

Unentangling nuclear magnetic resonance computing

Matthias Bechmann¹, John A. Clark², Angelika Sebald¹, Susan Stepney²

¹ Department of Chemistry,

² Department of Computer Science,
University of York, UK, YO10 5DD

Abstract. Nuclear magnetic resonance (NMR) is typically thought of as a possible technology for quantum computation. Here we instead outline how commercially available NMR spectrometers could be used to perform non-quantum computation: from addressable 3D memory, to a programmable 3D reaction-diffusion computer.

1 Introduction

Recently, exciting border-crossing joint activities between computer sciences and nuclear magnetic resonance (NMR) spectroscopy have emerged. From the point of view of NMR spectroscopy, one would probably first think of recent applications of genetic algorithms for finding improved pulse sequences in solution- and solid-state NMR [8] [9]. In the context of computer science, probably NMR first comes to mind as a tool for implementing and testing ideas in the area of quantum computing [5] [7]. NMR can play this role due to its uniquely well defined quantum mechanical nature and the relative ease of (some) experimental implementations. Commercial NMR spectrometers are widely available, for example providing the prime analytical technique in synthetic organic chemistry, or as a slightly more sophisticated method for the study of molecular structures in biochemistry.

Here we describe some ideas for mutually beneficial interplay between NMR and unconventional computation other than quantum computing. We first give a brief summary of the most basic features and concepts encountered in (solution-state) NMR spectroscopy. Then we propose various computational approaches, and outline how such concepts could be implemented in a real NMR spectrometer.

2 Background: The physics of NMR

Here we give a brief and qualitative sketch of the basic principles involved in NMR; we refer the reader to the literature [11] for further reading about the physical background, the theory and quantum mechanical description of NMR.

Nuclear magnetic resonance exploits an intrinsic property of most isotopes in the period table of the elements: the magnetic moments of atomic nuclei.

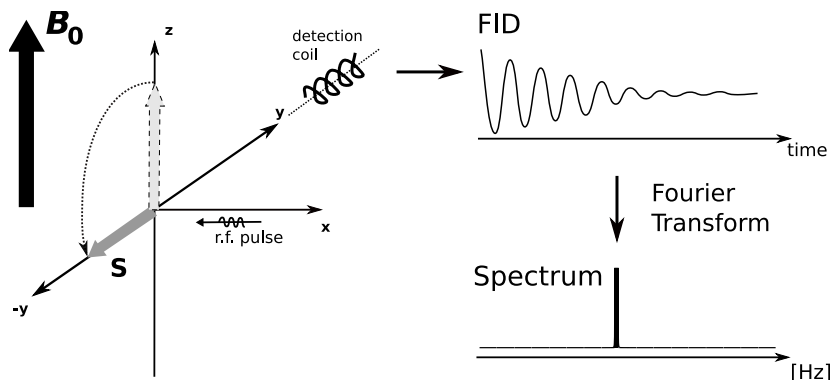


Fig. 1. Larmor precession in the rotating frame, resonance, FID detection and frequency-domain spectrum.

Associated with this magnetic moment is a property called spin. For the sake of simplicity we consider only the case of isotopes with spin quantum number $I = 1/2$. In addition, we initially assume that we have a homogeneous sample with only one type of spins present. The sample is placed in a very strong and homogeneous external magnetic field B_0 along the z -direction of the laboratory coordinate system (see figure 1). The spins start precessing around the B_0 direction at the so-called Larmor frequency (which is isotope specific) and a net magnetisation of the spins in the sample builds up along the z -direction. Adopting a reference frame rotating at the Larmor frequency around the z -direction yields a description where the spin dynamics appear static. Therefore, we can now easily inspect the effect of a radio frequency (r.f.) pulse along the x -direction of the rotating frame as the effect of an additional static magnetic field B_1 in the x -direction of the rotating frame on the magnetisation vector. During the pulse, the spins precess around the direction of the applied field B_1 : for example, an r.f. pulse of suitable amplitude and duration at the Larmor frequency tips the magnetisation vector in the $-y$ -direction (a so-called $(\pi/2)_x$ -pulse, where the subscript indicates the direction/phase of the r.f. pulse), giving the maximum possible signal in a detection coil aligned in the y -direction. (See figure 2 for examples of the effects of various pulses.)

