a space colony than there is on earth. As machines become more capable, the economics will favor an ever higher machine-to-people ratio. Humans will not necessarily become fewer at this stage; the machines will just multiply faster, becoming ever more competent with each new generation. Imagine the immensely lucrative robot factories that could be built in the asteroids. Solar-powered machines would prospect and deliver raw materials to huge, unenclosed, automatic processing plants. Metals, semiconductors, and plastics produced there would be converted by robots into components that would be assembled into other robots and into structural parts for more plants. Machines would be recycled as they broke. If their reproduction rate is higher than the wear-out rate, the factories will grow exponentially, like a colony of bacteria, on a Brobdingnagian scale. A harvest of a small fraction of the output of materials, components, and whole robots could make investors incredibly rich.

Eventually humans, whether workers, design engineers, managers, or investors, will become unnecessary in space enterprises, as the scientific and technical discoveries of self-reproducing superintelligent mechanisms are applied to making themselves smarter still. These new creations, looking quite unlike the machines we know, will explode into the universe, leaving us behind in a cloud of dust.

Robot Bushes

The human world has been shaped by human hands, which are still our most effective general-purpose tool. Yet many useful and easily described tasks are beyond human dexterity (pull tightly on both ends of the string, while holding the knot between your fingers, lift the bundle, wrap the ends around it tightly four times...). If such actions are attempted at all, it is with varying degrees of success using special tools and fixtures.

It is unlikely that our superintelligent descendants will be satisfied with mere stumpy fingers. Consider the following observations. Worms and other animals shaped like balls or sticks are unable to manipulate or even locomote very well. Animals with legs (a stick with smaller, movable sticks) locomote quite well but are still clumsy at manipulation. Animals like us, with fingers on their legs (sticks on sticks on a stick), can manipulate much better. Now generalize the concept—a robot that looks like a tree, with a big

stem repeatedly branching into thinner, shorter, and more numerous twigs, ultimately ending in an astronomical number of microscopic cilia. Each intermediate branch would be able to swing forward and backward and side to side while its top, where the next smaller branches are attached, rotates on the branch axis. Possibly the branch could also change its length like a telescope—the number of motions

