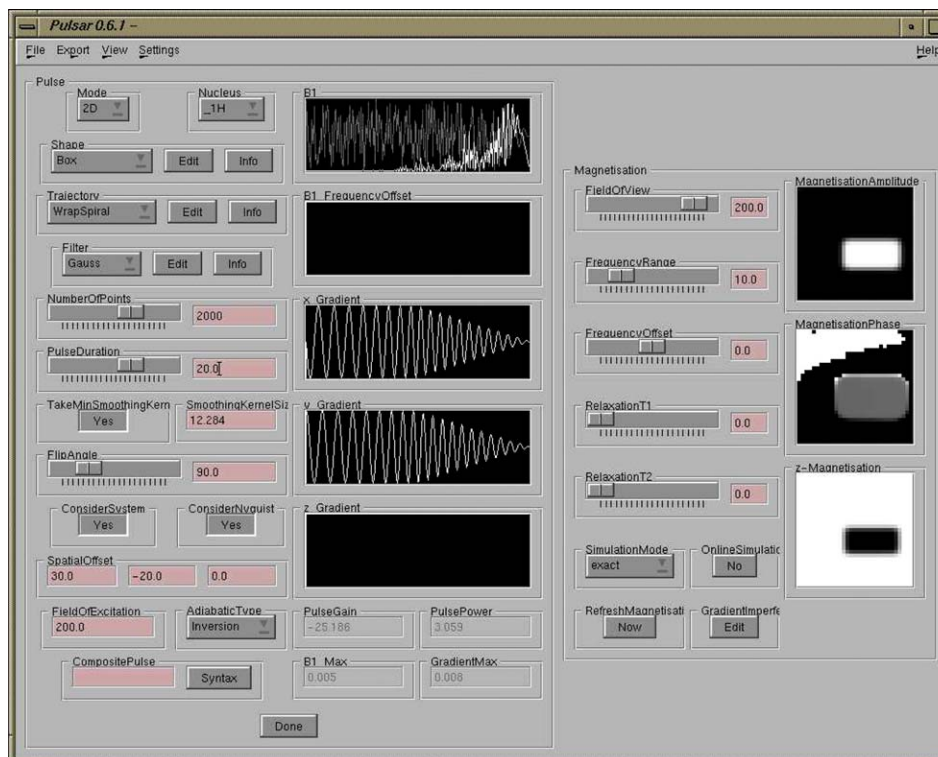


Fig. 4. The Pulsar user interface for interactive pulse design and simulation. The panel to the left allows editing of the pulse parameters and shows the time courses of the RF and gradient fields. The current settings show a 2D selective pulse, i.e., a pulse that restricts the excited spins in two dimensions. The right-hand side displays the result of a simulation with this pulse.

247 extend this list by storing sequence objects for each
248 repetition into the container. This emulates the behavior
249 of a built-in vector class.

250 To specify which parts of the sequence will be re-
251 peated and which vectors will be modified at each rep-
252 etition, loop objects play a central role in sequence
253 design with ODIN. They possess a function-like syntax
254 (functors) when used within a sequence:

```
255 loop (kernel) [vector1][vector2]...
```

256 With this line of source code, the loop object loop is
257 used to repeat the sequence part kernel while incre-
258 menting the properties of the vector objects vec-
259 tor1, vector2, ... that are located within kernel.
260 Instead of using a vector object, an integer number N can
261 also be given as an argument to the loop, which will then
262 repeat the sequence part N times unchanged. By using
263 this common notation for all variable aspects of a se-
264 quence, new applications can be implemented rapidly
265 without dealing with the specific aspects of the hardware.