After switching off the r.f. pulse, the signal is observed in the time domain in the form of a free induction decay (FID) in the receiver coil (see figure 1). Commonly, this time domain signal is transformed to the frequency domain by applying a complex Fourier transformation. By varying the phase, amplitude and/or duration of the r.f. pulse, the magnetisation vector can be tipped in any direction – and the corresponding NMR spectrum carries frequency as well as phase information. This is the simplest form of a common, one-dimensional (1D) NMR spectrum.

The ‘burst of radiation’ in the form of the r.f. pulse leaves the system in a non-equilibrium state. Accordingly, after the pulse the system returns to equi-




Fig. 2. Illustration of the effects of different r.f. pulses on magnetisation vectors S : (a) a $(\pi/2)_y$ -pulse acting on the $+z$ -magnetisation vector; (b) a $(\pi)_y$ -pulse acting on the $+x$ -magnetisation vector; (c) a $(\pi)_{45}$ -pulse acting on the $+x$ -magnetisation vector.

librium z -magnetisation. This happens with a characteristic time constant T_1 and measures the rate at which z -magnetisation is re-built. Another effect, with a characteristic time constant T_2 , concerns the xy -plane: initially, right after an r.f. pulse, all the spins are in phase. Small fluctuating local fields cause a dephasing in the xy -plane at a rate measured by T_2 . In non-viscous liquids at ambient conditions usually $T_1 \sim T_2$. As T_1 and T_2 are of the order of seconds, NMR resonances under these conditions tend to be very narrow, giving very high spectral resolution.

Now we add another feature. Suppose we add magnetic field gradients (in the x, y, z directions, either static or pulsed gradients) to the external static magnetic field. Such gradients alter the resonance frequencies of spins in different locations (voxels) of the sample. In our simple system of a homogeneous liquid sample with only one type of spins present, different voxels of the sample are now represented by different resonance frequencies – as the spins in different regions of the sample feel different local magnetic fields, depending on the gradients used. Thus, gradients add space (and further phase) encoding to the scenario. If we now take advantage of selective r.f. pulses, we have enormous degrees of freedom to select slices and specific regions in the sample for all kinds of ‘spin gymnastics’. A practical, and probably the most widely known, application of these principles is magnetic resonance imaging (MRI) in the medical context where NMR signals of ^1H nuclei in fat and water, our predominant soft body materials, deliver three-dimensional images. Although keeping these principles in mind, we do not aim at using them for image production, but rather aim to exploit them in a sense of localised spectroscopy.

Having generally available now frequency, phase and space encoding we need to expand the 1D NMR spectrum to include a second dimension. 2D NMR needs the application of a sequence of pulses and/or events in order to induce dependency of the NMR signal(s) on two time variables. These two time variables, along with different kinds of pulse sequences, are able to deliver 2D NMR spectra displaying a wide range of e.g. correlations amongst spins, after complex Fourier transformation of the time-domain data with regard to two time variables. In that respect, 2D NMR is not just a simple expansion of 1D NMR but adds a fundamentally important additional feature. (Note that here the ‘2D’ refers to two *temporal* dimensions, not to spatial dimensions.)

If, in gedankenexperiments, we now add more complexity to our samples (e.g. by including samples with more than one spin species present, by including interactions amongst spins by dipolar coupling interactions, by including samples with anisotropic properties leading to anisotropic signatures in NMR spectra) plus allow for virtually unlimited complexity of pulse (and gradient switching) sequences, it takes little imagination to see that we have a highly controllable, extremely well defined – and highly complex – system at our disposal.

We claim that this system can, and should, be exploited for further experiments, beyond the implementation of quantum-computing model systems. As noted earlier, NMR is one proposed approach to quantum computation [5] [7]. There, the interactions between nuclear spins (couplings) are used to encode

qubits, and entangling those spins gives qubits. However, like all quantum computation implementations, there are the usual engineering problems of maintaining coherence and scaling up to large numbers of qubits. Here we take a step back, and investigate using the nuclear spins to encode *classical* bits, and use liquid-state NMR spectroscopy to do classical (as in non-quantum) massively parallel computation.

