

COMPUTATIONAL TASKS IN MEDICAL NANOROBOTICS

Robert A. Freitas Jr.

Nanomedicine is the application of nanotechnology to medicine: the preservation and improvement of human health, using molecular tools and molecular knowledge of the human body. Medical nanorobotics is the most powerful form of future nanomedicine technology. Nanorobots may be constructed of diamondoid nanometer-scale parts and mechanical subsystems including onboard sensors, motors, manipulators, power plants, and molecular computers. The presence of onboard nanocomputers would allow *in vivo* medical nanorobots to perform numerous complex behaviors which must be conditionally executed on at least a semiautonomous basis, guided by receipt of local sensor data and constrained by preprogrammed settings, activity scripts, and event clocking, and further limited by a variety of simultaneously executing real-time control protocols. Such nanorobots cannot yet be manufactured in 2007 but preliminary scaling studies for several classes of medical nanorobots have been published in the literature. These designs are reviewed with an emphasis on the basic computational tasks required in each case, and a summation of the various major computational control functions common to all complex medical nanorobots is extracted from these design examples. Finally, we introduce the concept of nanorobot control protocols which are required to ensure that each nanorobot fully completes its intended mission accurately, safely, and in a timely manner according to plan. Six major classes of nanorobot control protocols have been identified and include operational, biocompatibility, theater, safety, security, and group protocols. Six important subclasses of theater protocols include locational, functional, situational, phenotypic, temporal, and identity control protocols.

15.1. INTRODUCTION

Nanotechnology is the engineering of molecularly precise structures and, ultimately, molecular machines. Nanomedicine [1, 2] is the application of nanotechnology to medicine: the preservation and improvement of human health, using molecular tools and molecular knowledge of the human body. Nanomedicine encompasses at least three types of molecularly precise structures: nonbiological nanomaterials, biotechnology materials and organisms, and nonbiological devices including inorganic nanorobotics. In the near term, the molecular tools of nanomedicine will include biologically active nanomaterials and nanoparticles having well-defined nanoscale features. In the mid-term (5–10 years), knowledge gained from genomics and proteomics will make possible new treatments tailored to specific individuals, new drugs targeting pathogens whose genomes have been decoded, and stem cell treatments. Genetic therapies, tissue engineering, and many other offshoots of biotechnology will become more common in therapeutic medical practice. We may also see biological robots derived from bacteria or other motile cells that have had their genomes re-engineered and reprogrammed, along with artificial organic devices that incorporate biological motors or self-assembled DNA-based structures for a variety of useful medical purposes.

In the farther term (2020s and beyond), the first fruits of molecular nanorobotics [3]—the most efficacious of the three classes of nanomedicine technology, though clinically the most distant and still mostly theoretical—should begin to appear in the medical field. These powerful therapeutic instrumentalities will become available once we learn how to design [4–10] and construct [3, 11–13] complete artificial nanorobots composed of diamondoid nanometer-scale parts [12a] and onboard subsystems including sensors [1a], motors [12b], manipulators [12b, 12c], power plants [1c], and molecular computers [1d, 12d]. The presence of onboard computers is essential because *in vivo* medical nanorobots will be called upon to perform numerous complex behaviors which must be conditionally executed on at least a semiautonomous basis, guided by receipt of local sensor data and constrained by preprogrammed settings, activity scripts, and event clocking, and further limited by a variety of simultaneously executing real-time control protocols.

The development pathway for diamondoid medical nanorobots will be long and arduous. First, theoretical scaling studies [4–10] are used to assess basic concept feasibility. These initial studies would then be followed by more detailed computational simulations of specific nanorobot components and assemblies, and ultimately full systems simulations, all thoroughly integrated with additional simulations of massively parallel manufacturing processes from start to finish consistent with a design-for-assembly engineering philosophy. Once nanofactories implementing molecular manufacturing capabilities become available, experimental efforts may progress from component fabrication and testing to component assembly and finally to prototypes and mass manufacture of medical nanorobots, ultimately leading to clinical trials. By 2007 there was some limited experimental work with microscale-component microrobots [14–18] but progress on nanoscale-

component nanorobots remains largely at the concept feasibility stage. Since 1998, the author has published seven theoretical nanorobot scaling studies [4–10], several of which are briefly summarized below. Such studies are not intended to produce an actual engineering design for a future nanomedical product. Rather, the purpose is merely to examine a set of appropriate design constraints, scaling issues, and reference designs to assess whether or not the basic idea might be feasible, and to determine key limitations of such designs, including the many issues related to biocompatibility of medical nanorobots as extensively discussed elsewhere [2].

Complex medical nanorobots probably cannot be manufactured using the conventional techniques of self-assembly. As noted in the final report [19] of the recently completed congressionally-mandated review of the U.S. National Nanotechnology Initiative by the National Research Council (NRC) of the National Academies and the National Materials Advisory Board (NMAB): “For the manufacture of more sophisticated materials and devices, including complex objects produced in large quantities, it is unlikely that simple self-assembly processes will yield the desired results. The reason is that the probability of an error occurring at some point in the process will increase with the complexity of the system and the number of parts that must interoperate.”

The opposite of self-assembly processes is positionally controlled processes, in which the positions and trajectories of all components of intermediate and final product objects are controlled at every moment during fabrication and assembly. Positional processes should allow more complex products to be built with high quality and should enable rapid prototyping during product development. Positional assembly is the norm in conventional macroscale manufacturing (e.g., cars, appliances, houses) but has not yet been seriously investigated experimentally for nanoscale manufacturing. Of course, we already know that positional fabrication will work in the nanoscale realm. This is demonstrated in the biological world by ribosomes, which positionally assemble proteins in living cells by following a sequence of digitally encoded instructions (even though ribosomes themselves are self-assembled). Lacking this positional fabrication of proteins controlled by DNA-based software, large, complex, digitally specified organisms would probably not be possible and biology as we know it would cease to exist.

The most important inorganic materials for positional assembly may be the rigid covalent or “diamondoid” solids, since these could potentially be used to build the most reliable and complex nanoscale machinery. Preliminary theoretical studies have suggested great promise for these materials in molecular manufacturing. The NMAB/NRC Review Committee recommended [19] that experimental work aimed at establishing the technical feasibility (or lack thereof) of positional molecular manufacturing should be pursued and supported: “Experimentation leading to demonstrations supplying ground truth for abstract models is appropriate to better characterize the potential for use of bottom-up or molecular manufacturing systems that utilize processes more complex than self-assembly.” Making complex nanorobotic systems requires manufacturing techniques that can build a molecular structure by positional assembly [3]. This will involve picking

and placing molecular parts one by one, moving them along controlled trajectories much like the robot arms that manufacture cars on automobile assembly lines. The procedure is then repeated over and over with all the different parts until the final product, such as a medical nanorobot, is fully assembled.

The positional assembly of diamondoid structures, some almost atom by atom, using molecular feedstock has been examined theoretically [12, 20] via computational models of diamond mechanosynthesis (DMS). DMS is the controlled addition of individual carbon atoms, carbon dimers (C_2), or single methyl (CH_3) groups to the growth surface of a diamond crystal lattice in a vacuum manufacturing environment. Covalent chemical bonds are formed one by one as the result of positionally constrained mechanical forces applied at the tip of a scanning probe microscope apparatus. Programmed sequences of carbon dimer placement on growing diamond surfaces *in vacuo* appear feasible in theory [20–23]. Diamond mechanosynthesis is being sought [11, 21] but has not yet been achieved experimentally; in 1999 Ho and Lee [24] demonstrated the first site-repeatable site-specific covalent bonding operation of a two diatomic carbon-containing molecules (CO), one after the other, to the same atom of iron on a crystal surface. In 2003, Oyabu et al. [25] vertically manipulated single silicon atoms from the Si(111)–(7×7) surface, using a low-temperature near-contact atomic force microscope to demonstrate removal of a selected silicon atom from its equilibrium position without perturbing the (7×7) unit cell and also the deposition of a single Si atom on a created vacancy, both via purely mechanical processes. Efforts are currently underway to achieve DMS with carbon atoms experimentally [11].

To be practical, molecular manufacturing must also be able to assemble very large numbers of medical nanorobots very quickly. Approaches under consideration include using replicative manufacturing systems [13] or massively parallel fabrication, employing large arrays of scanning probe tips all building similar diamondoid product structures in unison [13]. For example, simple mechanical ciliary arrays consisting of 10,000 independent microactuators on a 1 cm^2 chip have been made at the Cornell National Nanofabrication Laboratory for micro-scale parts transport applications and similarly at IBM for mechanical data storage applications [26]. Active probe arrays of 10,000 independently actuated microscope tips have been developed by Mirkin's group at Northwestern University for dip-pen nanolithography [27] using DNA-based ink. Almost any desired 2D shape can be drawn using 10 tips in concert. Another microcantilever array manufactured by Protiveris Corp. has millions of interdigitated cantilevers on a single chip [28]. Martel's group has investigated using fleets of independently mobile wireless instrumented microrobot manipulators called NanoWalkers to collectively form a nanofactory system that might be used for positional manufacturing operations [29]. Zyvex Corp. (www.zyvex.com) of Richardson, Texas received a \$25 million, five-year, National Institute of Standards and Technology (NIST) contract to develop prototype microscale assemblers using microelectromechanical systems.

Ultimately, medical nanorobots will be manufactured in nanofactories efficiently designed for this purpose. One possible rough outline for a combined

experimental and theoretical program to explore the feasibility of nanoscale positional manufacturing techniques, starting with the positionally controlled mechanosynthesis of diamondoid structures using simple molecular feedstock and progressing to the ultimate goal of a desktop nanofactory appliance able to manufacture macroscale quantities of molecularly precise product objects according to digitally defined blueprints, is available at the Nanofactory Collaboration website: <http://www.MolecularAssembler.com/Nanofactory/Challenges.htm>.

The purpose of this chapter is to examine the wide range of computational tasks that might need to be performed by various representative classes of medical nanorobots. Only a sampling of such tasks will be presented because time and space do not permit an exhaustive survey. The discussion starts with descriptions of several classes of medical nanorobots for which preliminary scaling studies have already been published in the literature. Basic computational tasks are described in each case. Following a more general discussion of the various major functions common to many or all complex medical nanorobots, we introduce the concept of nanorobot control protocols which are required to ensure that each nanorobot fully completes its intended mission accurately, safely, and in a timely manner according to plan.

15.2. EXEMPLAR MEDICAL NANOROBOT DESIGNS

Preliminary scaling studies of several classes of medical nanorobots have been published in the literature and include respirocytes (artificial mechanical red cells) [4], microbivores (artificial white cells) [5], clottocytes (artificial platelets) [6], and chromallocytes (chromosome exchanging nanorobots) [7]. These studies are summarized here, with emphasis on the computational tasks and requirements for such devices.

15.2.1. Respirocytes

15.2.1.1. Nanorobot Description. The first theoretical design study of a medical nanorobot ever published in a peer-reviewed medical journal (in 1998) described an artificial mechanical red blood cell or “respirocyte” [4] made of 18 billion precisely arranged atoms (Fig. 15.1)—a bloodborne, spherical 1-micron diamondoid 1000-atmosphere pressure vessel [1e] with active pumping [1f] powered by endogenous serum glucose [1g], able to deliver 236 times more oxygen to the tissues per unit volume than natural red cells and to manage carbonic acidity, controlled by gas concentration sensors [1h] and an onboard nanocomputer [1d, 12d].

In the exemplar design, onboard pressure tanks can hold up to three billion oxygen (O_2) and carbon dioxide (CO_2) molecules. Molecular pumps are arranged on the surface to load and unload gases from the pressurized tanks. Tens of thousands of these individual pumps, called molecular sorting rotors [1f], cover a large fraction of the hull surface of the respirocyte. Molecules of oxygen (O_2) or




Figure 15.1. An artificial red cell: the respirocyte [4]. Designer Robert A. Freitas Jr., illustration by E-spaces. (© 2002 E-spaces and Robert A. Freitas Jr.)

carbon dioxide (CO_2) may drift into their respective binding sites on the exterior rotor surface and be carried into the respirocyte interior as the rotor turns in its casing. Once inside, a small pin is inserted into the binding site, forcibly ejecting the bound molecule into the interior tank volume. There are 12 identical pumping stations laid out around the equator of the respirocyte, with oxygen rotors on the left, carbon dioxide rotors on the right, and water rotors in the middle. Temperature [1i] and concentration [1j] sensors tell the devices when to release or pick up gases. Each station has special pressure sensors [1k] to receive ultrasonic acoustic messages [1m], so doctors can tell the devices to turn on or off, or change the operating parameters of the devices, while the nanorobots are inside a patient.