266 2.4. Sequence parameters

267 Normally, each sequence has a set of parameters
268 which specify the actual experiment, for example, the
269 sampling rate for data acquisition or the duration of the
270 RF pulse. The sequence parameters are edited interac-
271 tively within the user interface of the measurement de-

vice, and the sequence is recalculated according to the
new settings. Within ODIN, these parameters are
members of the C++ sequence class, allowing trans-
parent access to their values in the member function that
prepares the experiment. Well-known data types (integer
numbers, floating point numbers, and Boolean values)
can be used as sequence parameters. They are designed
to be used exactly like built-in types of the C++ lan-
guage, resulting in understandable source code.

Whenever possible, the native user interface of the
measurement device is used to present the set of pa-
rameters specified by the sequence programmer. There-
by, the parameter values are exchanged between the
native user interface and the ODIN library. If no native
mechanism for parameter editing exists (e.g., on a stand-
alone platform), ODIN provides its own set of widgets
using the Qt library [16] to edit the parameters interac-
tively (Fig. 5). After the measurement, the parameters
are stored on disk in JCAMP-DX format [17] together
with the raw data. In the post-processing step, the pa-
rameters and the raw data are then read from disk.

3. Internal representation of the sequence

Any NMR sequence has a nested structure, that is,
basic sequence objects can be grouped together to form

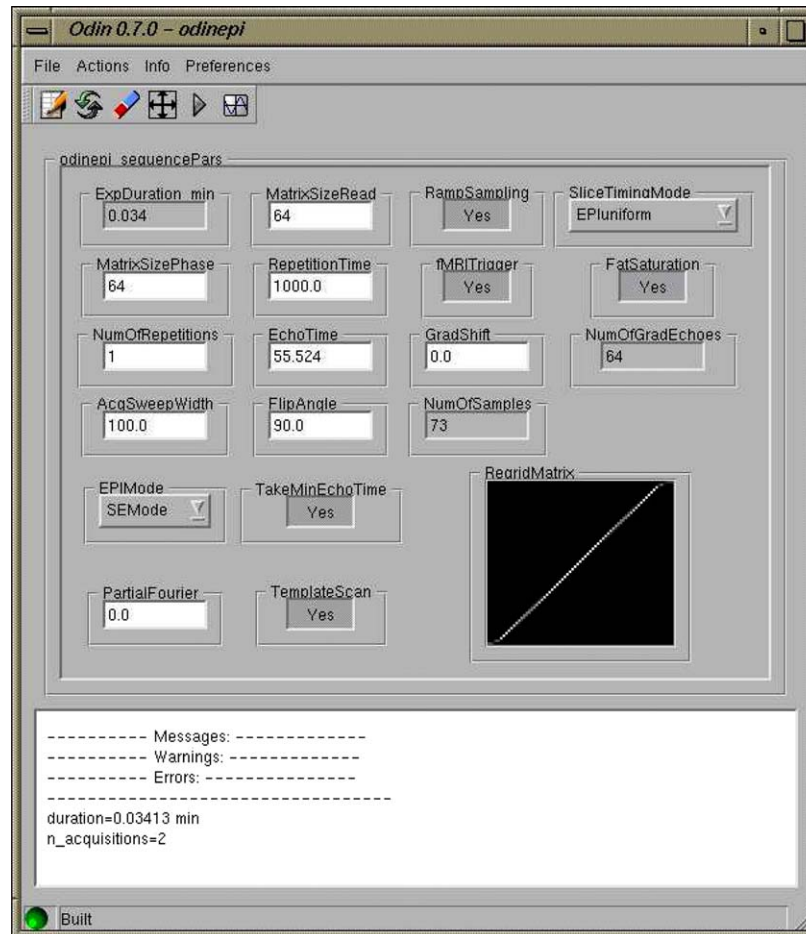


Fig. 5. User interface for rapid sequence design. It combines frequently used functionality to edit, compile, visualize, and simulate ODIN sequences. The set of widgets for the parameters is dynamically generated according to the specifications of the sequence module. The parameter set for an EPI sequence is shown here.