A Robot Bush

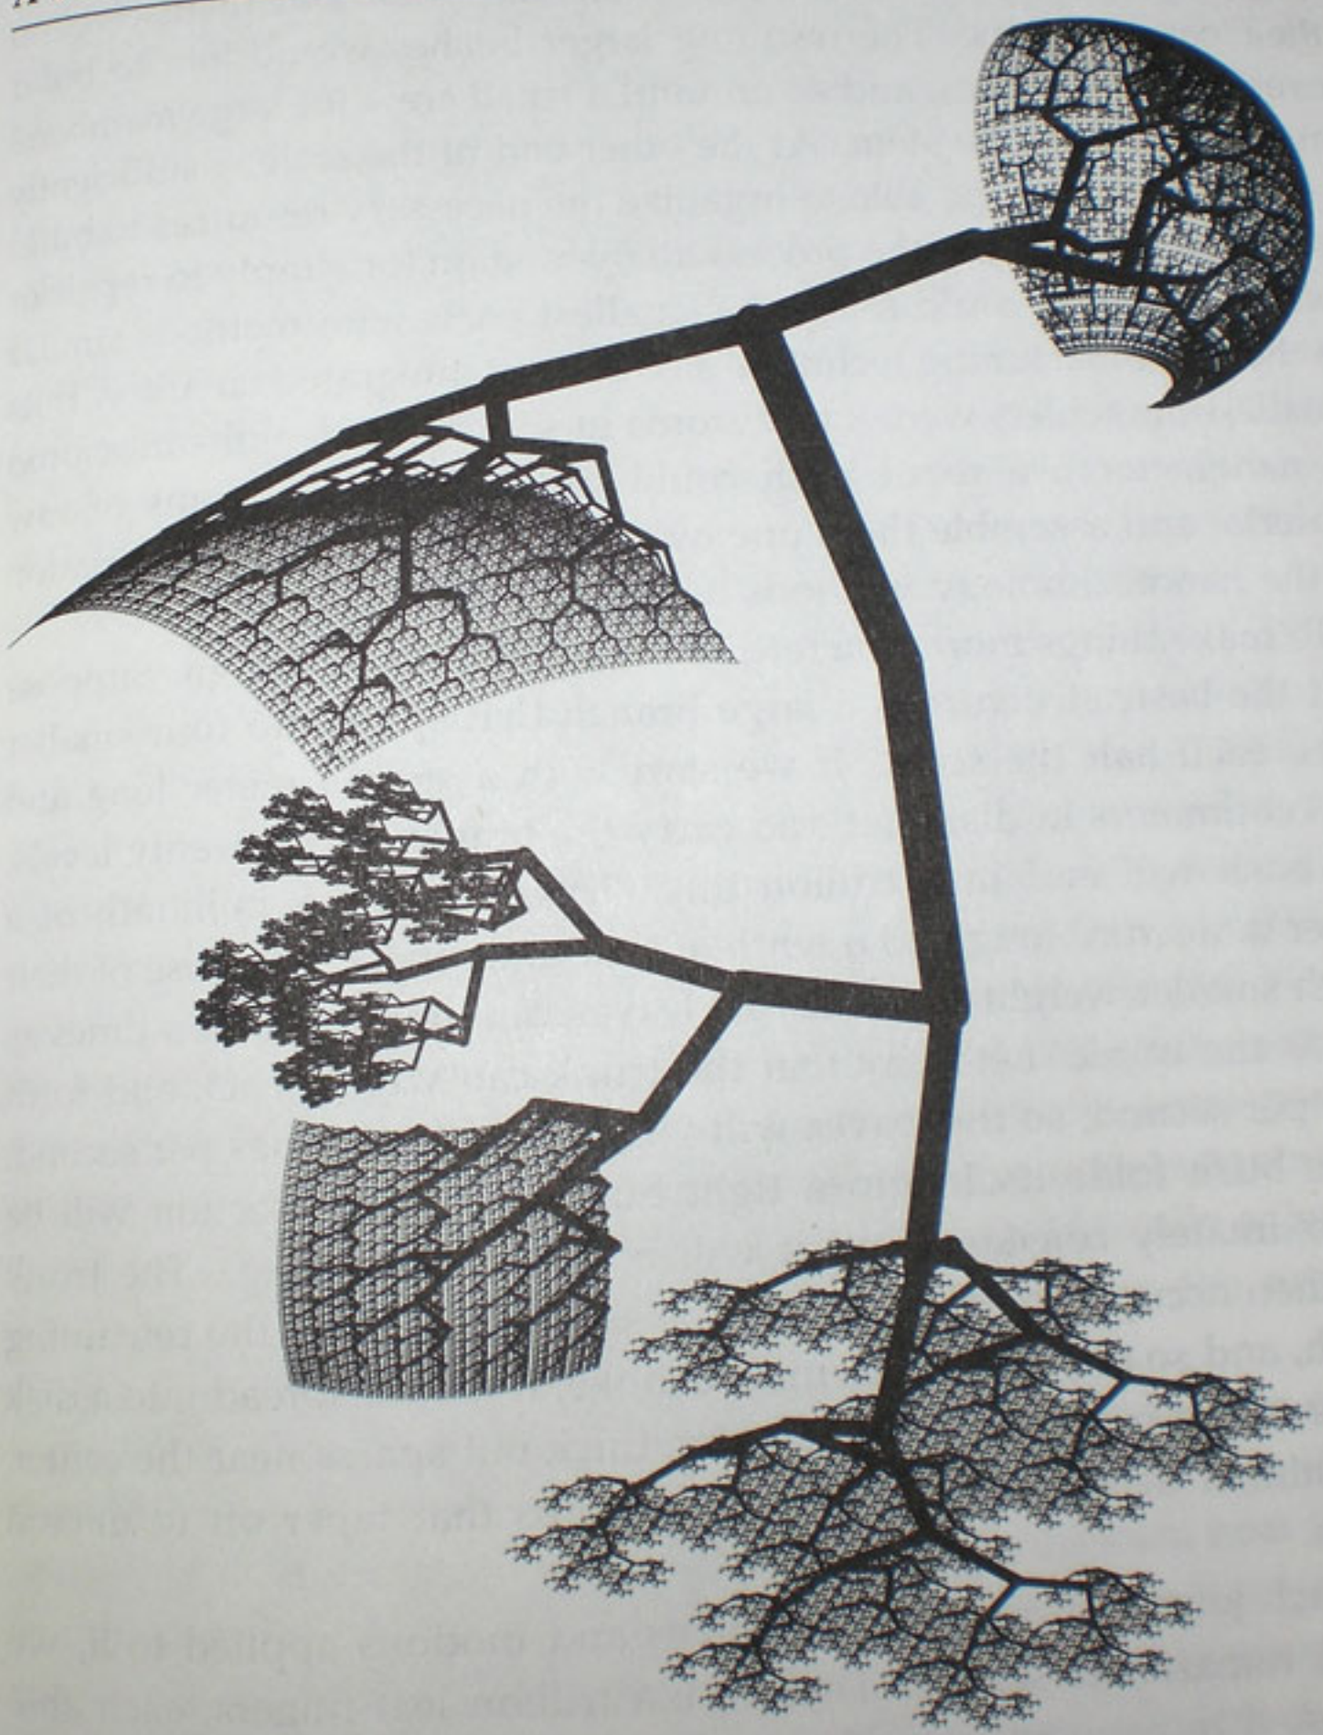

of each branch can be traded off for more levels. Each joint would have sensors to measure its position and also the force it exerts. Although made of branches, each with a rigid mechanical character, the overall structure would have an "organic" flexibility because of the great multitude of ways its parts could move.

A robot of this design could be self-constructing. Tiny bushes, only a few millionths of the weight of the final device, would be "seeded" to start the process. These would work in groups to build the next larger sprigs from available raw materials, then join themselves to their constructions. The resulting larger bushes would join to build even larger branches, and so on until a small crew (of large members) met to assemble the stem. At the other end of the scale, a sufficiently large bush should be able to organize the necessary resources to build the tiny seeds to start the process all over again (or simply to repair or extend itself). It could make the smallest parts with methods similar to the micromachining techniques of current integrated circuitry. If its smallest branchlets were a few atoms in scale (with lengths measured in nanometers), a robot bush could grab individual atoms of raw material and assemble them one by one into new parts, in a variation of the nanotechnology methods mentioned in Chapter 2.

To make things more concrete, we can do an actual design. Suppose that the basic structure is a large branch that splits into four smaller ones, each half the scale. If we start with a stem a meter long and ten centimeters in diameter and carry the branching to twenty levels, the bush will end in a trillion tiny "leaves," each a millionth of a meter (a micron) long and a tenth of that in diameter. Because of their much smaller weight and size, the leaves can move a million times as fast as the trunk. Let's say that the trunk can wiggle back and forth once per second; so the leaves will vibrate a million times per second. If the bush folds itself into a tight bundle, its cross section will be approximately constant, and it will be two meters long. The trunk will then occupy half that length, the second level half the remaining length, and so on. Unfolded, umbrellalike, it would spread into a disk a little under two meters in diameter, thick but sparse near the center, and thinner at the edge, with smaller gaps that taper off to micron spaces.