The basic operation of respirocytes is straightforward. These nanorobots, still entirely theoretical, would mimic the action of the natural hemoglobin-filled red blood cells. In the tissues, oxygen is pumped out of the device by the sorting rotors on one side. Carbon dioxide is pumped into the device by the sorting rotors on the other side, one molecule at a time. Half a minute later, when the respirocyte reaches the patient's lungs in the normal course of the circulation of the blood,

these same rotors reverse their direction of rotation, recharging the device with fresh oxygen and dumping the stored CO_2 , which diffuses into the lungs and can then be exhaled by the patient. The onboard nanocomputer and numerous chemical and pressure sensors enable complex device behaviors remotely reprogrammable by the physician via externally applied ultrasound acoustic signals. There is also a large internal void surrounding the nanocomputer which can be a vacuum, or can be filled or emptied with water. This allows the device to control its buoyancy very precisely (Section 2.1.2) and provides a crude but simple method for removing respiratory cells from the blood using a centrifuge, a procedure called nanapheresis [1n].

Various sensors are needed to acquire external data essential in regulating gas loading and unloading operations, tank volume management, and other special protocols. It is also convenient to include internal pressure sensors [1k] to monitor O_2 and CO_2 gas tank loading, ullage (container fullness) sensors [1o] for ballast and glucose fuel tanks, and internal/external temperature sensors [1i] to help monitor and regulate total system energy output. The attending physician can broadcast signals to molecular mechanical systems deployed inside the human body most conveniently using modulated compressive pressure pulses received by mechanical transducers embedded in the surface of the respiratory cell. Converting a pattern of pressure fluctuations into mechanical motions that can serve as input to a mechanical computer requires transducers that function as pressure-driven actuators [1p, 12e]. Data transmitted at ~ 10 MHz ($\sim 10^7$ bits/sec) using peak-to-trough 10-atm pressure pulses (the same as medical pulse-echo diagnostic ultrasound systems [30]) should attenuate only $\sim 10\%$ per 1 cm of travel [1q], so whole-body broadcasts should be feasible even in emergency field situations. Internal communications within the respiratory cell itself can be achieved by impressing modulated low-pressure acoustical spikes [1r] on the hydraulic working fluid of the power distribution system [1s], or via simple mechanical rods and couplings [1t].

Onboard power is provided by a mechanochemical engine [1u] or fuel cell [1v] that exoergically combines ambient glucose and oxygen to generate mechanical or electrical energy to drive molecular sorting rotors and other subsystems, as demonstrated in principle in a variety of biological motor systems [1w]. Sorting rotors absorb glucose directly from the blood and store it in an internal fuel tank. Oxygen is tapped from onboard storage. The power system is scaled such that each glucose engine can fill the primary O_2 tank from a fully empty condition in 10 seconds, requiring a peak continuous output of 0.3 pW. The glucose fuel tank is scaled such that one tankful of fuel drives the glucose engine at maximum output for 10 seconds, consuming 5% of the O_2 gas stored onboard and releasing a volume of waste water approximately equal to the volume of the glucose consumed. Power is transmitted mechanically or hydraulically using an appropriate working fluid, and can be distributed as required using rods and gear trains [1x] or using pipes and mechanically operated valves, controlled by the nanocomputer.

A 5 cc therapeutic dose of 50% respiratory cell saline suspension containing 5 trillion nanorobots would exactly replace the gas carrying capacity of the patient's entire 5.4 liters of blood. If up to 1 liter of respiratory cell suspension can

safely be added to the human bloodstream [2a], this could keep a patient's tissues safely oxygenated for up to 4 hours, even if a heart attack caused the heart to stop beating or if there was a complete absence of respiration or no external availability of oxygen. Primary medical applications of respirocytes would include emergency revival of victims of carbon monoxide poisoning at the scene of a fire; transfusable blood substitution; partial treatment for anemia and various lung and perinatal/neonatal disorders; enhancement of cardiovascular/neurovascular procedures, tumor therapies, and diagnostics; prevention of asphyxia; artificial breathing (e.g., underwater, high altitude, etc.); sports applications (e.g., ability to sprint at Olympic speed for up to 15 minutes without breathing); and a variety of other applications in veterinary medicine, military science, and space exploration.

15.2.1.2. Computational Tasks. An onboard nanocomputer is necessary to provide precise control of all basic respirocyte operations. These operations include respiratory gas loading and unloading, sensor field management and data handling, molecular rotor field and storage/ballast tank control, glucose engine throttling, management of glucose fuel supply, powerplant effluent regulation, control of power distribution, and interpretation of sensor data and commands received from the outside. The nanocomputer is also responsible for basic self-diagnosis monitoring and activation of failsafe shutdown protocols, and ongoing real-time revision or correction of protocols for *in vivo* devices.

With efficient programming, a 10^4 bit/sec computer can probably meet all computational requirements, given the simplicity of analogous chemical process control systems in factory settings [31, 32]. That's roughly the computing capacity of a transistor-based 1960-vintage IBM 1620 computer, or about 1/50th the capacity of a 1976-vintage Apple II microprocessor-based PC. Both the IBM 1620 and the Apple II used $\sim 10^5$ bits of internal memory, but even the early PCs typically had access to 10^6 bits (0.1 megabyte) of external floppy drive memory. Assuming ~ 500 bits/sec/nm³ and 10^{18} bits/sec/watt for nanomechanical computers [12d], and ~ 5 bits/nm³ for nanomechanical mass storage systems [1y, 12f], each 10^4 bit/sec CPU is allocated a volume of $\sim 10^4$ nm³ and consumes ~ 0.01 pW (3% of the power output of one glucose engine), while 500 kilobits of memory requires $\sim 10^5$ nm³. The use of reversible logic could significantly reduce power consumption [33–35].

Respirocyte behavior is initially governed by a set of default protocols which can be modified at any time by the attending physician. Basic protocols will exist for operating molecular sorting rotors at various speeds and directions in response to sensor data. For example, ballast water pumping will normally be driven by internal ullage and temperature sensors. O₂ rotors may load tanks at $P_{O_2} > 95$ mmHg (alveolar) and unload at $P_{O_2} < 40$ mmHg (tissues). CO₂ rotors may fill at $P_{CO_2} > 46$ mmHg and unfill at $P_{CO_2} < 40$ mmHg, and may incorporate other sensor data including P_{O_2} , temperature, etc., to fine-tune pressure thresholds and enhance reliability. Gas loading parameters may be precisely specified in an individualized onboard lookup table provided by the physician for his patient, as for instance to adjust for declining arteriovenous oxygen gradient at high altitudes [36].

Respirocytes, like natural hemoglobin, may also participate in the elimination of CO and in NO-mediated vascular control [37] if appropriate sorting rotors and onboard tankage are provided.

One important protocol described in the original paper [4] is the initial warmup procedure for empty respirocytes that must be infused into a patient. Respirocytes comprising a single augmentation dosage are stored as a dry powder, tanks empty, in sealed plastic drip bags with two hose couplings. With no batteries to run down, consumables to age, vapors to outgas or organic matter to decompose, the product should have a long shelf life. To use the product, the bag is filled with ice-cold 0.13 M glucose (23 gm/liter) solution, plus any salts, minerals, vitamins, proteins or other substances the physician deems appropriate. The powder floats on the liquid surface. An external source of O₂ gas (and CO₂, if required) is provided through the second coupling. Respirocyte sensors detect the presence of glucose and begin pumping fuel into the glucose tanks. As these tanks fill, each device loads its oxygen tank to rated capacity. The powder still floats on the surface. Finally, the respirocytes load ballast tanks and sink to the bottom of the bag. (Powder remaining on the surface evidences malfunction.) During this ~30 sec charging process, the respirocytes absorb 53.7 liters of O₂ @ STP, store 0.78 liters of CO₂ waste, release enough energy to warm bag contents to 42°C, and leave behind a 0.005 M glucose solution, closely matching normal blood concentration and temperature. Upon receiving the correct activation command, broadcast acoustically using an ultrasound transmitter device pressed against the bag, the respirocytes blow sufficient ballast water to achieve neutral buoyancy, creating a perfect suspension (after agitation) ready for IV drip. The suspension is ~300 times plasma viscosity (~castor oil or canola oil at 37°C), still permitting ready plug flow [38].

Once a therapeutic purpose is completed, it may be desirable to extract artificial devices from circulation, requiring activation of a respirocyte filtration protocol. During this protocol, called nanapheresis [1n], blood to be cleared may be passed from the patient to a specialized centrifugation apparatus where ultrasound acoustic transmitters command respirocytes to establish neutral buoyancy. No other solid blood component can maintain exact neutral buoyancy, hence those other components precipitate outward during gentle centrifugation and are drawn off and added back to filtered plasma on the other side of the apparatus. Meanwhile, after a period of centrifugation, the plasma, containing mostly suspended respirocytes but few other solids, is drawn off through a 1-micron filter, removing the respirocytes. Filtered plasma is recombined with centrifuged solid components and returned undamaged to the patient's body. The rate of separation is further enhanced either by commanding respirocytes to empty all tanks, lowering net density to 66% of blood plasma density, or by commanding respirocytes to blow a 5-micron O₂ gas bubble to which the device may adhere via surface tension, allowing it to rise at 45 mm/hour in a normal gravitational environment.

Respirocytes can be programmed with even more sophisticated behaviors. Detection of P_{CO2} < 0.5 mmHg and P_{O2} > 150 mmHg, indicating direct exposure to

atmosphere and a high probability that the device has been bled out of the body, should trigger a prompt gas venting and failsafe device shutdown procedure. This is particularly important in the case of oral bleeding, since respirocytes could explode if crushed between two hard planar surfaces with no avenue of escape as might occur during dental grinding [1z]—tooth enamel is the hardest natural substance in the human body, and a patient with an oral lesion could spread respirocyte-impregnated blood over the teeth. Of course, single-device explosions of a micron-size pressure tank caused by dental grinding would probably be undetectable by the patient.

Self-test algorithms monitoring tank filling rates, unaccounted pressure drops (indicating a leak), clutch responses, etc. may detect significant device malfunction, causing the respirocyte to place itself in standby mode ready to respond to an acoustic command to execute the filtration protocol for nanapheresis. Basic self-diagnostic routines should be able to detect simple failure modes such as jammed rotor banks, plugged flues, gas leaks, and so forth, and to use backup systems to place the device into a fail-safe dormant mode pending removal by filtration. Dormant mode can also be triggered by thermal shutoff protocols which activate upon self-detection of any of at least four failure scenarios: (1) A normally functioning respirocyte is deposited in a relatively dry (e.g., osteal or cartilaginous) location, losing contact with the aqueous heat sink; (2) one or more glucose engines become jammed at full open throttle; (3) inbound sorting rotors overload pressure vessel; or (4) device fragmentation or combustion.

Using onboard $(21\text{ nm})^3$ pressure transducers, respirocytes can detect, record, or respond to changes in heart rate or blood pressure, since arterial pressure is normally 0.1–0.2 atm ($\sim 100\text{ zJ}/(21\text{ nm})^3$ sensor) and systolic/diastolic differential is 0.05–0.07 atm ($40\text{--}60\text{ zJ}/(21\text{ nm})^3$ sensor), both well above the mean thermal noise limit of $k_B T \sim 4\text{ zJ}$ at human body temperature. Outmessaging [1aa] protocols could allow the population of respirocytes to communicate systemwide status (e.g., “low oxygen,” “low glucose,” “under immune attack,” “cyanide detected” [39]) directly with the patient by inducing recognizable physiological cues (fever, shivering, gasping) or by transmitting the information to implanted GUIs accessible to the patient [1bq], or with the physician by generating subtle respiratory patterns requiring diagnostic equipment to detect, either automatically or in response to an acoustically transmitted global inquiry initiated by patient or physician. Attempted phagocytosis of bloodborne respirocytes by the RES [2b], particularly Kupffer cells in the liver [2c], bloodborne white cells, macrophages, or other natural phagocytic cells could, upon detection by the respirocyte, activate appropriate phagocyte avoidance and escape protocols (Section 15.4.2) using techniques detailed elsewhere [2d].

Other useful control protocols may enable *in vivo* respirocytes to be commanded to cease or resume operating locally/globally, to enable/disable one or another class of sorting rotors, to alter sensor sensitivities, to blow tanks or to run all engines continuously, and so forth. Note that the long-term bloodstream presence of a therapeutic dose of working respirocytes would preclude tissue hypoxia, possibly reducing EPO secretion as low as 1% of normal levels [40], thus

suppressing erythropoiesis. It may be possible to avoid the resulting decimation of the natural erythrocyte population by adjusting respirocyte P_{O_2}/P_{CO_2} response thresholds so that these devices activate only when red cells are stressed. Command messages of these and similar types should require lengthy authorization codes (Section 15.4.5) to prevent accidental or malicious triggering of respirocyte behaviors having potentially harmful consequences.