3 Encoding bits

The liquid sample is divided into spatial regions, or *voxels*, by the field gradients and associated resonant frequencies: reading spins at discretised frequencies corresponds to a discretised space. Each voxel thus defined can encode a bit (or multiple bits from different spins in a complex molecule).

As described above, the basic ingredients for encoding bits are the frequency, duration/amplitude and phase of r.f. pulses, together with the magnetisation vector in the absence or presence of further r.f. pulses. We need to distinguish two different scenarios: a ‘quasi-static’ case which, in NMR terms would be referred to as a ‘on resonance’ condition, and a ‘quasi-dynamic’ case which, in NMR terms would be called a ‘off resonance’ case. The former condition permits a time-independent operation whereas the latter condition is more complicated in that time becomes an integral part of the logic.

3.1 The on-resonance, quasi-static condition

We start from equilibrium magnetisation (along the $+z$ -direction) with a $(\pi/2)$ -preparation pulse, for example (see figure 3) a $(\pi/2)_y$ -pulse giving $+x$ -magnetisation, or a $(\pi/2)_x$ -pulse giving $-y$ -magnetisation. We define the areas in the xy -plane covering $\pm 45^\circ$ around the $+x$ - and $-x$ -directions as representing 1, the remaining areas as 0 (see figure 3; this choice of areas will become obvious in the following).

Thus, the phase and the duration/amplitude of the preparation $(\pi/2)$ r.f. pulse permits us to prepare a 1 or a 0 state. Here we have chosen the r.f. frequency such that it is *precisely* on resonance, with the consequence that the magnetisation vectors in the xy -plane, again, appear static (‘they stay where we put them’, loosely speaking). We can now introduce the second bit (see figure 4): a $(\pi)_{45}$ -pulse P (1), or no pulse (0).

Application of the $(\pi)_{45}$ -pulse changes the first bit from 1 to 0, whereas in the absence of this pulse the first bit remains in its 1 state. If the presence or absence of the r.f. pulse is controlled by the value of another bit, the result is a CNOT gate on those two bits (see table 1). Currently this requires some computation external to medium itself, to physically control the r.f. pulse; in section 5.3 we also discuss computation intrinsic to the medium.




Fig. 3. Definition of the '1' and '0' regions in the xy -plane (left) and some spectra corresponding to various different $\pi/2$ preparation r.f. pulses (right).




Fig. 4. The effect of presence (top) and absence (bottom) of the $(\pi)_{45}$ -pulse on the $+x$ -magnetisation vector.

pulse (control bit)	target bit (spin)	final spin
off (0)	0	0
off (0)	1	1
on (1)	0	1
on (1)	1	0

Table 1. Components of a CNOT gate.




Fig. 5. The effect of an off-resonance preparation pulse as $f(\text{time})$.

3.2 The off-resonance, quasi-dynamic condition

We start again with equilibrium magnetisation and a $(\pi/2)$ -preparation pulse, flipping the magnetisation vector to the xy -plane, as is illustrated in figure 5. The only difference to the previous on-resonance condition is that now we chose a r.f. pulse not exactly on resonance but strong enough to flip or invert the magnetisation vector. The consequence of this frequency mismatch between Larmor frequency and r.f. frequency is the loss of time-independence in the rotating-frame picture (see figure 5): now the magnetisation vector appears either ‘slower’ or ‘faster’ than the reference frame, depending on the choice of r.f. frequency.

The amount of frequency mismatch obviously determines how fast the magnetisation vector rotates in the xy -plane. As a function of time, the magnetisation vector now periodically travels through all the ‘0’ and ‘1’ regions.