If each joint can measure the forces and motions applied to it, we have a remarkable sensor. There are a trillion leaf fingers, each able to sense a movement of perhaps a tenth of a micron and a force of

a few micrograms, at speeds up to a million changes per second. This is vastly greater than the sensing ability of the human eye, which has a million distinguishable points that can register changes at most a hundred times per second. If our bush puts its fingers on a photograph, it will "see" the image in immense detail simply by feeling the height variations of the developed silver on the paper. It could watch a movie by walking its fingers along the film as it screamed by at high speed. There is no reason the fingers could not also be sensitive to light and temperature and other electromagnetic effects; indeed, the smallest are the right size to be "antennas" for light. The bush could form an eye by holding up a lens and putting a few million of its fingers in the focal plane behind it. It may even be able to get by without the lens by holding a bunch of its fingers in a carefully spaced diffraction pattern, thus forming a holographic lens.

In addition to having a sensing capability to match that of the world's current human population, our bush would have the ability to affect its environment at the same prodigious rate. A well-trained human, using precise and well-timed hand and body motions, each able to change direction at most a few times in a second, with a precision no better than a few percent of the total movement, could conceivably affect the world at a net rate of a thousand bits per second—a fast typist, for instance, produces less than one hundred bits per second of text. The potential data rate of a robot with one trillion fingers, each able to move a million times per second, is more than a quadrillion (10^{15}) times greater. Such high data rates imply huge coordination of enormous processing power, but imagine the possibilities. The bush robot could reach into a complicated piece of delicate mechanical equipment—or even a living organism—simultaneously sense the relative position of millions of parts, some possibly as small as molecules, and rearrange them for a near-instantaneous repair. In most cases the superior touch sense would totally substitute for vision, and the extreme dexterity would eliminate the need for special tools.