15.2.2. Microbivores

15.2.2.1. Nanorobot Description. A second theoretical design study of a medical nanorobot describes an artificial mechanical white cell or “microbivore”—an oval-shaped device (Fig. 15.2) measuring a few microns in size and made of diamond and sapphire—that would seek out and digest unwanted bloodborne pathogens [5]. One main task of natural white cells is to absorb and consume microbial invaders in the bloodstream, called phagocytosis. Microbivore nanorobots would also perform phagocytosis, but would operate much faster, more reliably, and under human control.

The microbivore is an oblate spheroidal nanomedical device measuring 3.4 microns in diameter along its major axis and 2.0 microns in diameter along




Figure 15.2. An artificial white cell: the microbivore [5]. Designer Robert A. Freitas Jr., illustration and additional design Forrest Bishop. (© 2001 Zyvex Corp. and Robert Freitas.)

its minor axis, consisting of 610 billion precisely arranged structural atoms in a gross geometric volume of 12.1 micron^3 and a dry mass of 12.2 picograms. This size helps to ensure that the nanorobot can safely pass through even the narrowest of human capillaries and other tight spots in the spleen (e.g., the interendothelial splenofenestral slits [2e]) and elsewhere in the human body [2f]. The microbivore has a mouth with an iris-like door, called the ingestion port, where microbes are fed in to be digested. The microbivore also has a rear end, or exhaust port, where the completely digested remains of the pathogen are harmlessly expelled from the device. The rear door opens between the main body of the microbivore and a tail-cone structure. The device may consume up to 200 pW of continuous power while completely digesting trapped microbes at a maximum throughput of 2 micron^3 of organic material per 30-second cycle, which is large enough to internalize a single microbe from virtually any major bacteremic species in a single gulp. This “digest and discharge” protocol [1ab] is conceptually similar to the internalization and digestion process practiced by natural phagocytes, except that the artificial process should be much faster and cleaner. For example, it is well-known that macrophages release biologically active compounds during bacteriophagy [41], whereas well designed microbivores need only release biologically inactive effluent.

The first task for the bloodborne microbivore is to reliably acquire a pathogen to be digested. If the correct bacterium bumps into the nanorobot surface, reversible species-specific binding sites on the microbivore hull can recognize and weakly bind to the bacterium. A set of 9 distinct antigenic markers should be specific enough [1ac], since all 9 must register a positive binding event to confirm that a targeted microbe has been caught. There are 20,000 copies of these 9-marker receptor sets, distributed in 275 disk-shaped regions across the microbivore surface. Inside each receptor ring are more rotors to absorb ambient glucose and oxygen from the bloodstream to provide nanorobot power. At the center of each 150-nm diameter receptor disk is a grapple silo. Once a bacterium has been captured by the reversible receptors, telescoping robotic grapples [1ad] rise up out of the microbivore surface and attach to the trapped bacterium, establishing secure anchorage to the microbe’s plasma membrane. The microbivore grapple arms are about 100 nm long and have various rotating and telescoping joints that allow them to change their position, angle, and length. After rising out of its silo, a grapple arm can execute complex twisting motions, and adjacent grapple arms can physically reach each other, allowing them to hand off bound objects as small as a virus particle. Grapple handoff motions can transport a large rod-shaped bacterium from its original capture site forward to the ingestion port at the front of the device. The captive organism is rotated into the proper orientation as it approaches the open microbivore mouth, where the pathogen cell is internalized into a 2 micron^3 morcellation chamber.

There are two concentric cylinders inside the microbivore. The bacterium is minced into nanoscale pieces in the morcellation chamber [1ae], the smaller inner cylinder, then the remains are pistoned into a separate 2 micron^3 digestion chamber, a larger outer cylinder. In a preprogrammed sequence, ~ 40 different engineered digestive enzymes are successively injected and extracted six times

during a single digestion cycle, progressively reducing the morcellate to mono-residue amino acids, mononucleotides, glycerol, free fatty acids and simple sugars, using an appropriate array of molecular sorting rotors. These basic molecules are then harmlessly discharged back into the bloodstream through the exhaust port at the rear of the device, completing the 30-second digestion cycle. When treatment is finished, the doctor may transmit an ultrasound signal to tell the circulating microbivores that their work is done. The nanorobots then exit the body through the kidneys and are excreted with the urine in due course.

The microbivore needs a variety of external and internal sensors to complete its tasks. External sensors include chemical sensors for glucose, oxygen, carbon dioxide, and so forth, up to 10 different molecular species with 100 sensors per molecular species. Pressure sensors for acoustic communication are mounted within the nanorobot hull to permit the microbivore to receive external instructions from the attending physician during the course of *in vivo* activities. Ten (redundant) internal temperature sensors capable of detecting 0.3°C temperature change [1i] are positioned near each of the 10 independent powerplants.

Microbivores use bloodstream glucose for power like respiocytes, but also use blood-dissolved oxygen which is available only at much lower concentrations than in the lung capillary bed, so many more molecular sorting rotors for O₂ must be present on the microbivore hull. The microbivore is scaled for a maximum power output of 200 pW, about 1000 times higher than for the respiocyte. Diamondoid mechanical cables may transmit internal mechanical energy at power densities of $\sim 6 \times 10^{12} \text{ W/m}^3$ [1x]. To connect every powerplant with each of its 9 neighbors via power cables, permitting rapid load sharing among any pair of powerplants inside the device, requires 45 power cables; assuming 1000 internal power cables to accommodate additional power distribution tasks and for redundancy, total power cable volume is 0.05 micron³. By varying the cable rotation rate, the same power cables can simultaneously be used to convey necessary internal operational information [1t] including sensor data traffic and control signals from the computers.

A human neutrophil, the most common type of leukocyte or white cell, can capture and engulf a microbe in a minute or less, but complete digestion and excretion of the organism's remains can take an hour or longer. Thus our natural white cells—even when aided by antibiotics—can sometimes take weeks or months to completely clear bacteria from the bloodstream. By comparison, a single terabot (10¹²-nanorobot) dose of microbivores should be able to fully eliminate bloodborne pathogens in just minutes or hours, even in the case of locally dense infections. This is accomplished without increasing the risk of sepsis or septic shock because all bacterial components (including all cell-wall LPS) are internalized and fully digested into harmless nonantigenic molecules prior to discharge from the device. No matter that a bacterium has acquired multiple drug resistance to antibiotics or to any other traditional treatment—the microbivore will eat it anyway. Microbivores would be up to ~ 1000 times faster-acting than antibiotic-aided natural leukocytes. The nanorobots would digest ~ 100 times more microbial material than an equal volume of natural white cells could digest

in any given time period, and would have far greater maximum lifetime capacity for phagocytosis than natural white blood cells. Besides intravenous bacterial, viral, fungal, and parasitic scavenging, microbivores or related devices could also be used to help clear respiratory or cerebrospinal bacterial infections, or infections in other nonsanguinous fluid spaces such as pleural [42], synovial [43], or urinary fluids; eliminate bacterial toxemias and biofilms [44]; eradicate viral, fungal, and parasitic infections; disinfect surfaces, foodstuffs, or organic samples; and help clean up biohazards and toxic chemicals. Related nanorobots with enhanced tissue mobility could be programmed to quickly recognize and digest even the tiniest aggregates of early tumor cells with unmatched speed and surgical precision, eliminating cancer. Similar nanorobots of this class could be programmed to remove circulatory obstructions in just minutes, quickly rescuing even the most compromised stroke victim from near-certain ischemic brain damage, and could have other uses in various veterinary and military applications.

15.2.2.2. Computational Tasks. An onboard nanocomputer is required to provide precise control of all basic microbivore operations. These operations include most of the basic computational tasks already described for the respirocyte (Section 2.1.2) such as onboard tank volume control, sensor coordination, data and power management, and nanapheresis protocols. Other tasks unique to the microbivore include management of multiple reversible microbial binding sites to ensure accurate identification of the targeted microbial species, control of hundreds of independent grapple elevator and segment rotation mechanisms, coordination of the grapple motion field to transport trapped microbes in a controlled manner across the nanorobot surface or for nanorobot locomotion, and sequence control for morcellation and digestion activities (ingestion port opening/closing, mincing, transfer pistoning, enzyme cycling, effluent ejection pistoning, and exhaust port opening/closing).

The microbivore computer is scaled as a 0.01 micron^3 mechanical nanocomputer [1y, 12d], in principle capable of >100 megaflops but normally operated at ~ 1 megaflop to hold power consumption to $\sim 60 \text{ pW}$. Assuming $\sim 5 \text{ bits/nm}^3$ for nanomechanical data storage systems [12d] and a read/write cost of $\sim 10 \text{ zJ/bit}$ at a read/write speed of $\sim 10^9 \text{ bits/sec}$ [1af, 12d], then a comparison with other software systems of comparable complexity (Table 15.1) suggests that the ~ 5 megabits of mass nanomechanical memory needed to hold the microbivore control system displaces a volume of 0.001 micron^3 and draws $\sim 10 \text{ pW}$ while in continuous operation. The baseline microbivore design includes 10 duplicate computer/memory systems for redundancy (with only one of the ten computer/memory systems in active operation at a time), displacing a total of 0.11 micron^3 and consuming $\leq 70 \text{ pW}$.

As with respirocytes, microbivore behavior is initially governed by a set of default protocols many of which can be modified at any time by the attending physician. Basic protocols have already been described for the respirocyte (Section 2.1.2). Since virtually every medical nanorobot placed inside the human body will encounter natural phagocytic cells many times during its mission [2d],

TABLE 15.1. Lines of Compactly Written Low-Error Software Code Required to Control Complex Semiautonomous Machines.

Control software for device	Estimated lines of code	Bits of code (~ 100 bits/line)	Ref.
Voyager Spacecraft Software	3,000	300,000	[45]
Viking Lander Software	—	432,000	[46]
Respirocyte Control System (est.)	—	~ 500,000	[4]
Galileo Spacecraft Software	8,000	800,000	[45]
Cassini Spacecraft Software	32,000	3,200,000	[45]
Microbivore Control System (est.)	—	~ 5,000,000	[5]
Ariane Flight Control Software	90,000	9,000,000	
Airbus 340 Flight Warning System	100,000	10,000,000	[47]
Mars Pathfinder Spacecraft	160,000	16,000,000	[48]
Space Shuttle Software	500,000	50,000,000	[45]
Boeing 777 and Airbus 340	3,000,000	300,000,000	[49]
			[50]

microbivores, like respirocytes, may incorporate any of several possible phagocyte avoidance and escape techniques [2d], possibly including, for example, surface-tethered phagocyte chemorepellent molecules [2g, 51] or phagocyte engulfment inhibitors [2h, 52], as well as more proactive approaches to phagocytosis avoidance [2d].

Many special protocols will be needed by microbivores but a complete enumeration is beyond the scope of this chapter. Biocompatibility-related protocols (Section 15.4.2) provide examples ranging from the simple to the more complex. At the simple end of the range is the microbial tail protocol [5]. Free releases of bacterial flagella into the bloodstream could produce inflammation or elicit various immune system responses and thus should be avoided. Complete internalization of tail may be ensured by specialized operational routines (e.g., forced end-over-end rotation of an internalized microbe while inside the morcellation chamber, thus completely spooling the tail into the microbivore before fully sealing the ingestion port door) or by specialized mechanical tools or jigs (e.g., a counterrotating interdigitated-knobbed capstan-roller pair).

An example of a more complex protocol arises from the fact that all bloodborne nanorobots larger than ~ 1 micron in all three physical dimensions are subject to possible geometrical trapping in the fenestral slits of the splenic sinusoids in the red pulp of the spleen [2e]. A small percentage of blood is forced to circulate through a physical filter in the spleen requiring passage through slits measuring 1–2 microns in width and ~ 6 microns in length [2e]. Microbivores which become pinned to a slit face-on, or which become stuck edge-on during an attempted passage, can detect that they have become trapped by measuring various blood component concentration and pressure differentials across their surfaces. The nanorobot then activates its automatic splenofenestral escape protocol [5], which involves the extension and patterned ciliation of surface

grapples until sensor readings reveal that passage through the slit has been achieved, after which the grapples are retracted.