296 logical units, which in turn can be collected to build
 297 more complex units. This leads to an internal repre-
 298 sentation of the sequence as an ordered tree of sequence
 299 objects. The leaves of this sequence tree are the basic
 300 sequence objects (RF pulses, gradients, acquisition
 301 windows, and evolution delays). The sequence contain-
 302 ers are represented by nodes of the tree. They contain a
 303 list of references to their members in the same order as
 304 given by the sequence programmer. The nodes can
 305 contain additional information, e.g., a loop object con-
 306 tains the number of repetitions besides the elements of
 307 the sequence that are repeated.

308 The tree is constructed during the preparation phase
 309 of the experiment according to the instructions of the
 310 sequence programmer. Each sequence has its special
 311 tree. As an example, Fig. 6 depicts the sequence tree
 312 structure for the sequence of Fig. 3. The created se-
 313 quence tree is the central data structure that is used in
 314 further steps of the experiment. If a certain operation
 315 has to be performed for the sequence, e.g., calculating
 316 the total duration of the experiment, the sequence tree is
 317 traversed recursively, querying each object for a value

(in this case its duration), or requesting a certain operation
 318 from the object. Thereby the starting point is the
 319 root of the sequence tree. At each node that contains an
 320 ordered list of other sequence objects, these sub-objects
 321 are in turn requested to perform the operation. This
 322 recursion in each branch terminates at the leaves, if a
 323 basic sequence object is reached. The two following
 324 sections describe how this technique of traversing the
 325 sequence tree is used to control the measurement device
 326 or to visualize and simulate the sequence.
 327

The whole sequence (i.e., the root of the tree) is in
 328 itself a container object, represented by a C++ class,
 329 which is implemented by the sequence programmer.
 330 This class is derived from a base class that acts as an
 331 interface between the sequence and the ODIN library.
 332 By the mechanism of virtual functions in C++, a set of
 333 sequence-specific member functions must be provided
 334 by the sequence class that will be called during initial-
 335 ization, preparation, and data processing of the experi-
 336 ment. With this technique, all sequence modules share a
 337 common interface which can be used by the library in a
 338 uniform manner.
 339

Label	Type	Duration[ms]	Properties
simple			
└─unnamedSeqObjLoop0	SimpleSequence	1100.130	NumOfObjects=1
└─pulse+phase/depth_acq_read/acq_middelay+acq_acq_a...	SeqObjLoop	1100.130	Times=11, NumOfVectors=1, NumOfObjects=1
└─pulse	SeqObjList	100.0	NumOfObjects=4
└─pulse_handler	SeqPulsarSinc	4.76688	Shape=Sinc, Trajectory=Const, Filter=Triangle
└─pulse_gz	SeqGradChanParallel	4.76688	
└─gy_dummy	SeqGradChanList	4.76688	
└─gx_dummy	SeqGradChanList	3.14667	
└─pulse_rtrain	SeqObjList	2.62937	NumOfObjects=2
└─pulse_shift_delay	SeqDelay	0.07250	
└─pulse_rf	SeqPuls	2.55687	Samples=326, B1=0.02143
└─phase/depth	SeqParallel	2.18621	
└─phase	SeqGradChanParallel	2.18621	
└─phase	SeqGradPhaseEnc	2.18621	
└─phase_grad	SeqGradVector	0.86121	Strength=0.00750, Channel=phase
└─phase_off	SeqGradDelay	0.72500	Strength=0.0, Channel=phase
└─depth	SeqParallel	2.070	
└─depth	SeqAcqDepth	2.070	
└─depth_grad	SeqGradConstRampPulse	2.070	
└─depth_grad_onramp	SeqGradRamp	0.0950	Strength=-0.00514, Channel=read
└─depth_grad_constgrad	SeqGradConst	1.280	Strength=-0.00514, Channel=read
└─depth_grad_offramp	SeqGradRamp	0.0950	Strength=-0.00514, Channel=read
└─acq_read/acq_middelay+acq_acq_acq_tozero	SeqAcqRead	3.54250	
└─acq_read	SeqGradChanParallel	3.340	
└─acq_read	SeqGradConstRampPulse	3.340	
└─acq_read_onramp	SeqGradRamp	0.090	Strength=0.00534, Channel=read
└─acq_read_constgrad	SeqGradConst	2.560	Strength=0.00534, Channel=read
└─acq_read_offramp	SeqGradRamp	0.090	Strength=0.00534, Channel=read
└─acq_middelay+acq_acq_acq_tozero	SeqObjList	3.53250	NumOfObjects=3
└─acq_middelay	SeqDelay	0.13250	
└─acq_acq	SeqAcq	2.710	NumOfObjects=1
└─acq_tozero	SeqDelay	0.690	
└─unnamedSeqDelay	SeqDelay	89.50441	