An astronomical amount of thinking would be required to control this wonderful machine. Much of it might be handled by what in animals are called *reflex arcs*, small bits of nervous system near the site being controlled. Each of the small branchlets could contain enough of a computer to control most of the routine activity; only exceptional situations would require intervention from larger computers nearer

the stem. If the branches also contained their own power source (think rechargeable battery) and a way of communicating remotely (radio, or sound vibrations of a few thousand synchronized cilia, would do), the bush could break into a coordinated swarm of smaller bushes. The smaller the individual bush, the less intelligent and less powerful it would be. It would be preprogrammed, and charged up, by its home stem to perform some function and then return as soon as possible to report and receive new instructions.

Small size would frequently be an advantage: a smaller robot can squeeze into smaller spaces. A tiny machine has a greater surface-area-to-weight ratio: while a large bush could walk securely along the floor, using its branches as so many nimble toes, a smaller machine should be able to walk on ceilings like a fly, with the tiny cilia holding onto microscopic cracks in the paint or sticking by molecular adhesion. Bushes could burrow by loosening particles of dirt and passing them backward, and swim efficiently by assuming a tight, streamlined shape, with the cilia forming a skin that pumps fluid to propel and also responds to the flow to prevent turbulence. Extremely small machines will have so much surface area for their weight that they can fly like insects, beating their cilia in patterns optimal for moving air.

The contribution of the bigger branches to the power and intelligence of the smaller ones can be visualized as a kind of reverse pyramid scheme. Each level (counting all the twigs at that level) of our robot has twice the volume of the next smaller level and thus room for twice the power supply and twice the computer. Two levels down, the ratio is four to one, and at three levels it is eight. If the control and power for each level are piped up from the branches three steps closer to the stem, the small branches can be four times as vigorous as otherwise. Only the stem and the first few branches radiating from it would be shortchanged. Since most of their power and attention is directed higher up in the tree, they may be incapable of much motion and may be relegated to providing a stable framework if they passively lock their joints.

A big question is how the control programs for such a beast would work. In the extreme case one could imagine a program that would try to determine the combination of actions of each individual joint that would best accomplish the desired task. This is almost certainly an example of a colossal NP (nondeterministic polynomial) problem that

can be optimally solved only by essentially examining every possible combination of motions and picking the best (see Appendix 3). Such solutions are intractable these days even for simple manipulators that have only a handful (!) of fingers. Though computers will be vastly more powerful in the future, the problem posed by a system with a large number of fingers is much, much bigger still. From time to time an especially clever strategy for coordinating thousands or millions of fingers to accomplish a particular task may be discovered, and collections of clever strategies will be passed between individuals—the manual skills of the superintelligent era. But while humans teach skills as simple as how to tie a shoelace, the lessons of superintelligent machines may be more comparable to instructions for assembling an airliner. Finding the best possible solution for routine problems will usually be out of the question, but finding a good enough one may not be too hard. I imagine a divide-and-conquer strategy, where the stem considers the overall problem and generates plausible subtasks for each of the four subtrees immediately connected to itself. These further subdivide the problem and pass those fragments on, and so on. The smallest branches would receive simple commands like move to a certain position, or move until resistance is met. A command from the stem might be something like: *North bush—stay on left side of plane A, and right side of plane B, and apply net force vector V to object; East bush—stay on right of A and B, and resist any motion more than 10 cm from the axis; South bush—right of A, left of B, apply force negative V; West bush—left of A and B, and resist.* If a subproblem, as passed to a small bush, cannot be solved, a complaint would be sent back to the originating branch, which would then go back to the drawing board to try something else.

A bush robot would be a marvel of surrealism to behold. Despite its structural resemblance to many living things, it would be unlike anything yet seen on earth. Its great intelligence, superb coordination, astronomical speed, and enormous sensitivity to its environment would enable it to constantly do something surprising, at the same time maintaining a perpetual gracefulness. Two-legged animals have three or four effective gaits; four-legged animals have a few more. Two-handed humans have two or three ways to hold an object. A trillion-limbed device, with a brain to match, is an entirely different order of being. Add to this the ability to fragment into a cloud of coordinated tiny fliers, and the laws of physics will seem to melt in

the face of intention and will. As with no magician that ever was, impossible things will simply *happen* around a robot bush. Imagine inhabiting such a body.