15.2.3. Clottocytes

15.2.3.1. Nanorobot Description. Another theoretical design study describes an artificial mechanical platelet [6] or “clottocyte” that would allow complete hemostasis in as little as ~ 1 second, even in moderately large wounds. This response time is on the order of 100–1000 times faster than the natural hemostatic system. The baseline clottocyte is conceived as a serum oxyglucose-powered spherical nanorobot ~ 2 microns in diameter (~ 4 micron³ volume) containing a fiber mesh that is compactly folded onboard. Upon command from its control computer, the device promptly unfurls its mesh packet in the immediate vicinity of an injured blood vessel—following, say, a cut through the skin. Soluble thin films coating certain parts of the mesh dissolve upon contact with plasma water, revealing sticky sections (e.g., complementary to blood group antigens unique to red cell surfaces [1ag]) in desired patterns. Blood cells are immediately trapped in the overlapping artificial nettings released by multiple neighboring activated clottocytes, and bleeding halts at once.

The required blood concentration n_{bot} of clottocyte nanorobots required to stop capillary flow at velocity $v_{\text{cap}} \sim 1$ mm/sec [1ah] in a response time $t_{\text{stop}} = 1$ sec, assuming $n_{\text{overlap}} = 2$ fully overlapped nets each of area $A_{\text{net}} = 0.1$ mm², is $n_{\text{bot}} \sim n_{\text{overlap}} / (A_{\text{net}} t_{\text{stop}} v_{\text{cap}}) = 20$ mm⁻³, or just ~ 110 million clottocytes in the entire 5.4-liter human body blood volume possessing ~ 11 m² of total deployable mesh surface. This total dose is ~ 0.4 mm³ of clottocytes, which produces a negligible serum nanocrit [1ai] or “Nct” of $\sim 0.00001\%$ (nanorobot/blood volume ratio). During the 1 second hemostasis time, an incision wound measuring 1 cm long and 3 mm deep would lose only ~ 6 mm³ of blood, less than one-tenth of a single droplet. There are 2–3 red cells per deployed 1 micron² mesh square, more than enough to ensure that the meshwork will be completely filled, allowing complete blockage of a breach.

Total natural bleeding time, as experimentally measured from initial time of injury to cessation of blood flow, may range from 2–5 minutes [53] up to 9–10 minutes [54, 55] if even small doses of anticoagulant aspirin are present [56], with 2–8 minutes being typical in clinical practice. With a ~ 1 sec response time, artificial mechanical platelets appear to permit the halting of bleeding 100–1000 times faster than natural hemostasis. While 1–300 platelets might be broken and still be insufficient to initiate a self-perpetuating clotting cascade, even a single clottocyte, upon reliably detecting a blood vessel break, can rapidly communicate this fact to its neighboring devices, immediately triggering a progressive controlled mesh-release cascade. Clottocytes may perform a clotting function that is equivalent in its essentials to that performed by biological platelets—but at only $\sim 0.01\%$ of the bloodstream concentration of those cells. Hence clottocytes appear to be $\sim 10,000$ times more effective as clotting agents than an equal volume of natural platelets.

15.2.3.2. Computational Tasks. Besides the nanocomputers and control systems similar to those previously described for other nanorobots, clottocytes crucially require special control protocols to ensure that these nanorobots cannot release their mesh packets in the wrong places inside the body, or at an inappropriate time. These protocols will demand that carefully specified constellations of sensor readings must be observed before device activation is permitted. Reliable communications protocols will be required to control coordinated mesh releases from multiple neighboring devices and to regulate the maximum multidevice-activation radius within the local clottocyte population.

Detection of the “bled out of body” condition will be an especially important component of these protocols. Atmospheric concentrations of gases such as carbon dioxide and oxygen are different than the concentrations of those same gases in blood serum. As clottocyte-rich blood enters a breach in a blood vessel, nanorobot onboard sensors can rapidly detect the change in partial pressures, often indicating that the nanodevice is being bled out of the body. At a nanorobot whole-blood concentration of 20 mm^{-3} , mean device separation is 370 microns. If the first device to be bled from the body lies 75 microns from the air–serum interface, oxygen molecules from the air can diffuse through serum at human body temperature (310 K) from the interface to the nanodevice surface in ~ 1 second [1aj]. Detection of this change can be rapidly broadcast to neighboring clottocytes using $\sim 1500 \text{ m/sec}$ waterborne acoustic pulses that are received by devices up to several millimeters away in times on the order of microseconds, allowing rapid propagation of a carefully controlled device-enablement cascade. Similarly, air temperature is normally lower than body temperature. The thermal equilibration time [1ak] across a distance L in serum at 310 K is $t_{\text{EQ}} \sim (6.7 \times 10^6) L^2$, hence a device that lies 75 microns from the air–plasma interface can detect a change in temperature in $t_{\text{EQ}} \sim 40$ millise. Other relevant sensor readings may include blood pressure profiles, bioacoustic monitoring, bioelectrical field measurements, optical and ultraviolet radiation detection, and sudden shifts in pH or other ionic concentrations. At some cost in rapidity of response, clottocytes also could eavesdrop [1am] on natural biological platelet control signals, using sensors with receptors for the natural prostaglandins produced by endothelial cells that normally induce or inhibit platelet activation, and then take appropriate action upon receipt of those natural biochemical signals.

Biocompatibility [2] requirements engender additional needs for specialized clottocyte control protocols. For instance, the rapid mechanical action of clottocytes could interfere with the much slower natural platelet adhesion and aggregation processes, or disturb the normal equilibrium between the clotting and fibrinolytic systems [2i]. Thus it may be necessary for artificial platelets to release quantities of various chemical substances that will encourage the remainder of the coagulation cascade to proceed normally or at an accelerated pace, including timed localized vasodilation and vasoconstriction, control of endothelial cell modulation of natural platelet action, and finally clot retraction and fibrinolysis much later during tertiary hemostasis.

Yet another biocompatibility-related complication is that the bare tissue walls of a wound will continue to exude fluid [2j], and may begin to desiccate, if only the capillary termini are sealed but the rest of the tissue is left exposed to open air. Since clottocytes may remain attached to their discharged nets, and can communicate with each other via acoustic channels [1m], it should be possible to precisely control the development of a larger artificial mesh-based clot via coordinated mesh extensions or retractions within the clot. Alternatively, clottocytes could allow blood fluids to flood small incised or avulsed wound volumes, allowing exposed tissue walls to be bathed in fluids but casting a watertight sealant net across the wound opening flush with the epidermal plane of the wound cavity. The operational protocols that trigger and direct these behaviors will need to be rather complex to ensure safe and fully biocompatible operation.

Additional protocols are required to allow clottocytes to prevent accidental but potentially fatal catastrophic natural clotting cascades such as disseminated intravascular coagulation or DIC [57]. One solution is to equip clottocytes with sensors to detect decreased serum levels of fibrinogen, plasminogen, alpha₂-antiplasmin, antithrombin III, factor VII, and protein C, and elevated levels of thrombin and various fibrin/fibrinogen-derived degradation products [2i]. If DIC conditions arise, nanorobots might respond by absorbing and metabolizing the excess thrombin (which trigger clotting), or by releasing thrombin inhibitors such as antithrombin III, hirudin, argatroban, or lepirudin [58] or anticoagulants that reduce thrombin generation such as danaparoid [58] to interrupt the cascade [58, 59]. For example, a $\sim 0.02\%$ Nct concentration of nanorobots, suitably activated according to physician-approved parameters, could replace the entire depleted natural bloodstream content of antithrombin III from onboard stores in seconds.

Clottocytes will require still more sophisticated operational protocols if they are intended to assist platelets participating in the sealing of internal blood vessel lesions, in order to avoid inadvertently blocking the lumen of the entire vessel, e.g., in the case of minor internal bleeding. Similarly, prevention of bleeding at vascular anastomoses, hemarthroses, internal bruising, “blood blisters,” and larger tissue hematomas, as well as forced local coagulation in tumors or in intracerebral aneurysms, may also require more advanced protocols, possibly including integration with preexisting *in vivo* navigation systems [1an]. For some of these applications, modified clottocytes possessing the mobility of microbivores may be required in tandem with a graduated recruitment response depending upon how many (intercommunicating) devices appear to be involved in the event.

15.2.4. Chromalloyocytes

15.2.4.1. Nanorobot Description. The chromalloyocyte [7] is a hypothetical mobile cell-repair nanorobot whose primary purpose is to perform chromosome replacement therapy (CRT). In CRT, the entire chromatin content of the nucleus in a living cell is extracted and promptly replaced with a new set of prefabricated chromosomes which have been artificially manufactured as defect-free copies of




Figure 15.3. Artist's conceptions of the basic chromalloyte [7] design: devices walking along luminal wall of blood vessel (left); telescoping funnel assembly and proboscis manipulator in extended position (right). (Image © 2006 Stimulacra LLC (www.stimulacra.net) and Robert A. Freitas Jr. (www.rfreitas.com))

the originals. The chromalloyte (Fig. 15.3) will be capable of limited vascular surface travel into the capillary bed of the targeted tissue or organ, followed by diapedesis (exiting a blood vessel into the tissues) [1ao], histonotation (locomotion through tissues) [1ap], cytopenetrantion (entry into the cell interior) [1aq], and complete chromatin replacement in the nucleus of the target cell. The CRT mission ends with a return to the bloodstream and subsequent extraction of the device from the body at the original infusion site. Replacement chromosomes are manufactured in a desktop *ex vivo* chromosome sequencing and manufacturing facility, then loaded into the nanorobots for delivery to specific targeted cells during CRT. A single lozenge-shaped 69 micron³ chromalloyte measures 4.18 microns and 3.28 microns along cross-sectional diameters and 5.05 microns in length, typically consuming 50–200 pW in normal operation and a maximum of 1000 pW in bursts during outmessaging, the most energy-intensive task. Treatment of an entire large human organ such as a liver, involving simultaneous CRT on all 250 billion hepatic tissue cells, might require the localized infusion of a ~ 1 terabot (10^{12} devices) ~ 69 cm³ chromalloyte dose in a 1-liter 7% saline suspension during a ~ 7 hour course of therapy.

The chromalloyte includes an extensible primary manipulator 4 microns long and 0.55 microns in diameter called the Proboscis that is used to spool up chromatin strands via slow rotation when inserted into the cell nucleus. After spooling, a segmented funnel assembly is extended around the spooled bolus of DNA, fully enclosing and sequestering the old genetic material. The new chromatin is then discharged into the nucleus through the center of the Proboscis by pistoning from internal storage vaults, while the old chromatin that is sequestered inside the sealed watertight funnel assembly is forced into the storage vaults as space is vacated by the new chromatin that is simultaneously being pumped out. The chromalloyte includes a mobility system similar to the microbivore grapple system, along with a solvation wave drive [1ar] that is used to ensure smooth passage through cell plasma and nuclear membranes.

15.2.4.2. Computational Tasks. Besides the usual sensor, mobility, data and power management tasks common to other nanorobots, chromalloy cells must also control Proboscis extension and variable rotation rates, adjustment of adhesioregulatory surfaces [2k] on the Proboscis, pumping and fluid gating through the Proboscis, funnel assembly extension and retraction motions, and valving of materials into and out of the storage vaults. Locomotion also requires protocols for diapedesis through capillary walls, ECM (extracellular matrix) brachiation [1as] or grapple-mediated ciliary swimming [1at] through acellular tissues spaces (including emergency use of the Proboscis), transit through plasma and nuclear membranes [1aq], and nanapheresis [1n] or other nanorobot extraction procedures. (Chromalloy cells must also be capable of emergency auto-excretion through the renal tubules as long as external acoustic power is still being supplied to the patient, or they must have a capture protocol in which they can allow themselves to be harmlessly phagocytosed and transported to lymph nodes [2m] for subsequent removal by other nanorobots in a separate procedure.) The position of the apical terminus of the Proboscis must be controlled during new chromatin discharge to allow placement of new chromosomes near their optimum territorial regions inside the nucleus, requiring integration of Proboscis position information with sensor-derived information about the nanorobot's position relative to the cell nucleus. Precise control, timing, and coordination with sensor data is required to manage the release and subsequent recovery, via molecular sorting rotors, of dozens of small-molecule engineered biological reagents that are deployed in waves by the nanorobot to temporarily suppress various natural processes such as mechanotransduction, apoptosis, and inflammation. The computer must also guide nanorobot navigation [1an] to and from the target cell by coordinating locomotion with real-time positional information possibly received from a navigational microtransponder network [1au] that has been preinstalled in the patient's body. (The navigational system itself, including its installation and management, are a separate nanorobotic instrumentality that is operated independently of the chromalloy cells.)