Fig. 6. The sequence tree of the example sequence from Fig. 1 visualized within the ODIN framework. The first column depicts the structure of the tree whereby the basic sequence objects can be found at the end of each branch and the container objects at the nodes, indicated by small boxes to the left. The second and third column show the C++ type and the duration of each object. Properties that are specific to each object are shown in the last column, e.g., the selected RF object *pulse_rf* has a waveform of 326 samples with the given amplitude B1.

340 4. Hardware-specific implementation

341 In this section, two examples show how the ODIN
342 sequence tree can be utilized to drive the hardware of
343 two scanners from different manufacturers:

344 Platform A (Bruker Medspec, 3T) is driven by a
345 pulse program which is an ASCII file that contains a list
346 of sequential instructions for the hardware and controlling
347 structures (loops, jumps) to repeat certain parts
348 of the sequence. To perform an experiment, a set of
349 parameters must be provided that contains the detailed
350 settings for the measurement. The pulse program and
351 the parameter set cover all characteristics of the exper-
352 iment on this platform. ODIN maps its internal repre-
353 sentation of the sequence to the device by traversing the
354 sequence tree and generating an entry in the pulse pro-
355 gram for each sequence object. In addition, each se-
356 quence object is asked to make an entry into the
357 parameter set. After transferring the generated files to
358 the system controller, the sequence can be executed.
359 Because the pulse program is generated externally on the
360 workstation, the limited memory and speed of the sys-
361 tem controller is not an issue. Even better, different
362 variations of the pulse program, which would usually be
363 implemented by conditional statements in the pulse
364 program itself (if-then-else instructions), are handled by
365 ODIN. Therefore, a minimal pulse program is generated
366 for each experiment containing only the necessary in-
367 structions, thereby reducing the code size which is ac-
368 tually processed by the system controller.

369 On platform B (Siemens Trio, 3T), the system com-
370 ponents are driven directly by a C++ program in real
371 time. The corresponding source code must be provided
372 by the sequence programmer. It contains instructions to
373 trigger hardware events (RF pulses, gradients) at speci-
374 fied points in time. The experiment is performed during
375 run-time of this program. On this platform, ODIN ex-
376 ecutes a sequence by traversing the sequence tree at run-
377 time, querying each sequence object for a corresponding
378 event. An internal counter takes care of the correct
379 starting time of each event. Although the additional lev-
380 el of indirection when using ODIN to trigger the
381 hardware events decreases execution speed a little, it was
382 still fast enough to execute all ODIN sequences, which
383 were tested so far, in real time. An additional amount of
384 memory is required for the ODIN library (typically
385 5 MB) which can be easily accommodated in the free
386 memory (approximately 18 MB) of the used system.