If, for example, we record the NMR response after a period τ sufficient for the vector to rotate by 90° , a previously 1 bit will have become a 0 bit, and vice versa. Hence, this time shift (clock shift) corresponds to a NOT operation. From the bottom part of figure 6 we can further see that at specific points in time we can create exactly those conditions we had previously considered in the on-resonance case, and in full analogy to that situation we can again apply (or omit) $(\pi)_{45}$ -pulses in order to selectively negate bits – but not only at these specific points in time. The off-resonance situation adds an element of choice (time / clock) to the previous on-resonance condition, in addition to a further degree of (potential) complication as the rotation of the magnetisation vector in the xy -plane can be of either positive or negative sign, and the sign is retained after a $(\pi)_{45}$ -pulse.




Fig. 6. The events in the off-resonance scenario as a $f(\text{time})$ without (top) and with (bottom) application of a $(\pi)_{45}$ -pulse.

We need to take the off-resonance situation into account for a variety of reasons. We might have a sample with more than one spin species present – in that case, only one spin species can be in the on-resonance condition at a time. Alternatively, we may have a sample with only one spin species present but may alter the resonance frequency in different voxels of the sample by applying field gradients. Or we may be interested in implementing logics which require time dependence of the spin dynamics in a voxel.

3.3 Implementation

Any of the above operations and bits based on phases of pulses and magnetisation vectors are straightforward to implement on contemporary NMR spectrometers; they can be carried out on liquid as well as on solid samples; and the necessity of having isolated spin species available does not impose any severe restrictions on the choice of suitable samples for practical implementations.

4 3D Memory

Conventional solid state memory chips have a limitation: one can extract memory values from the physical edges only, so memory values stored in the middle of the chip need to be moved to the edges before they can be read. This limits i/o bandwidth. This is a well-known limitation of CCD chips, for example.

Mechanisms that can directly address ‘internal’ memory bits usually involve electromagnetic radiation (such as light), as this requires no wires. For example, the bacteriorhodopsin protein is being developed as a three dimensional optical storage memory [3] [4] [16]. This protein has two stable states that can be set by pulses of different frequency laser light, with a switching time of ~ 0.5 ns. So a protein molecule can store a bit: in practice several thousand are used per bit. The theoretical storage density is also limited by the wavelength of the laser

light. The quoted theoretical limit is 10^{18} bits m^{-3} , with practical limits being a few orders of magnitude lower. Each bit can be addressed directly by suitably focussing the lasers. Parallel access to a 2D slice of the memory is achieved by arrays of read/write lasers.

The NMR spin states in voxels could similarly be used as a directly 3D addressable RAM. The molecules involved are smaller and simpler than the bacteriorhodopsin protein. The spins equilibrate on the (long) T_1 timescale, and so the memory would need to be refreshed periodically.

The potential data density is determined by the NMR voxel size, determined by the spectroscopy resolution available. Current resolution is on the order of 1 mm^3 for medical imaging applications (MRI). Localised spectroscopy can achieve resolutions $1 \mu\text{m}^3$; current sample sizes of $\sim 100 \text{ mm}^3$ could store $\sim 10^8$ bits at these resolutions. The total sample size is limited by the volume over which the magnetic field B_0 can be made homogeneous. This resolution implies a data density of 10^{18} bits m^{-3} , comparable to the quoted bacteriorhodopsin figures. Currently it is significantly less than that, of course: the sample that actually stores the data is dwarfed by the vast size of the surrounding superconducting magnet ($\sim 1 \text{ m}^3$) and r.f. pulse generator ($\sim 1 \text{ m}^3$). (Smaller non-cryogenic NMR technologies also exist, such as the portable “NMR mouse” [6], which is moved over a surface to analyse its composition, providing “mobile NMR”.)

Memory voxels are written by r.f. pulses, which can be delivered every $1 - 5 \mu\text{s}$. The number of pulses required to write multiple voxels depends on the complexity of the data being written: very regular patterns of data can be written with a single pulse, more complex patterns require a pulse sequence. The writing time scales with the data complexity (limited by the ability to generate the complex r.f. pulse sequences), not with the sample size (except for the increased power requirements of the pulse). The reading can be performed in a constant time ($\sim 1 \text{ ms}$) per slice of memory. The read/write r.f. system can be easily modulated to address only parts of the memory, with consequent gains in speed.