Chromalloy cell onboard computation and control is provided by a computer system similar to that employed in the microbivore [5]. This includes a tenfold redundant 0.01 micron^3 CPU throttled back to a ~ 1 megaflop processing rate to conserve energy, giving a total computer volume of 0.1 micron^3 . However, the chromalloy cell incorporates a tenfold redundant mass memory system that is ten times larger (50 megabits, 0.01 micron^3) than for the microbivore (5 megabits, 0.001 micron^3), giving a total data storage volume of 0.1 micron^3 . This increased memory allocation is justified by (1) the increased complexity of a CRT mission as compared to an antimicrobial mission, and (2) the need for greater reliability, safety, and certainty of result in the case of CRT, where a mission failure could have more serious medical consequences. There is sufficient unallocated volume in the current chromalloy cell design to permit significantly increased onboard data storage if required.

The overall CRT mission requires completing a five-phase procedure. One of these phases, called "chromosome replacement," involves at least

semiautonomous nanorobot activities and includes a specific 26-step sequential process for performing whole-genome chromosome exchange nanosurgery on a living cell using the chromalloyocytes [7]. Each step must be verified as complete before the next step can be initiated. Failure to reliably and provably complete any given step should trigger a safe mission abort protocol in which the nanorobot must appropriately reverse any steps already completed, thus restoring the cell to its original condition, then withdraw safely from the cell and subsequently from the patient's body, while reporting all details of this failure to the attending physician. Autonomous control of a 26-step chromosome exchange operation will include numerous sensor-driven checkpoints. Motions and speeds of sensed mobile components must also be restrained to within safe operating envelopes. Other possible operating protocols of this highly sophisticated medical nanorobot are outlined in Section 15.4.

Specialized protocols required for chromalloyocytes could be particularly complex. Significantly modified procedures will be required for numerous unusual cases including (1) proliferating, pathological, multinucleate, and karyolobate cells, (2) cells in locations where access is difficult such as brain, bone, or mobile cells, (3) cells expressing genetic mosaicism, and (4) alternative missions such as mitochondrial DNA replacement [7].

15.3. COMMON FUNCTIONS REQUIRING ONBOARD COMPUTATION

Functions which may be common to many different classes of medical nanorobots and must be controlled by the onboard nanocomputer include those listed below.

1. *Pumping.* Single-molecule recognition, sorting, and pumping via molecular sorting rotors [1f] to allow molecule-by-molecule exchanges with the environment. The typical ~ 1 micron nanorobot might have $\sim 10^4$ – 10^5 sorting rotors requiring individual control, though many can probably be operated in banks or clusters of 10–100 rotors. Molecular pumps would be a primary system in nanorobots such as pharmacytes [10] which are intended solely to dispense drugs or other reagents.
2. *Sensing.* Chemical sensors [1j], pressure [1k] and temperature [1i] sensors, ullage sensors (include onboard pressure tank management) [1o], electrical, magnetic and optical sensors [1av], position/orientation sensors [1aw], gravity sensors [1ax], and molecular recognition sites [1ay, 1az]; and coordination and interpretation of macrosensing data providing onboard information from external acoustic [1ba], proprioceptive [1aw], electric/magnetic [1bb], or optical [1bc] sources, or direct neural traffic eavesdropping [1bd]. The typical ~ 1 micron nanorobot might employ $\sim 10^4$ – 10^5 sensors of various kinds requiring individual control.
3. *Configuration.* Control of device shape [1be]; gas-driven extensible bumpers to maintain physical contact among adjacent devices during

stationkeeping [1bf]; control of internal ballasting for nanapheresis [1n]; and control of chemical ligands for hull displays, flags, or semaphores [1bg], or for controlled adhesioregulation of external surfaces [2k].

4. *Energy.* Control of onboard power generation or power receiver systems including thermal, mechanical, acoustic, chemical, electrical, photonic, or nuclear sources; management of onboard energy storage; controlling the transduction, conditioning, and conversion of rf or tethered energy sources; and control of internal power distribution and load balancing throughout a nanorobotic device [1c].
5. *Communication.* Control of communications hardware including receivers and transmitters, whether chemical, acoustic, electromagnetic, or other modality; interpretation of received signals as new commands from the physician; replacement of existing operating parameters with new ones (while the devices are *in vivo*), such as the changing of respirocyte operating parameters at high altitudes or deep underwater at high pressures; coordination of internal fiber or mobile networks; control of inmessaging (signals from external sources that are directed to *in vivo* nanorobots) and outmessaging (signals from *in vivo* nanorobots that are directed to external recipients), including coordination of nanorobot populations to accurately transfer information directly to or from the patient (e.g., directly via somesthetic, kinesthetic, auditory, gustatory, olfactory, ocular or dermal [1bq] displays, or indirectly via artificial symptoms); and routing of data signals internally throughout the nanorobot [1bh].
6. *Navigation.* Establishing absolute or relative physical position across many regimes including bloodstream, tissues, organs, and cells; positional navigation by dead reckoning, cartotaxis, macro/microtransponder networks (including stationkeeping protocols among neighboring intercommunicating devices and self-correcting calibration protocols allowing ~ 3 micron positioning accuracies in ~ 1000 -device stacks maintaining ~ 100 -micron interdevice separations across 10-cm tissue columns); functional navigation using thermographic, barographic, chemographic, or microbiotic modalities; and validation of the identity of encountered tissue types [1an].
7. *Manipulation.* Deployment and actuation of manipulators including ciliary, pneumatic, or telescoping systems; stowage, retrieval, selection, installation, use, and detachment of tooltips and other end-effectors; management of tool and manipulator garages; management of coordinated manipulator arrays; and control of onboard disposal or disassembly systems including morcellation, grinding, sonication, thermal or chemical decomposition systems [1b].
8. *Locomotion.* Control of specific *in vivo* locomotion systems [1bi] including ciliary or grapple systems, surface deformation, inclined planes/screws, volume displacement, and viscous anchoring systems; control of locomotion across cell-coated tissue surfaces (e.g., vascular lumen) via

legged ambulation, tank-tread rolling, amoeboid motion or inchworm locomotion; control of histonataion (tissue swimming) and cytopenetration as required; cytocarriage (nanorobots controlling movement of natural cells such as leukocytes from intracellular berths); and *ex vivo* locomotion [1bj] including dental walking, epidermal locomotion, and airborne nanoflight.

9. *Computation.* Control of all onboard computer systems [1d] including CPUs, memory, and internal data transmission lines, and procedures for switching between redundant computers or computer components; and the management, synchronization and calibration of onboard clocks and calendars [1bk]. Nanocomputers that employ standard dissipative architectures (as opposed to energy-saving reversible computing architectures [33–35]) will generate a high power density of waste heat, hence nanocomputers in larger nanorobots will commonly be throttled down from maximum processing speeds during normal operations both to save energy and to avoid excessive localized device heating. This is a principal design limitation on medical nanorobot computers. Consequently, a good design philosophy is to offload as many computational tasks as possible to *in vivo* data processing devices located external to the nanorobot (e.g., fixed-location tissue-embedded *in vivo* nanorobotic implanted nodes), or preferably to large *ex vivo* computers embedded in the physician's office, the surgical suite, or the hospital infrastructure, with data transferred in and out of the patient's body through fixed or mobile communications networks [1bm].
10. *Redundancy Management.* Acceptable system reliability for populations of trillions of cooperating medical nanorobots will require extensive subsystem and component redundancy. Typically, tenfold redundancy among mission-critical components appears sufficient to ensure acceptable mission reliability [12g] in therapeutic applications. For example, a good nanorobot design [4–8] may specify 10 times more sorting rotors, sensors, or appendages for locomotion than are strictly necessary, with the large surplus held in reserve as spares and backups. This implies that another important computational function will be redundant systems management [60, 61] and modeling [62]. Onboard computers must continuously monitor the performance of all redundant components and subsystems to determine whether or not any have failed, and if so, to decide which backup system to swap in to replace the function performed by the failed system. Studies are needed to define optimal control protocols for redundant systems management. As a simple example, in the case of the respiocytes which employ tenfold-redundant onboard glucose-oxygen power stations, reliability simulations may be useful to determine whether all ten power stations ideally should be run: (a) at peak power on a rotating schedule, (b) at partial power on a continuous basis, or (c) one at a time until failure, thereafter switching to

the next backup. During any component switchovers, the respirocyte computer must control the real-time distribution of power that is transmitted hydraulically to local station subsystems and also along a dozen independent interstation trunk lines that allow stations to pass hydraulic power among themselves as required, permitting load shifting and balancing. Redundancy management also applies to onboard computers, which are themselves multiply redundant. Further analysis is needed to determine the best techniques for safely switching among them without sacrificing system reliability.

11. *Flawless Compact Software.* Another fundamental limit to onboard computing is the minimum practical volume that may be occupied by registers or other physical media capable of storing or executing onboard data or instructions. For example, a 1 micron³ storage volume could contain $\sim 10^{10}$ bits (~ 1 GB assuming 8-bit words) on a hydrofluorocarbon memory tape with a maximum ~ 10 sec access time, or $\sim 10^7$ bits (~ 1 MB) in a tightly packed block of mechanical registers comprised of nanoscale diamond rods with a maximum ~ 1 millisecc access time [1af]. Given the presumed requirements for tenfold redundancy and the need for non-memory computer components, maximum allowable onboard data storage may often be an order of magnitude smaller than the above figures, or less. Given such limited availability of onboard nanorobot computer memory, it seems inevitable that efficient ultracompact software and space-efficient algorithms will again come into vogue. Such software was commonplace in the early days of PC development when total RAM memories were often limited to 48 KB or less—a meager allocation that was nonetheless sufficient to hold an entire word processing program. High software reliability must also be a major design criterion [63] because medical nanorobots must be extremely reliable. People have already been killed or almost killed [64] by software bugs in conventional medical devices. Policies for medical and product liability insurance will almost certainly incorporate strict requirements for provable software reliability and system noncrashability as measured by formalized industry metrics because of the potential significant risks to human health if a nanorobot control malfunction due to software error occurs during therapeutic use—a stringent requirement that should, at long last, make the formalized production of “error-free” code economically viable. The related supervisory software must also be designed with a simple user interface to minimize the possibility of error even when operated by weary, distracted, and fallible human medical personnel [65].

15.4. NANOROBOT CONTROL PROTOCOLS

A nanorobot control protocol is a specialized sequence of nanorobot actions that may be executed in carefully specified circumstances in order to ensure that the

defined mission is successfully completed in an accurate, safe, and timely manner according to plan. Protocol-driven actions may be obligate (mandatory) or facultative (permissive) as required. Such actions may be initiated in response to sensor input (ranging from simple inputs such as ambient temperature or chemical concentration, to more complex inputs such as detection of engulfment by a phagocyte), clock or calendar datemarks, completion of intermediate mission benchmarks, or external commands.

Six general classes of nanorobot control protocols have been identified and are described below. All nanorobots will require some of these protocols, and some may require most or all of them.

15.4.1. Operational Protocols

Operational protocols control the simplest and most basic procedures that nanorobots must perform. Such protocols may control the operation of individual or groups of sensors or sorting rotors, or may regulate internal power transfer and energy storage using logic-branched feedback loops. In the case of the respirocyte, a proper analysis of operational protocols would involve drawing up an operational flowchart for its normal behaviors, including alternate modes, operational options, parameter lookup tables, communications interfaces, load levels, switching among backups, internal state monitoring, sensor integration with sorting rotor function, engine throttling, and device-specific operational modes—e.g., resting, exercising, asphyxiating, underwater, high altitude, and so forth. Also included here are standardized user interfaces and user protocols to make it as convenient as possible for doctors to access and reprogram the nanodevices while ensuring patient safety.

The nanorobot must also be able to follow its activity script and recognize when its mission has ended, and then be conveniently shut down. In the case of the microbivore, the end of the mission is not clearly defined because another target microbe might always be found, hence the termination of this nanorobot's open-ended mission most likely will be determined by a clock or will be externally commanded by the attending physician. In other cases, like the chromalloocyte, the mission sequence is clearly defined and has a specific end state because the device can track its progress through the preprogrammed activity script checklist as a function of completed tasks and external events. Activity scripts may be the most complex form of operational protocol.

15.4.2. Biocompatibility Protocols

Most nanorobots may require active processes to maintain biocompatibility [2]. These processes will require computer control and may range from simple protocols such as controlling alterable surface chemical modifications (e.g., adhioregulation [2k]) and the microbial tail protocol (Section 2.2.2) to relatively complex processes such as phagocyte avoidance and escape protocols [2d] or the splenofenestral escape protocol [5] (for bloodborne nanorobots above a certain size), activation of which will initiate a stereotyped set of locomotive movements

that result in breaking free from splenic captivity and returning to the bloodstream (Section 2.2.2).