387 In the procedures described above, the connection
388 between the ODIN library and the current platform is
389 realized by a set of so-called *hardware drivers*, as illus-
390 trated in Fig. 2. These hardware drivers are implemented
391 by C++ classes. Each basic sequence object uses a
392 hardware driver to execute itself on the current plat-
393 form. Thereby, the hardware drivers are interchangeable,
394 depending on the hardware to operate. For example,
395 an RF pulse uses different hardware drivers on each
396 platform: the driver for platform A is responsible for
397 an entry in the pulse program and for the pulse-
398 specific parameter settings. The driver for platform B is

399 responsible for preparing and triggering hardware
400 events to execute the RF pulse. The internals of the
401 drivers are hidden behind a common interface (abstract
402 virtual base class in C++) so that there is little coupling
403 between the drivers and the rest of the library. With this
404 design, the code to deal with the peculiarities of each
405 platform is located only within a small set of C++
406 classes. In the case of porting ODIN to a new platform
407 or in the case of a software update by the manufacturer
408 which is accompanied by a considerable change of the
409 sequence programming interface, only these driver
410 classes have to be implemented or updated, the rest of
411 the library and the ODIN sequences remain unchanged.
412 The benefit is straightforward portability to new system
413 types and minimum effort in case of a software update.

414 Usually, the sequence programmer is responsible for
415 manually adding code to calculate the total duration of
416 the sequence or estimating the RF power deposition for
417 safety control in human or animal studies. With the
418 sequence tree in ODIN, which holds all information
419 about the sequence, this tedious process can be com-
420 pletely automated by the library which traverses the
421 sequence tree and queries the objects at each branch for
422 their properties (duration, power deposition). Thereby,
423 simple but error-prone programming tasks are trans-
424 ferred to the ODIN library, allowing the sequence pro-
425 grammer to concentrate on the important features of the
426 sequence.

427 5. Sequence visualization and simulation

428 Even on computers where no NMR device is at-
429 tached, the ODIN framework can be useful for devel-
430 oping sequences. On a stand-alone platform, the time
431 courses of the different channels (RF, gradients, and
432 receiver) can be displayed, or a simulation of the se-
433 quence acting on a virtual sample can be performed.
434 This is achieved by giving all basic sequence objects the
435 capability to generate a digitized version of themselves,
436 i.e., a function that returns the values of each channel
437 for equally spaced points in time.

438 To generate a digitized version of the whole sequence
439 for visualization, the container objects can combine them
440 recursively, traversing the sequence tree until the whole
441 sequence is processed. The result can then be displayed
442 graphically. For simplicity, this is currently realized by
443 generating a multi-channel audio file which is then dis-
444 played using conventional sound editors. In addition,
445 predefined functions exist which calculate important as-
446 pects of the sequence numerically using the digitized se-
447 quence, for example gradient moments, the strength of
448 diffusion weighting or the k -space encoding of different
449 coherence pathways in a multi-pulse sequence.

450 For the simulation, a virtual sample that holds spa-
451 tially resolved NMR-specific properties (spin density,

relaxation rates T_1 and T_2 , and frequency offset) is re- 452
quired. It can be created by means of a graphic editor or 453
a special ODIN sequence that measures these properties 454
of a real sample with a high resolution. The latter will be 455
used in the experimental section of this work to compare 456
the simulation to actual measurements. The digitized 457
version of each sequence object is then used to simulate 458
its effect on the sample. By traversing the sequence tree, 459
the simulation is performed in the same order as the 460
sequence objects would be played out on a real NMR 461
device. An exact solution of the Bloch equations for 462
piecewise constant fields [18] is utilized for the calcula- 463
tion: It transforms the magnetization vector at each 464
point of the sample recursively according to the set of 465
values within the digitized arrays of the sequence object. 466
During acquisition periods, a virtual NMR signal is 467
generated by integrating over the transverse component 468
of the magnetization vector for all points within the 469
virtual sample. The result of the simulation is then a 470
synthetic NMR signal that can be post-processed with 471
the same algorithm as the real signal would be pro- 472
cessed. 473

This simulation strategy is most useful for analyzing 474
imaging sequences. Because it is limited to ensembles of 475
isochromatic spins with single-quantum coherences and 476
interactions simplified by T_1 and T_2 (e.g., quadrupolar 477
coupling, spin-spin coupling), other tools [19–21] are 478
more appropriate to generate virtual spectra of samples 479
with different nuclei, to simulate higher-order quantum 480
coherences or explicit interactions. Another limitation is 481
given by the finite spatial size of the volume elements: 482
The simulation does not account for static intra-voxel 483
dephasing due to field inhomogeneities (T_2^*). 484