5 Computation

Beyond storing data in the medium, we would like to be able to perform computations on that data using the physical properties of the medium. This requires the data encoded in different bits to interact in a data-dependent manner. For interactions between different bits encoded in a single voxel (for example, interactions between spins in a complex molecule) this can be done using the same encodings as NMR quantum computation (see for example [15]), but without quantum entanglement being used to give superpositions of multiple configurations.

We would also like the the data in different voxels to interact in a data-dependent manner. We have seen an example earlier (the CNOT gate), where the value of one bit controlled the application of an r.f. pulse on another bit. That design requires some external decision making (read one bit, and on the




Fig. 7. Modulating the voxel cell size between steps




Fig. 8. Modulating the voxel cell position between steps

basis of its value, choose whether or not to fire the r.f. pulse at another bit.) We would also like an *intrinsic* mechanism for the bits to interact directly within the medium. This is difficult to do in general (hence the dearth of, say, NMR-based quantum cellular automata). However, there is a variety of methods for allowing data in adjacent voxels to interact through non-linear interactions of adjacent phases. These include modulating the logical grid, and physical movement (diffusion) of the molecules.

5.1 Grid modulation

The size and positions of the voxels in the grid can be programmatically changed by modulating the magnetic field.

One can alter the size of the grid cells between timesteps, alternately enlarging the cells (thereby combining the different information contents of several smaller cells) and reducing them (thereby spreading the information across several (smaller) cells (see figure 7).

One can alter the position of the grid between timesteps, changing the offsets. For example, in classical cellular automata (CAs) the Margolus neighbourhood [14] is used to simulate particle dynamics. Here the entire grid is partitioned into neighbourhoods of 2×2 cells. The CA rule is applied to each neighbourhood. Then, on the following timestep, the neighbourhoods are offset by $(1, 1)$, allowing a different set of four cells to interact (see figure 8).

Given each NMR voxel is bulk matter, one could employ an analogous process with half-integer grid movements, allowing corners of adjacent cells in one timestep to become members of the same cell in subsequent timesteps. The re-

sulting bulk spin (and hence bit value) would depend on some combined or average property of the individual spins.

Then combinations of size and position changes can be used to provide programmatic control over information interactions.

5.2 Field gradients as parameters

Varying the gradient of the magnetic field allows vertical slices to experience a different value of the field. This can be used to provide a parameter value, and perform an array of 2D calculations in parallel, with identical computations, but different parameter values.

5.3 Diffusion

In solids, spins of neighbouring molecules interact by dipolar coupling, with the result that the spin state diffuses through the system. Also, ions can physically diffuse through crystal lattices, at a much slower rate than the spin diffusion rate. The resultant bulk spin associated with the voxel is a non-linear combination of the original spin, the diffusing spin states, and the diffusing molecules.

The result is a **programmable 3D reaction-diffusion computer**. The ‘reaction’ is that of spins in combination, so is of a strictly limited kind, unlike chemical reaction-diffusion (RD) computers, which have a wide range of non-linear chemical reactions available [2]. Chemical RD computers have relatively limited interactivity: they tend to be programmed through the initial distribution of chemical reagents, with some opportunity of adding spots of chemicals as the reaction proceeds, or (usually passively) illuminating patterned areas over the light-sensitive reactions [10] [13]. The NMR 3DRD computer has the potential for detailed programmability: a ‘program’ consists of a sequence of r.f. pulses; each timestep a programmed r.f. pulse can manipulate each of the bits in the medium. The choice of paths through the program can be a function of the current state of the system (read out in the receiver coil). Clearly, if this function of current state is being calculated out of the medium, in the conventional computer controlling the r.f. pulse program, it needs to be a relatively simple calculation, such as a property of a few voxels, or certain bulk properties of the entire medium. The computation performed is an emergent property of all these actions. See figure 9.

One particular use of the programmability would be to isolate certain volumes by the use of continually zeroed voxels while part of a computation was being performed, then reprogramming these voxels to allow interactions to occur between the regions in another part of the computation.

6 Variations on a theme

6.1 Other materials

So far we have restricted most of the discussion to isotropic liquids with simple spin states (for example, the ^1H spins in water), and simple solids.