Numerous other biocompatibility situations requiring specialized protocols are identified elsewhere [2]. One interesting example involves the commandeering of natural motile cells by medical nanorobots, a procedure known as cytocarriage [1bn], as an alternative mode of *in vivo* transport. During cytocarriage, one or more medical nanorobots may enter a motile cell, ride or steer the cell to a desired destination inside the human body, then vacate the cell upon arrival. When macrophages and other leukocytic cells become infected, they express B7 on their membrane surface which can be recognized by a T-cell CD28 receptor protein, triggering an immunologic response. In the “infected cytovehicle protocol,” nanorobot pilots would inspect the cell surface of a prospective cytovehicles for B7 and similar flags prior to cytopenetration to avoid choosing an infected cell for cytocarriage that could then spread the infection. If the cell subsequently becomes infected and begins expressing B7 or other warning substances during the journey, the protocol would direct the nanorobot pilot to abort the mission and steer the cytovehicle to a nearby disposal site, or implement immediate therapeutic measures; failing this, the pilot should abandon the vehicle at once.

15.4.3. Theater Protocols

A theater protocol is a control process by which the nanorobot verifies that it is performing its tasks in the intended location or in the desired circumstances, as prescribed by the attending physician. If the nanorobot determines that its present location is inappropriate or the necessary conditions to enable action no longer exist, the protocol triggers an appropriate corrective or defensive response. When a nanorobot detects an out-of-theater condition (e.g., bled out of body), the most appropriate response may often be to save the device and then shut down. Subsequent reentry into theater after shutdown (e.g., bled-out devices find their way back into the bloodstream) normally should not allow reactivation of the device because during its absence from theater, the nanorobot may have been subjected to unknown forces, chemicals, radiation, or even reprogramming that could render it harmful to the patient upon reentry, if reactivated. Even if a full pre-restart self-diagnostic routine could be performed after an out-of-theater excursion, unwanted modifications may be too subtle to detect via internal sensors and the risk of malicious tampering is too great. One possible exception is where the violated theater of operation is a functional condition such as “patient is asleep,” “patient is drunk,” or “patient is sexually aroused.” In these cases, it is less likely that switching on or off in response to theater protocols would pose significant risk either to the devices or to the patient because these conditions are normally cyclical volitional states that define inherently ephemeral or periodic theaters of action.

The longer a therapeutic nanorobotic medical procedure takes to complete, the more times a theater protocol should be executed during the mission. Protocols might be checked more or less frequently depending upon their mission criticality,

their impact on safety, the likelihood of a rapid change in status, their computational overhead (the more often the protocol is run, the fewer computer resources are available for primary tasks), and data acquisition time (e.g., detecting a very rare molecule in the bloodstream might take many seconds or minutes; the reprotocol cycle time should be significantly longer than the maximum practical sensing period). Some protocols might be checked continuously—biochemical processes may occur in milliseconds, other biological responses may require seconds or minutes (e.g., automobile fuel gauges are rechecked electronically about once every second.)

It is worth noting that theater protocols are found in natural human cell biology. For example, most tissue cells require attachment to the ECM and subsequent spreading on the ECM, for proper growth, function, and even survival, a feature known as anchorage dependence [66]. Upon losing this anchorage and departing from the desired theater of operation, cells often die by undergoing apoptosis or programmed cell death. (Only the cells circulating in the blood are designed to survive without attachment and spreading—some tumor cells acquire this ability and leave their original tissue site to form metastases, a pathological state.) In this theater protocol, the cell is designed to default to suicide, but defers this action as long as a survival signal is received, indicating continued residence in the permitted theater of operations. The survival signal is generated as long as the cell maintains integrin attachment to a basement surface or ECM. This protocol trigger is called an apoptotic switch [67]. T-cells employ a similar theater protocol via the “immunological synapse” [68].

There are at least six basic classes of theater protocols.

1. *Locational.* Nanorobot activities may be locationally restricted to certain specific regions of the body (e.g., heart, liver, epidermis, skull, ear canal, right arm, big toe, or bloodstream). They may be restricted negatively—allowed to operate only as long as they are not present in a particular location—or may have a combination such that they exhibit two distinct modes of behavior (e.g., location A behavior vs. location B behavior, or location A versus location not-A behaviors). Thus a nanorobot that has been bled out of the body or has been surreptitiously syringed from one body to another would automatically deactivate. Since nanorobots can contain reliable onboard geographic location and gravity sensors [1ax], nanorobots could vary their activities as the patient moved across the surface of the Earth, and astronaut-implanted nanorobots could vary their behaviors according to whether the human is (a) on the ground, (b) experiencing high-acceleration during boost phase, or (c) experiencing on-orbit weightlessness.
2. *Functional.* Nanorobots may include theater protocols to restrict their actions according to functional criteria, such as where a set of specified conditions exists (e.g., only in a cancer tumor secreting certain biochemical factors), or only where blood composition lies within certain preset limits, or only in bloodstreams in which large amounts of alcohol, nicotine,

hallucinogenic drugs, or other bioactive substances are present (or absent). Nanorobots might be restricted to operations only when blood pressure moves above or below certain preset thresholds, or when respiration or heart rate is too fast or too slow, or within a defined range. Measured progesterone level could serve as one indicator of pregnancy, and a high blood concentration of the hormone relaxin could indicate that the adult patient has recently undergone childbirth. Other devices might be programmed to operate only during periods of sexual arousal (as determined operationally by a set of critical body-wide physiological parameters), or only when the patient is hungry, or is having an epileptic seizure, or is in REM-phase sleep, or only when the patient is talking, sneezing, defecating, or urinating (all of which have unique and distinct physiological adjuncts that can be detected, measured, and acted upon by a sufficiently discriminating acoustic macrosensing system [1ba]). Another functional protocol that might be useful in some cases is a determination as to whether the patient was functionally “alive” or “dead”—as provisionally determined by a molecular analog to the Karnofsky Performance Score [69] using a constellation of sensor readings from an *in vivo* nanorobot population including blood velocity, tissue oxygenation or carbonation levels, glucose depletion, monitoring of mechanical body noises and body temperature, presence of chemical poisons or forensic biochemical death markers, absence of electrical or other neural activity, dysfunctional ionic balances, and so forth.

3. *Situational.* Nanorobots could be programmed with narrowly specialized situational protocols driven by complex multiple sensor data designed to detect rare events. For example, one situational theater might be defined as the simultaneous receipt of sensor data from nanorobots stationed throughout the body indicating (1) high concentrations of adrenalin in the bloodstream, (2) a sudden transition from a gravity-detected to a low-gravity (e.g., free-fall) environment as recorded by acceleration and gravity sensors, (3) whistling noises detected by nanorobots stationed at the cochlear nerves, (4) continuous rotational movements detected by the vestibular apparatus, (5) drop in ambient air pressure in the lungs, and (6) kinesthetic sensors detecting flailing motions of arms and legs. The simultaneous occurrence of all these sensory events might suggest that the patient was falling through the air, and might provide the unique theater of operation of nanorobots designed to automatically respond to this event. Other more common situations might also be recognized using situational protocols. For instance, determination that the patient is wearing clothing, and the approximate extent of coverage of the epidermis by that clothing, may be detectable by a sophisticated whole-body nanorobot population by sampling dermal, neural, pressure, and temperature receptor output, by measuring variations in transdermal light transmission, or by mapping variations in sweat gland cell metabolic rate [1bo].

4. *Phenotypic.* Nanorobots might be allowed to operate only in patients that have (or lack) specific detectable phenotypes. Membership in such classifications might be assessed biochemically by identifying specific sets of bloodborne substances. For example, determining the ratio of estrogen to testosterone might be one of several coordinated tests for gender. Biological age might be roughly estimated by combining measurements of glucosamine levels (which increase with age), testosterone levels (which often decrease with age), racemization of proteins, and other biochemical markers of aging. Obesity might be detectable by sampling the ratios of specific classes of lipid carriers or monitoring leptin and other biochemicals in the blood. It might be possible to detect certain behavioral traits such as thrill-seeking, shyness, psychopathology, or schizophrenia, all of which are known or suspected to be associated with significant biochemical or hormonal variations that are in principle detectable by bloodborne nanorobots. Hair or skin color might also be measurable, if the pigmentation generates detectable bloodborne adjuncts. Recent research indicating possible physiological and biochemical correlates of sexual preference [70–72] may put this into the class of phenotypes that could be selectable as permissible or restrictive theaters of nanorobotic operation, though reliability may be a significant challenge in the implementation of certain phenotypic protocols.
5. *Temporal.* The theater of nanorobot activity may be further limited by temporal restrictions to operations performed at certain specific times or with measured durations. Examples of temporal protocols might include restrictions to operation only at night (e.g., when the patient is likely to be asleep) or at a specific time of day (e.g., between 4 AM and 5 AM in the morning). Activity may be event-sequenced—permitted only at specific points in the daily circadian rhythm, or in synchronization with other natural rhythms of the body such as brain waves, gastric oscillations, cardiac impulses, respiratory cycling, daily cortisol variations, menstrual cycles, or coughing staccata. Nanorobots could be restricted to operating for only one minute out of every hour, whether simply to save power or alternatively to perform a lengthy procedure in smaller increments for medical reasons. Temporal protocols might apply only for a set period of time followed by termination of activity, or may be keyed to the calendar, activating, for example, only at midnight on New Years Eve or annually on one's marital anniversary. (These latter instances may presume a longer-lived nanorobotic residence *in vivo*.)
6. *Identity.* Nanorobots may also be subject to identity protocols in which activities are permitted only if the device can confirm that it is present within the body of a specific person and no one else. This sort of control may be enabled via biochemical biometrics in which onboard sensors are used to detect specific cytochemical or plasma membrane self-markers [1ag] such as HLA specificities in MHC Class I molecules on cell

membrane, blood group antigens on red cells surfaces, and other individual biochemical signatures other than genetic differences (which are not directly accessible to most nanorobot species). An alternate method for uniquely identifying individual people is to insert an artificial gene into all cells of each person from birth whose sole purpose is to cause cells to synthesize a harmless simple chemical marker, perhaps a digitally encoded short carbohydrate chain [73]—even a small hexasaccharide can encode $\sim 10^{12}$ unique combinations [74]—that is released into the bloodstream and is readily detectable and decoded by theater-restricted medical nanorobots. The deployment of such self-renewable *in vivo* fiducial chemical markers could serve as a convenient substitute for less-intrusive phenotypic measurements but might be regarded as controversial much like the 1999 USPTO-granted patent for barcoding humans [75].

15.4.4. Safety Protocols

Safety engineering involves making sure things do not fail in the presence of random faults. As Dorner [76] notes that

learning theory tells us [that] breaking safety rules is usually reinforced, which is to say, it pays off. Its immediate consequence is only that the violator is rid of the encumbrance the rules impose and can act more freely. Safety rules are usually devised in such a way that a violator will not be instantly blown sky high, injured, or harmed in any other way but will instead find that his life is made easier. The positive consequences of violating safety rules reinforce our tendency to violate them, so the likelihood of a disaster increases. And when one does in fact occur, the violator of safety rules may not have another chance to modify his behavior in the future.

Since the risk of harm is great if medical nanorobots are misused, future designers should try to make it as hard as possible to disable the safety features. To the greatest degree practical, these features should be permanently embedded in hardware to minimize the probability of circumvention.

Nanorobots control systems should employ failsafe designs which may incorporate parallelism (dividing tasks among a large number of simple systems), specialization (individual systems optimized for particular tasks), and redundancy (comparing the output of multiple systems to improve reliability of the results). These built-in safeguards are then enhanced by the use of safety protocols which are active device behaviors designed to further enhance and reinforce safety. Failsafe designs must ensure that even a total system failure will not lead to death or serious injury of the patient. Safety protocols should be able to recognize various catastrophic internal failure states including compromised physical structure or software/data corruption, necessitating localized, intermediate, or even whole system shutdown, with entry into a safe-harbor mode that may permit

the nanorobot to engage in actions that will safely remove it from the body since it can no longer fulfill its mission. One simple failsafe mechanism widely employed in high-reliability programmable devices is the watchdog timer—a counter that can shut down the computer if it ever reaches zero, but which is continually reset by a correctly operating program so that it never reaches zero as long as the program continues functioning (analogous to the biological “apoptotic switch” described earlier). In the case of medical nanorobots that require active supervision, the reset command could be periodically rebroadcast to the nanodevices *in vivo* by the attending physician via ultrasound messaging.

Safety protocols may range from relatively simple procedures, such as self-diagnostic routines, to very complex procedures such as dental protocols which instruct mobile nanorobots bled into the mouth to avoid the hard grinding surfaces by retreating to lower positions on the teeth [1z], or flight protocols for aerial nanorobots including self-enforcement of no-fly zones near nose and mouth along with other active anti-inhalation, inhale-safe, and post-inhalation extraction protocols [1bp]. Other safety protocols may incorporate a wide variety of user-set locks and limits, command blocking, limits on access to functions by non-physicians, and perhaps some equivalent to the humorously termed “shame blocker” (a recent telephone gimmick in which the user dials 333 and a number, and then his phone won’t let him call that number, e.g., his ex-wife when he’s drunk).