6. Data processing

485

In a typical NMR experiment, the RF signal that is 486
induced by the magnetization of the sample and received 487
by the coil is post-processed to obtain interpretable 488
data. This can be a frequency analysis for spectroscopic 489
applications or the reconstruction of spatially resolved 490
parameter maps for imaging. In general, the data pro- 491
cessing algorithm is specific to the NMR sequence which 492
was used to acquire the raw data. This step is supported 493
by a software layer that integrates external numerical 494
libraries consistently into ODIN. 495

After the measurement, the raw data is processed by a 496
function of the same sequence module that was used for 497
the experiment. Because this function is implemented as 498
a C++ member function, all parameters of the mea- 499
surement are directly accessible. The external numerical 500
libraries can be used within this function. After the 501
processing step, the final data is written back to disk. 502

When dealing with large datasets, e.g., for fMRI, the 503
problem arises that the whole record cannot be held in 504

505 memory for analysis at once. ODIN addresses this by
 506 the use of memory paging mechanisms of the underlying
 507 operating system (mmap/munmap functions under Li-
 508 nux/UNIX) so that the array can be accessed transpar-
 509 ently, even if it is too large for the main memory.

510 6.1. Integration of external libraries

511 As a basis for further integration of external libraries
 512 into ODIN, the expression-template based multidimen-
 513 sional array type provided by the Blitz++-library [22] is
 514 used to hold the NMR data during the different pro-
 515 cessing steps. Many useful functions that operate on
 516 multidimensional arrays are already made available by
 517 Blitz++. However, more complex numerical operations
 518 must be added separately as they are not part of
 519 Blitz++. Therefore, an interface to the following li-
 520 braries has been implemented so that they always op-
 521 erate on the array type of Blitz++ and add the described
 522 functionality to it:

- 523 • NewMat [23]: Supports various matrix types and ma-
 524 trix calculations.
- 525 • GSL (GNU Scientific Library) [24]: Non-linear least-
 526 square fitting, interpolation.
- 527 • FFTW (Fastest Fourier Transform in the West) [25]:
 528 Fourier transform for multidimensional arrays.

529 For example, an FFT of arrays with arbitrary dimen-
 530 sionality can be programmed in one line of C++-code
 531 with this integration of external libraries:

```
532 blitz_fftw(data(all, 0, all));
```

533 This instruction will perform a complex in-place FFT
 534 over the first and third dimension of the array `data` for
 535 all values with index 0 in the second dimension.

536 7. Experiments

537 Two sequences were executed with the same subject
 538 and the same settings on platform A and B. Fig. 7 shows
 539 the reconstructed images from a power-reduced variant
 540 of the modified driven equilibrium Fourier transform
 541 (MDEFT) sequence [26]. Although the position of the
 542 brain within the slice differs due to different positioning
 543 of the subject within the magnet, both images show the
 544 same spatial pattern and comparable contrast with a
 545 signal-to-noise ratio of 30.5 (platform A) and 25.1
 546 (platform B) in white matter.

547 In Fig. 8, spin-echo EPI [27] experiments are com-
 548 pared with the result of a simulation which was per-
 549 formed using high-resolution maps of the NMR
 550 parameters (spin density, T_1 , T_2 , and frequency offset).
 551 These maps were acquired on platform A during the
 552 same session. The simulation was then carried out off-
 553 line on a Linux PC to generate a synthetic signal using
 554 the same sequence code that was used for the measure-
 555 ments. The images are similar in terms of contrast and




Fig. 7. MDEFT images from platform A (top) and B (bottom) with a matrix size of 252×252 pixels, $FOV = 220$ mm and a sweep-width of 25 kHz. This sequence type is highly sensitive to the T_1 relaxation time. Therefore it is well-suited to display anatomical structures.