Fig. 9. Schematic architecture of the programmable 3D NMR reaction-diffusion computer

Molecules with more complex spin states exist, particularly the ones used for quantum NMR computing, where the different spin states per molecule are used to encode different qubits. However, even in a relatively simple two-spin state molecule, the Larmor frequencies are different, and the spins interact in a complex non-linear manner that could potentially be exploited for computation.

Inhomogeneous materials with complex structures, for example, layers with different properties, can be constructed, with gels or solids. These can have anisotropic spin diffusion rates, that again could be exploited for different kinds of computation.

6.2 Dynamically programmable matter

The consideration of complex materials leads on to the possibility of dynamically programmable matter. One requires a material where the spin state changes some interesting properties of the material. Clearly it changes the r.f. properties, as this is what is used to read out the current state: there may be other properties that are also usefully changed.

One could then program the medium to have the desired properties to perform some task, with those properties being a function of time. That is, the matter could be programmed to react, or adapt, to changing circumstances. For example, consider adaptive optics. These are used in astronomy to remove the

twinkling caused by atmospheric fluctuations, which otherwise produce blurred images from ground-based observations. This is achieved by computer controlled flexible mirrors, or variable refraction (for an introduction, see for example [12]). The longer the wavelength, the easier this is to do, so the r.f. domain may be an ideal application.

Other programmable matter applications will become apparent as the capability of the NMR computer, particularly with complex materials, is explored further.

6.3 Breaking the model

One of the major advantages of the proposed approach is the availability of commercial off the shelf NMR spectrometers, and the associated laboratory expertise. However, some of the design properties of those NMR spectrometers are to enable conventional NMR spectroscopy experiments to be performed effectively in a homogeneous environment. These design constraints might not be suitable for moving to full NMR computers that could exploit information processing capabilities of an inhomogeneous environment, such as inhomogeneous magnetic fields (there is a whole sub-branch of NMR dealing with inhomogeneous magnetic fields, applied to, for example, oil exploration), or space and time dependent temperature fluctuations in the medium. Much of the time, effort, and expense, of current instruments goes into controlling these properties: if NMR computers can operate (and operate more effectively) in less homogeneous environments, then there is even greater potential for smaller, lower cost devices.

In our initial proposal, we are looking at discretised 3D space (voxels) and time (Larmor period clocking). However, we will also need to look at the possibilities offered by considering the medium to be a continuum existing in continuous time. This could potentially offer a simpler design of the r.f. pulse structure. However, all the usual problems of classical analogue computing will also become apparent, particularly the sensitivity to noise. So initially we are concentrating on the discretised spatial and temporal domain.

7 Next Steps

It is important to recognise that we are not proposing the NMR computer as a *universal* computer in the Turing sense. We want to understand what the physics of NMR offers computation: what the computation physical system can do in a ‘natural’ manner. We want to exploit its physical properties, not to shoe-horn it into an unnatural (for it) computational paradigm. Hence we have *not* done the usual first step in a novel computational domain of demonstrating how the device can be made to emulate a variety of logic gates. The exception here is the NOT gate implementation (section 3), since there are very natural ways to achieve this in the NMR setting.

Given that we are not striving for universality, what is needed is a classification of the computational capabilities of an NMR computer, and the application areas for which it is best suited.

NMR is an incredibly rich domain. We propose to start simply, with the following prototypes in simple media:

1. prototype an implementation of a 3D memory in water, and measure:
 - the memory capacity (achievable spatial resolution)
 - the write rate (r.f. pulse programmability, complexity and frequency)
 - the read rate (read cycles, exploitation of phase)
 - the required refresh rate (relaxation timescale)
 - the error rates
2. prototype an implementation of a programmable 3D reaction-diffusion computer in a simple solid, and measure:
 - the spin diffusion rates in various media
 - the ion diffusion rates in various media
 - the logical operations implemented by these diffusions, with different grid modulation options
 - the constraints on the computation in the controlling r.f. pulse generator
 - the computational tradeoffs between the restricted spin reaction types, and the flexible data programmability
3. prototype an implementation of a hybrid NMR-silicon computer, balancing the computation in the medium, and in the external computer, exploiting the strengths of each approach
4. investigate suitable programming styles to encompass these various uses (including the use of evolutionary algorithms to program an unfamiliar and complex device)
5. prototype a range of applications to help determine computational and application classes (based on existing chemical RD examples [2], such as Voronoi diagram calculation, and on amorphous computing [1] applications, and how these can be extended to exploit the 3D and programmability capabilities of the NMR computer)