15.4.5. Security Protocols

Even the best security cannot prevent all harm but can help avoid significant harm while allowing the system to continue operating normally. While most features are useful for what they do, security features within products are useful because of what they don’t allow to be done. Security engineering helps to ensure that the nanodevice will not fail “in the presence of an intelligent and malicious adversary who forces faults at the worst time and in the worst way.” Security protocols are required to ensure that incoming commands originate from trusted and authorized sources, and may employ checksums, signed and certified programs, and formally proofed systems without trapdoors. These protocols should normally refuse to accept commands that could cause the device to exhibit behaviors that would harm the patient. Communication protocols might include TCP/IP protocol stacks combined with typical security technologies such as firewalls, packet filtering, intrusion detection, and secure procedures for flow control and authentication (e.g., passwords, biometrics, and public-key cryptography) to prevent “body-hacking.” Depending on circumstances, nanorobots may need blockers for viruses, worms, and spam, and protocols linked to sensors that monitor structural integrity to ensure tamper resistance.

Other security protocols may be required to authorize reading and writing data into personal medical record caches implanted in the patient’s body—multiply redundant caches $\sim 1\text{ mm}^3$ in size could hold $>1000\text{ TB}$ of fast-access mechanical memory, capacious enough to store a lifetime of detailed medical data,

thus simplifying and perfecting medical recordkeeping while making lifesaving data more quickly accessible to emergency medical personnel. Security protocols must safeguard not just these internal data stores, but also access to the patient's implanted internal communications network and any other *in vivo* nanorobotry that might be present. Other security protocols enforcing digital rights management (DRM), audit trails, process accounting, expiration dates and related schemes may be employed on commercial systems to enforce patent rights. However, any regulations requiring implementation of embedded government-sponsored controls or content filtering should be regarded with deep suspicion.

15.4.6. Group Protocols

Group protocols may be required to control the collective behaviors of large populations of simultaneously interacting *in vivo* medical nanorobots. It is hoped that relatively simple individual nanorobot behaviors [77] can be programmed that will give rise to more complex desired group behaviors. Multirobot control algorithms are a major research field today with precedents in agoric algorithms [78–81], stigmergy [82], swarm computing [83, 84], and agent-based systems [RX80].

15.5. CONCLUSIONS

Medical nanorobots may be constructed of diamondoid nanometer-scale parts and subsystems including onboard sensors, motors, manipulators, power plants, and molecular computers. The presence of onboard nanocomputers will allow *in vivo* medical nanorobots to perform numerous complex behaviors which must be conditionally executed on at least a semiautonomous basis, guided by receipt of local sensor data, constrained by preprogrammed settings, activity scripts, and event clocking, and further limited by a variety of simultaneously executing real-time control protocols.

Such nanorobots cannot yet be manufactured, but preliminary scaling studies for several classes of medical nanorobots including respiocytes, microbivores, clottocytes and chromalloyocytes have been published in the literature. These designs allow an analysis of basic computational tasks and a summation of major computational control functions common to all complex medical nanorobots. These functions include the control and management of pumping, sensing, configuration, energy, communication, navigation, manipulation, locomotion, computation, and the use of redundancy management and flawless compact software.

Nanorobot control protocols are required to ensure that each nanorobot completes its intended mission accurately, completely, safely, and in a timely manner according to plan. Six major classes of nanorobot control protocols have been identified and include operational, biocompatibility, theater, safety, security, and group protocols. Six important subclasses of theater protocols include

locational, functional, situational, phenotypic, temporal, and identity control protocols.

ACKNOWLEDGMENTS

Thanks are due to Forrest Bishop for the respirocyte and microbivore images, Zyvex Corp. for the use of the microbivore image, and John Luu and his colleagues at Stimulacra LLC for the chromalloyocyte images. Research grants from Alcor Foundation, Life Extension Foundation, Kurzweil Foundation, and the Institute for Molecular Manufacturing supported this work.

REFERENCES

1. R. A. Freitas, Jr. *Nanomedicine, Volume I: Basic Capabilities*. Georgetown, TX: Landes Bioscience, 1999; <http://www.nanomedicine.com/NMI.htm>; (a) Chapter 4, (b) 9.3, (c) Chapter 6, (d) 10.2, (e) 10.3, (f) 3.4.2, (g) 6.3.4, (h) 4.2.1, (i) 4.6, (j) 4.2, (k) 4.5, (m) 7.2.2, (n) 10.3.6, (o) 4.5.4, (p) 6.3.2, (q) 6.4.1, (r) 7.2.5.3, (s) 6.4.3, (t) 7.2.5.4, (u) 6.3.4.4, (v) 6.3.4.5, (w) 6.3.4.2, (x) 6.4.3.4, (y) 10.2.1, (z) 9.5.1, (aa) 7.4.6, (ab) 10.4.2.4.2, (ac) 8.5.2.2, (ad) 9.3.1.4, (ae) 9.3.5.1, (af) 7.2.6, (ag) 8.5.2.1, (ah) 8.2.1.1, (ai) 9.4.1.4, (aj) 3.2, (ak) 10.5.4, (am) 7.4.5.2, (an) Chapter 8, (ao) 9.4.4.1, (ap) 9.4.4, (aq) 9.4.5, (ar) 9.4.5.3, (as) 9.4.4.2, (at) 9.4.2.5.1, (au) 8.3.3, (av) 4.7, (aw) 4.9.2, (ax) 4.9.2.4, (ay) 4.2.8, (az) 8.5.2, (ba) 4.9.1, (bb) 4.9.3, (bc) 4.9.4, (bd) 4.9.5, (be) Chapter 5, (bf) 5.4, (bg) 5.3.6, (bh) Chapter 7, (bi) 9.4, (bj) 9.5, (bk) 10.1, (bm) 7.3, (bn) 9.4.7, (bo) 8.4.1.3, (bp) 9.5.3.6, (bq) 7.4.6.7.
2. R. A. Freitas, Jr. *Nanomedicine, Volume IIA: Biocompatibility*. Georgetown, TX: Landes Bioscience, 2003; <http://www.nanomedicine.com/NMIIA.htm>; (a) 15.6.2, (b) 15.4.3.1, (c) 15.4.3.2.3, (d) 15.4.3.6, (e) 15.4.2.3, (f) 15.4.2, (g) 15.4.3.6.1, (h) 15.4.3.6.4, (i) 15.2.5, (j) 15.5.2.2, (k) 15.2.2.4, (m) 15.4.3.4.
3. R. A. Freitas, Jr. Current status of nanomedicine and medical nanorobotics (Invited Survey). *Journal of Computational and Theoretical Nanoscience*, 2: pp 1–25, Mar 2005; <http://www.nanomedicine.com/Papers/NMRevMar05.pdf>.
4. R. A. Freitas, Jr. Exploratory design in medical nanotechnology: A mechanical artificial red cell. *Artificial Cells, Blood Substitutes, and Immobilization Biotechnology*, 26: pp 411–430, 1998; <http://www.foresight.org/Nanomedicine/Respirocytes.html>.
5. R. A. Freitas, Jr. Microbivores: Artificial mechanical phagocytes using digest and discharge protocol. *Journal of Evolution and Technology*, 14: pp 1–52, Apr 2005; <http://www.jetpress.org/volume14/freitas.html>.
6. R. A. Freitas, Jr. Clottocytes: Artificial Mechanical Platelets. Foresight Update No. 41: pp 9–11, June 30, 2000; <http://www.imm.org/Reports/Rep018.html>.
7. R. A. Freitas, Jr. The ideal gene delivery vector: chromalloyocytes, cell repair nanorobots for chromosome replacement therapy. Submitted 2007.
8. R. A. Freitas, Jr. Christopher J. Phoenix. Vasculoid: a personal nanomedical appliance to replace human blood. *Journal of Evolution and Technology*, 11: pp 1–139, Apr 2002; <http://www.jetpress.org/volume11/vasculoid.pdf>.

9. R. A. Freitas, Jr. Nanodentistry. *The Journal of the American Dental Association*, 131: pp 1559–1566, Nov 2000; <http://www.rfreitas.com/Nano/Nanodentistry.htm>.
10. R. A. Freitas Jr. Pharmacytes: An ideal vehicle for targeted drug delivery. *Journal of Nanoscience and Nanotechnology*, 6: pp 2769–2775, Sep–Oct 2006; <http://www.nanomedicine.com/Papers/JNNPharm06.pdf>.
11. Nanofactory Collaboration website. Accessed Jan 2007. <http://www.MolecularAssembler.com/Nanofactory>.
12. K. E. Drexler. *Nanosystems: Molecular Machinery, Manufacturing, and Computation*. New York: Wiley, 1992; (a) Chapter 10, (b) Chapter 11, (c) 13.4, (d) Chapter 12, (e) 16.3.2, (f) 12.4, (g) 6.7.2.
13. R. A. Freitas Jr. and R. C. Merkle. *Kinematic Self-Replicating Machines*. Georgetown, TX: Landes Bioscience, 2004; <http://www.molecularassembler.com/KSRM.htm>.
14. K. Ishiyama, M. Sendoh, and K. I. Arai. Magnetic micromachines for medical applications. *Journal of Magnetism and Magnetic Materials*, 242–245: pp 1163–1165, 2002.
15. D. D. Chrusch, B. W. Podaima, and R. Gordon. Cytobots: intracellular robotic micromanipulators. In: W. Kinsner, A. Sebak, editors. Conference Proceedings, 2002 IEEE Canadian Conference on Electrical and Computer Engineering. Winnipeg, Canada: 2002.
16. J. B. Mathieu, S. Martel, L. Yahia, G. Soulez, and G. Beaudoin. MRI systems as a mean of propulsion for a microdevice in blood vessels. *Biomedical Materials and Engineering*, 15: pp 367, 2005.
17. K. B. Yesin, P. Exner, K. Vollmers, and B. J. Nelson. Biomedical micro-robotic system. 8th International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI 2005), Palm Springs, California, Oct 26–29, 2005; www.miccai2005.org; p 819.
18. Micro-robots take off as ARC announces funding. Monash University, Oct 11, 2006; <http://www.monash.edu.au/news/newline/story/1038>.
19. Committee to Review the National Nanotechnology Initiative, National Materials Advisory Board (NMAB), National Research Council (NRC). *A Matter of Size: Triennial Review of the National Nanotechnology Initiative*. Washington DC: The National Academies Press, 2006; <http://www.nap.edu/catalog/11752.html#toc>.
20. R. C. Merkle and R. A. Freitas Jr. Theoretical analysis of a carbon-carbon dimer placement tool for diamond mechanosynthesis. *Journal of Nanoscience and Nanotechnology*, 3: pp 319–324, Aug 2003; <http://www.rfreitas.com/Nano/JNNDimerTool.pdf>.
21. R. A. Freitas Jr., inventor. A simple tool for positional diamond mechanosynthesis, and its method of manufacture. U.S. Provisional Patent Application No. 60/543,802, filed Feb 11, 2004; U.S. Patent Pending, Feb 11, 2005; <http://www.MolecularAssembler.com/Papers/DMSToolbuildProvPat.htm>.
22. J. Peng, R. A. Freitas Jr., R. C. Merkle, J. R. Von Ehr, J. N. Randall, and G. D. Skidmore. Theoretical analysis of diamond mechanosynthesis. Part III. Positional C2 deposition on diamond C(110) surface using Si/Ge/Sn-based dimer placement tools. *Journal of Computational and Theoretical Nanoscience*, 3: pp 28–41, Feb 2006; <http://www.MolecularAssembler.com/Papers/JCTNPengFeb06.pdf>.
23. B. Temelso, C. D. Sherrill, R. C. Merkle, and R. A. Freitas Jr. High-level *ab initio* studies of hydrogen abstraction from prototype hydrocarbon systems. *Journal of*