556 image quality, but show slightly different field-of-views
 557 in phase encoding direction which is very sensitive to
 558 frequency offsets due to the small bandwidth. The mis-
 559 match may therefore be caused by non-optimal com-
 560 pensation of the field inhomogeneities (shimming) or
 561 eddy-currents modifying the phase encoding blips. This
 562 otherwise undesired discrepancy could be used here to
 563 study the effects of field variations and gradient imper-
 564 fections. However, the general similarities between the
 565 result of the simulation and the actual experiments in-
 566 dicate that the simulation can be used to reproduce the
 567 measurement and that it is feasible to develop and test
 568 sequences on a stand-alone platform.

569 8. Availability and licensing

570 The software package is published under the terms of
 571 the GNU General Public License. It can be obtained as
 572 source code and binary packages for different platforms
 573 (Linux, IRIX, Windows, and VxWorks) from the web




Fig. 8. Spin-echo EPI images from platform A (top), platform B (middle), and the simulation (bottom) with a matrix size of 64×64 , 100 accumulations, 100 kHz sweep-width and the same slices as in Fig. 7. The phase encoding direction is aligned vertically.

574 [28]. The online manual for the class hierarchy can also
575 be found at this location.

576 9. Conclusion

577 A cross-platform environment for developing NMR
578 sequences has been presented. The sequence program-
579 ming interface provides a concise C++ class hierarchy to
580 set up an NMR experiment within a short time. Without
581 changing the source code, the sequence can be visual-
582 ized, simulated, and executed on different NMR hard-
583 ware. This is particularly useful in laboratories where

more than one scanner exists, or to exchange sequences 584
between research facilities with different hardware in- 585
frastructure. With the ODIN data processing frame- 586
work, a consistent interface to reliable open-source 587
libraries for calculating the final data is provided. The 588
internal representation of the experiment by the se- 589
quence tree is adequately matched to the application 590
domain and allows easy extensibility when porting the 591
framework to new platforms. 592

Acknowledgments 593

The authors thank Robert Trampel, Markus Körber, 594
and Andreas Schäfer for improving the software and 595
Harald E. Möller for helping with the manuscript. 596

References 597

- [1] B. Stroustrup, The C++ Programming Language, Addison- 598
Wesley, Boston, 2000. 599
- [2] J. Debbins, K. Gould, V. Halleppanavar, J. Polzin, M. Radick, G. 600
Sat, D. Thomas, S. Trevino, R. Haworth, Novel software 601
architecture for rapid development of magnetic resonance ap- 602
plications, *Conc. Magn. Reson. (Magn. Reson. Engin.)* 15 (3) (2002) 603
216–237. 604
- [3] Available from <<http://java.sun.com/java2/whatis>>. 605
- [4] T.H. Jochimsen, M. von Mengershausen, T. Mildner, D.G. 606
Norris, Separating the contributions to the intravascular signal 607
change in spin-echo fMRI at 3 Tesla, *Eur. Soc. Mag. Reson. Med.* 608
Biol. 20 (2003) 146. 609
- [5] T.H. Jochimsen, H.E. Möller, D.G. Norris, Is there a change in 610
spin density associated with fMRI? *Proc. Intl. Soc. Mag. Reson.* 611
Med. (2004) (in press). 612
- [6] A. Schäfer, T.H. Jochimsen, H.E. Möller, fMRI with intermolec- 613
ular double-quantum coherences (iDQC) at 3T, *Proc. Intl. Soc.* 614
Mag. Reson. Med. (2004) (in press). 615
- [7] R. Trampel, T.H. Jochimsen, T. Mildner, D.G. Norris, H.E. 616
Möller, Efficiency of flow-driven adiabatic spin inversion under 617
realistic experimental conditions: a computer simulation, *Magn.* 618
Reson. Med. (2004) (in press). 619
- [8] N.P. Davies, P. Jezzard, Selective arterial spin labeling (SASL): 620
perfusion territory mapping of selected feeding arteries tagged 621
using two-dimensional radiofrequency pulses, *Magn. Reson. Med.* 622
49 (2003) 1133–1142. 623
- [9] S. Conolly, D. Nishimura, A. Macovski, G. Glover, Variable-rate 624
selective excitation, *J. Magn. Reson.* 78 (1988) 440–458. 625
- [10] M.S. Silver, R.I. Joseph, D.I. Hoult, Highly selective $\pi/2$ and π 626
pulse generation, *J. Magn. Reson.* 59 (1984) 347–351. 627
- [11] E. Kupče, R. Freeman, Adiabatic pulses for wideband inversion 628
and broadband decoupling, *J. Magn. Reson. A* 115 (1995) 273– 629
276. 630
- [12] C.H. Meyer, J.M. Pauly, A. Macovski, D.G. Nishimura, Simul- 631
taneous spatial and spectral selective excitation, *Magn. Reson.* 632
Med. 15 (1990) 287–304. 633
- [13] J. Pauly, D. Nishimura, A. Macovski, A k -space analysis of small- 634
tip-angle-excitation, *J. Magn. Reson.* 81 (1989) 43–56. 635
- [14] M. Levitt, Symmetrical composite pulse sequence for NMR 636
population inversion. I. Compensation of radiofrequency field 637
inhomogeneity, *J. Magn. Reson.* 48 (1982) 234–264. 638
- [15] T.H. Jochimsen, D.G. Norris, Single-shot curved slice imaging, 639
MAGMA 14 (2002) 50–55. 640