8 Conclusions

We have described how a commercially available NMR spectrometer could be used to perform (non-quantum) computation: from addressable 3D memory, to a programmable 3D reaction-diffusion computer. We make no claims as to its Turing universality, or otherwise: rather we seek to exploit the computational abilities naturally occurring in interacting (but not necessarily entangled) switchable protons spins in bulk media.

Acknowledgements

Financial support of our work (MB and AS) by the Deutsche Forschungsgemeinschaft is gratefully acknowledged.

References

1. Harold Abelson, Don Allen, Daniel Coore, Chris Hanson, George Homsy, Thomas F. Knight, Radhika Nagpal, Erik Rauch, Gerald J. Sussman, and Ron Weiss. Amorphous computing. *CACM*, 43(5):74–82, 2001.
2. Andrew Adamatzky, Benjamin De Lacy Costello, and Tetsuya Asai. *Reaction-Diffusion Computers*. Elsevier, 2005.
3. Robert R. Burge. Protein-based optical computing and memories. *Computer*, 25(11):56–67, November 1992.
4. Robert R. Burge. Protein-based computers. *Scientific American*, 272(3):66–71, March 1995.
5. D. G. Cory, A. F. Fahmy, and T. F. Havel. Ensemble quantum computing by NMR spectroscopy. *Proc. Natl. Acad. Sci. USA*, 94:1634–1639, 1997.
6. G. Eidmann, R. Savelsberg, P. Blümner, and B. Blümich. The NMR MOUSE[®]: A mobile universal surface explorer. *J. Magn. Reson. A*, 122:104–109, 1996.
7. N. A. Gershenfeld and I. L. Chuang. Bulk spin-resonance quantum computation. *Science*, 275:350–356, 1997.
8. C. Kehlet, T. Vosegaard, N. Khaneja, S. J. Glaser, and N. C. Nielsen. Low-power dipolar recoupling in solid-state NMR developed using optimal control theory. *Chem. Phys. Lett.*, 414:204–209, 2005.
9. K. Kobzar, T. Skinner, N. Khaneja, S. J. Glaser, and B. Luy. Exploring the limits of broadband excitation and inversion pulses. *J. Magn. Reson.*, 170:236–243, 2004.
10. L. Kuhnert, K. I. Agladze, and V. I. Krinsky. Image processing using light-sensitive chemical waves. *Nature*, 337:244–247, 1989.
11. Malcolm H. Levitt. *Spin Dynamics: Basics of Nuclear Magnetic Resonance*. Wiley, 2001.
12. Jason Porter, Hope Queener, Julianna Lin, and Karen Thorn. *Adaptive Optics for Vision Science: Principles, Practices, Design and Applications*. Wiley, 2006.
13. N. G. Rambidi and D. Yakovenchuk. Chemical reaction-diffusion implementation of finding the shortest paths in a labyrinth. *Phys.Rev.E*, 63:026607, 2001.
14. Tommaso Toffoli and Norman H. Margolus. *Cellular Automata Machines: a new environment for modeling*. MIT Press, 1987.
15. Lieven M. K. Vandersypen, Costantino S. Yannoni, and Isaac L. Chuang. Liquid state NMR quantum computing. In David M. Grant and Robin K. Harris, editors, *Encyclopedia of Nuclear Magnetic Resonance, Volume 9: Advances in NMR*, pages 687–697. Wiley, 2002.
16. B. Xi, Kevin J. Wise, Jeffrey A. Stuart, and Robert R. Birge. Bacteriorhodopsin-based 3D optical memory. In V. Renugopalakrishnan and Randolph V. Lewis, editors, *Bionanotechnology: Proteins to Nanodevices*, pages 39–59. Springer, 2006.