- Physical Chemistry A*, 110: pp 11160–11173, Sep 28, 2006; <http://www.Molecular-Assembler.com/Papers/TemelhoHAbst.pdf>.
24. H. J. Lee and W. Ho. Single bond formation and characterization with a scanning tunneling microscope. *Science*, 286: pp 1719–1722, 1999; <http://www.physics.uci.edu/%7Ewilsonho/stm-iets.html>.
25. N. Oyabu, O. Custance, I. Yi, Y. Sugawara, and S. Morita. Mechanical vertical manipulation of selected single atoms by soft nanoindentation using near contact atomic force microscopy. *Physical Review Letters*, 90: p 176102, May 2, 2003; <http://link.aps.org/abstract/PRL/v90/e176102>.
26. P. Vettiger, G. Cross, M. Despont, U. Drechsler, U. Duerig, B. Gotsmann, W. Haeberle, M. Lantz, H. Rothuizen, R. Stutz, and G. Binnig. The millipede: nanotechnology entering data storage. *IEEE Transactions on Nanotechnology*, 1: pp 39–55, Mar 2002.
27. D. Bullen, S. Chung, X. Wang, J. Zou, C. Liu, and C. Mirkin. Development of parallel dip pen nanolithography probe arrays for high throughput nanolithography. Symposium LL: Rapid Prototyping Technologies, Materials Research Society Fall Meeting; Dec 2–6, 2002; Boston, Proceedings of the MRS, Vol. 758, 2002; <http://mass.micro.uiuc.edu/publications/papers/84.pdf>.
28. Microcantilever Arrays. Protiveris Corp., 2003; http://www.protiveris.com/cantilever_tech/microcantileverarrays.html.
29. S. Martel and I. Hunter. Nanofactories based on a fleet of scientific instruments configured as miniature autonomous robots. Proceedings of the 3rd International Workshop on Microfactories; Sep 16–18 2002; Minneapolis, pp 97–100, 2002.
30. G. R. Ter Haar. Biological effects of ultrasound in clinical applications. In: K. S. Suslick, editor. *Ultrasound: Its Chemical, Physical and Biological Effects*. New York: VCH Publishers, 1988.
31. E. Ayers. An automatic chemical plant. *Scientific American*, 187: pp 82–88; 1952.
32. H. D. Luke. *Automation for Productivity*. New York: Wiley, 1972.
33. C. Bennett. The thermodynamics of computation – A review. *International Journal of Theoretical Physics*, 21: pp 905–940 1981.
34. R. C. Merkle. Reversible electronic logic using switches. *Nanotechnology*, 4: pp 21–40 1993.
35. J. S. Hall. Nanocomputers and reversible logic. *Nanotechnology*, 5: pp 157–167, 1994.
36. W. H. Weihe, editor. *Physiological Effects of High Altitude*. New York: Macmillan, 1964.
37. L. Jia, C. Bonaventura, J. Bonaventura, and J. S. Stamler. S-nitrosohaemoglobin: a dynamic activity of blood involved in vascular control. *Nature*, 380: pp 221–226, 1996.
38. H. L. Goldsmith and V. T. Turitto. Rheological aspects of thrombosis and haemostasis: basic principles and applications. *Journal of Thrombosis and Haemostasis*, 55: pp 415–435 1986.
39. M. H. Smit and A. E. G. Cass. Cyanide detection using a substrate-regenerating, peroxidase-based biosensor. *Analytical Chemistry*, 62: pp 2429–2436, 2006.
40. A. J. Erslev. *In vitro* production of erythropoietin by kidneys perfused with a serum-free solution. *Blood*, 44: pp 77–85, 1974.
41. E. F. Fincher, L. Johannsen, L. Kapas, S. Takahashi and J. M. Krueger. Microglia digest staphylococcus aureus into low molecular weight biologically active compounds. *American Journal of Physiology*, 271: pp R149–R156, July 1996.

42. C. Strange and S. A. Sahn. The definitions and epidemiology of pleural space infection. *Seminars in Respiratory and Critical Care Medicine*, 14: pp 3–8, Mar 1999.
43. L. Carreno Perez. Septic arthritis. *Baillieres Best Practice and Research Clinical Rheumatology*, 13: pp 37–58, Mar 1999.
44. J. W. Costerton, P. S. Stewart, and E. P. Greenberg. Bacterial biofilms: a common cause of persistent infections. *Science*, 284: pp 1318–1322, May 21, 1999.
45. D. Goldin. Remarks at the 15th Annual NASA Continual Improvement and Reinvention Conference. Apr 27, 2000; http://rk.gsfc.nasa.gov/richcontent/Speeches/goldin_remarks_5ps.pdf.
46. N. Leveson. Information and Computer Science Department, University of California, Irvine. Viking Lander. *The Risks Digest* 3: Oct 1, 1986; <http://128.240.150.127/Risks/3.72.html#subj1>.
47. P. Marks. Dependence day. *New Scientist*: July 31, 1999; <http://www.newscientist.com/article/mg16321971.000-dependence-day.html>.
48. Ada Information Clearinghouse. Ada in Airbus 340 Flight Warning System. ITT Research Institute, 1994; <http://www.adahome.com/Ammo/Success/aerofws.html>.
49. H. Lin. The development of software for ballistic-missile defense. *Scientific American*, 253: p 52, Dec 1985.
50. V. Cortellessa, B. Cukic, D. Del Gobbo, A. Mili, M. Napolitano, M. Shereshevsky, and H. Sandhu. Certifying adaptive flight control software. Proceedings of ISACC 2000: The Software Risk Management Conference. Reston, VA, Sep 2000; <http://www.isacc.com/presentations/3c-bc.pdf>.
51. D. E. Van Epps and B. R. Andersen. Streptolysin O inhibition of neutrophil chemotaxis and mobility: nonimmune phenomenon with species specificity. *Infection and Immunology*, 9: pp 27–33, Jan 1974.
52. A. Odegaard and J. Lamvik. The effect of phenylbutazone and chloramphenicol on phagocytosis of radiolabeled candida albicans by human monocytes cultured *in vitro*. *Acta Pathologica et Microbiologica Scandinavica. Section C, Immunology*, 84: pp 37–44, Feb 1976.
53. R. Kumar, J. E. Ansell, R. T. Canoso, and D. Deykin. Clinical trial of a new bleeding-time device. *American Journal of Clinical Pathology*, 70: pp 642–645, Oct 1978.
54. L. R. Hertzendorf, L. Stehling, A. S. Kurec, and F. R. Davey. Comparison of bleeding times performed on the arm and the leg. *American Journal of Clinical Pathology*, 87: pp 393–396, Mar 1987.
55. S. E. Lind. Chapter 33. The Hemostatic System In: R. I. Handin, T. P. Stossel, and S. E. Lux, editors. *Blood: Principles and Practice of Hematology*. Philadelphia: Lippincott, 1995 pp 949–972.
56. L. Ardekian, R. Gaspar, M. Peled, B. Brener, and D. Laufer. Does low-dose aspirin therapy complicate oral surgical procedures? *Journal of the American Dental Association*, 131: pp 331–335, Mar 2000.
57. R. L. Bick, B. Arun, and E. P. Frenkel. Disseminated intravascular coagulation. Clinical and pathophysiological mechanisms and manifestations. *Haemostasis*, 29: pp 111–134, 1999.
58. T. E. Warkentin. Heparin-induced thrombocytopenia: a ten-year retrospective. *Annual Review of Medicine*, 50: pp 129–147, 1999.

59. D. B. Brieger, K. H. Mak, K. Kottke-Marchant, and E. J. Topol. Heparin-induced thrombocytopenia. *Journal of the American College of Cardiology*, 31: pp 1449–1459, June 1998.
60. J. R. Sklaroff. Redundancy management technique for space shuttle computers. *IBM Journal of Research and Development*, 20: pp 20–28, Jan 1976; <http://www.research.ibm.com/journal/rd/201/ibmrd2001E.pdf>.
61. M. Malek. Survivable algorithms and redundancy management in NASA's distributed computing systems. Final Technical Report, NAS 1.26:189827, NASA/NTIS, 1992.
62. G. Huszler and I. Majzik. Modeling and analysis of redundancy management in distributed object-oriented systems by using UML statecharts. 27th Euromicro Conference, 2001: A Net Odyssey (Euromicro'01): p 200, 2001.
63. C. Fishman. They write the right stuff. *Fast Company*, no. 6: p 95, Dec 1996; <http://www.fastcompany.com/magazine/06/writestuff.html>.
64. M. A. Coppess, J. M. Miller, D. P. Zipes, and W. J. Groh. Software error resulting in malfunction of an implantable cardioverter defibrillator. *Journal Cardiovascular Electrophysiology*, 10: pp 871–873, June 1999.
65. T. P. Ozahowski, M. L. Greenberg, P. Mock, P. T. Holzberger, B. Gerling, C. Zalinger, and C. Perry. Implantable cardioverter defibrillator clinic casualties: inadvertent reprogramming during routine implantable cardioverter defibrillator follow-up. *Pacing and Clinical Electrophysiology*, pp 1524–1525, Oct 19, 1996.
66. E. Ruoslahti. Stretching is good for a cell. *Science*, 276: pp 1345–1346, May 30, 1997.
67. C. S. Chen, M. Mrksich, S. Huang, G. M. Whitesides, and D. E. Ingber. Geometric control of cell life and death. *Science*, 276: pp 1425–1428, May 30, 1997.
68. A. Grakoui, S. K. Bromley, C. Sumen, M. M. Davis, A. S. Shaw, P. M. Allen, and M. L. Dustin. The immunological synapse: a molecular machine controlling T cell activation. *Science*, 285: pp 221–227, July 9, 1999.
69. D. A. Karnofsky and J. H. Burchenal. The clinical evaluation of chemotherapeutic agents in cancer In: C. M. MacLeod, editor. *Evaluation of Chemotherapeutic Agents*. New York: Columbia University Press, 1949, p 196.
70. D. F. Swaab and M. A. Hofman. Sexual differentiation of the human hypothalamus in relation to gender and sexual orientation. *Trends in Neurosciences*, 18: pp 264–270, June 1995.
71. W. H. James. Biological and psychosocial determinants of male and female human sexual orientation. *Journal of Biosocial Science*, 37: pp 555–567, Sep 2005.
72. M. J. Sergeant, T. E. Dickins, M. N. Davies, and M. D. Griffiths. Women's hedonic ratings of body odor of heterosexual and homosexual men. *Archives of Sexual behavior*: Dec 22 2006.
73. N. Sharon and H. Lis. Carbohydrates in Cell Recognition. *Scientific American*, 268: pp 82–89, Jan 1993.
74. R. A. Laine. A calculation of all possible oligosaccharide isomers, both branched and linear yields 1.05×10^{12} structures for a reducing hexasaccharide: the isomer barrier to development of single-method saccharide sequencing or synthesis systems. *Glycobiology*, 4: pp 1–9, 1994.
75. T. W. Heeter, inventor. Method for verifying human identity during electronic sale transactions. United States Patent 5,878,155, Mar 2 1999; <http://patents.uspto>.

- gov/cgi-bin/iftetch4?ENG+PATBIB-ALL+0+967198+0+7+25907+OF+1+1+1+PN%2f5878155.
76. D. Dorner. *The Logic of Failure*, Metropolitan Books. (English translation). New York: Henry Holt, 1996.
 77. A. Cavalcanti and R. A. Freitas Jr. Autonomous multi-robot sensor-based cooperation for nanomedicine. *International Journal of Nonlinear Sciences and Numerical Simulation*, 3: pp 743–746, Aug 2002; <http://lipari.usc.edu/~pal/cs5xx/Cavalcanti02.pdf>.
 78. M. S. Miller and K. E. Drexler. Markets and computation: agoric open systems. In: B. A. Huberman, editor. *The Ecology of Computation*. Amsterdam: Elsevier, 1988; <http://www.agorics.com/Library/agoricpapers/aos/aos.0.html>.
 79. O. Guenther, T. Hogg, and B. A. Huberman. Market organizations for controlling smart matter In: R. Conte, R. Hegselmann, and P. Terna, editors. *Simulating Social Phenomena, Lecture Notes in Economics and Mathematical Systems*. Berlin: Springer, 1997, pp 241–257.
 80. J. S. Hall, L. Steinberg, and B. D. Davison. Combining agoric and genetic methods in stochastic design. *Nanotechnology*, 9: pp 274–284, Sep 1998; <http://www.cs.rutgers.edu/~davison/pubs/chsmith/chsmith.html>.
 81. T. Toth-Fejel. Agents, assemblers, and ANTS: scheduling assembly with market and biological software mechanisms. *Nanotechnology*, 11: pp 133–137, June 2000; <http://www.erim.org/cec/papers/NanoAgents.pdf>.
 82. E. Bonabeau, S. Guerin, D. Snyers, P. Kuntz, and G. Theraulaz. Three-dimensional architectures grown by simple ‘stigmergic’ agents. *BioSystems*, 56: pp 13–32, 2000.
 83. E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Santa Fe Institute Sciences of Complexity Series. New York: Oxford University Press, 1999.
 84. G. Theraulaz and E. Bonabeau. Modeling the collective building of complex architectures in social insects with lattice swarms. *Journal of Theoretical Biology*, 177: pp 381–400, 1995.