- 641 [16] Available from <<http://www.trolltech.com>>. 655
- 642 [17] A.N. Davies, P. Lampen, JCAMP-DX for NMR, Appl. Spectr- 656
- 643 soc. 47 (8) (1993) 1093–1099. 657
- 644 [18] H.C. Torrey, Transient nutation in nuclear magnetic resonance, 658
- 645 Phys. Rev. 76 (8) (1949) 1059–1068. 659
- 646 [19] S.A. Smith, T.O. Levante, B.H. Meier, R.R. Ernst, Computer 660
- 647 simulations in magnetic resonance. An object-oriented program- 661
- 648 ming approach, J. Magn. Reson. A 106 (1994) 75. 662
- 649 [20] P. Nicholas, D. Fushman, V. Ruchinsky, D. Cowburn, The virtual 663
- 650 NMR spectrometer: a computer program for efficient simulation 664
- 651 of NMR experiments involving pulsed field gradients, J. Magn. 665
- 652 Reson. 145 (2000) 262–275. 666
- 653 [21] W.B. Blanton, BlochLib: a fast NMR C++ tool kit, J. Magn. 667
- 654 Reson. 162 (2003) 269–283. 668
- [22] T.L. Veldhuizen, Arrays in Blitz++, in: Proceedings of the Second 655
- International Scientific Computing in Object-Oriented Parallel 656
- Environments (ISCOPE'98), Lecture Notes in Computer Science, 657
- Springer-Verlag, 1998. 658
- [23] R. Davies, Writing a matrix package in C++, in: The Second 659
- Annual Object-Oriented Numerics Conference, 1994, pp. 207–213. 660
- [24] Available from <<http://www.gnu.org/software/gsl>>. 661
- [25] M. Frigo, S.G. Johnson, FFTW: An Adaptive Software Archi- 662
- ture for the FFT, 1998, pp. 1381–1384. 663
- [26] D.G. Norris, Reduced power multislice MDEFT imaging, J. 664
- Magn. Reson. Imag. 11 (4) (2000) 445–451. 665
- [27] P. Mansfield, Multi-planar image formation using NMR spin 666
- echoes, Solid State Phys. 10 (1977) L55–L58. 667
- [28] Available from <<http://od1n.sourceforge.net>>. 668

UNCORRECTED PROOF