

Damien Broderick

Terrible Angels

The Singularity and Science Fiction

In the twenty-first century, we are ghosts haunting a futurist dream a century old. Reminded of the brash functionalist lines of ‘tomorrow’s city’, the clean Mondrian Formica kitchen surfaces, the sparkling electric railroads, we find that we accelerated by mistake into Gothic mode: Ground Zero at the heart of the world capital of capital, jittery crack zombies in the skyscrapers’ shadows, the broken power plants of Iraq, the desolation of much of Africa, the last Space Shuttle hauled off to a museum, global economies a hairsbreadth from disaster. And yet this dull headache reality still might be just a momentary blip in a curve of change that is even now continuing to remake the future, fashioning it into forms we cannot even begin to imagine until we’re smacked behind the ear.

What if (the basic science fictional opening move) there’ll be *no* human history, no ensemble of possible histories to be modeled by humanist fiction or futurism? What if the third millennium sees a phase change into some altogether unprecedented state of being, neither Apocalypse nor Parousia but some secular mix of both? What if, as Vernor Vinge proposed, exponentially accelerating science and technology are rushing us into a Singularity (Vinge, 1986; 1993), what I have called the Spike? Technological time will be neither an arrow nor a cycle (in Stephen Jay Gould’s phrase), but a series of upwardly accelerating logistical S-curves, each supplanting the one before it as it flattens out. Then there’s no pattern of reasoned expectation to be mapped, no knowable Chernobyl or Fukushima Daiichi to deplore in advance. Merely — opacity.

Correspondence:
Email: thespike@satx.rr.com

Journal of Consciousness Studies, 19, No. 1–2, 2012, pp. 20–41

Copyright (c) Imprint Academic 2011
For personal use only -- not for reproduction

The general idea of exponential accelerating change goes back further than that, though. In *The Spike* (Broderick, 1997/2001), the first book-length consideration of the technological Singularity, I cited an article by G. Harry Stine, published in the sf magazine *Analog* in 1961.¹ Yet Stine's insight had already been canvassed in February, 1952, in another sf magazine, *Galaxy*, by Robert A. Heinlein. His essay 'Where To?' (1952/1980) not only described the path toward Singularity (without calling it that; he dubbed it 'the curve of human achievement'), but even displayed a graph showing four possible trend lines stretching from the upward curve of change between 1900 to 1950, extended conjecturally toward 2000. Trend 1 quickly slowed and went flat-line. Trend 2 also slowed, without quite turning stationary. The other two kept going upward:

What is the correct way to project this curve into the future? Despite everything, there is a stubborn 'common sense' tendency to project it along dotted line number one.... Even those who don't expect a slowing up at once tend to expect us to reach a point of diminishing returns (dotted line number two).

Very daring minds are willing to predict that we will continue our present rate of progress (dotted line number three — a tangent).

But the proper way to project the curve is dotted line number four — for there is no reason, mathematical, scientific, or historical, to expect that curve to flatten out, or to reach a point of diminishing returns, or simply to go on as a tangent. The correct projection, by all the facts known today, is for the curve to go on up indefinitely with *increasing* steepness.

Sf writers are rarely silenced by this futurological *mise en abîme*. At the very least, tales can be told of the runaway ride up the slope of the Singularity. By reaching into myth and the most expansive vistas of advanced physics, neuroscience, biology and even cosmology, perhaps we can guess at some of the immense prospects that loom beyond the veil. Runaway change can be represented as a singularity because it is a spike on the graph not just of human progress but of human reality in its entirety. The strangest feature of such a graph, taken literally, is that the higher you rise on its curve, the faster it climbs ahead of you. The slope is worse than Sisyphean, because we can't even get to the top and then slide despairingly back to the base.

[1] G. Harry Stine, 'Science Fiction is Too Conservative' (1961); this was followed by G. Harry Stine, 'Science Fiction is Still Too Conservative!' (1985). In 1996, I asked Stine for his current assessment. Without commenting on his wildly optimistic and erroneous earlier projections, he replied gamely: 'Science fiction is STILL too conservative!' (Personal communication).

‘As we move closer to this point, it will loom vaster and vaster over human affairs till the notion becomes a commonplace,’ Vinge (1993) pointed out. ‘Yet when it finally happens it may still be a great surprise and a greater unknown.’ Vinge is often denied due priority in his pioneering attack on this problematic, usually (as Professor Chalmers does, while acknowledging John W. Campbell’s visionary sf from the thirties) by citing a briefly described conversation between John von Neumann and Stanislaw Ulam (1958) and I. F. Good’s ‘intelligence explosion’ (1965) leading to cascading self-bootstrapping improvement. As a science fiction writer and futurist, I find it pleasing (and dazzling!) that Heinlein was there, sketching the Spike, some 60 years ago.

Sf’s enterprise is both quixotic and impossible. It is — to invoke the inevitable tedious religious comparisons — akin to the futility of a theologian or a physicist attempting to understand ‘the Mind of God’ (as atheist Stephen Hawking rhetorically and misleadingly dubbed his own scientific efforts). Such utter dislocation of knowledge and feeling alike is the traditional literary figure for the incomprehensibly sublime.

Perhaps a less odious comparison, to both believers and nonbelievers (although regrettably bearing its own kitsch New Age freight) is a Rilkean Angel. Rainer Maria Rilke, in the *Duino Elegies*, makes our blood shudder with his vision of the beings (although he did not know this) that, in all truth, might stand beyond the veil of the Singularity. Pressed against the heart of such a being, its beauty inseparable from the terror it instils in us, we would simply perish, for, in Rilke’s great and tolling words, ‘Every Angel is terrible’ (Rilke, 1989). The Singularity, for Vinge, Michael Swanwick, Iain M. Banks, Charles Stross and other sf pioneers, usually begins with artificial intelligence: rogue, benign, or simply — *there*. Evolution takes such a *long*, agonizing time to ratchet itself up a notch. Break free from that mindless DNA process, via machine self-programming or genetic engineering or both, and everything changes *fast*.

How fast? Vinge, in nonfiction mode, is uncompromising: ‘The precipitating event will likely be unexpected — perhaps even to the researchers involved. (“But all our previous models were catatonic! We were just tweaking some parameters...”) If networking is widespread enough (into ubiquitous embedded systems), it may seem as if our artefacts as a whole had suddenly wakened’ (Vinge, 1993). After that, after the first machines have awoken — or the global Internet, as occurs in Robert Sawyer’s *WWW* trilogy (Sawyer, 2009–11)? Nobody knows. ‘For all my rampant technological optimism,’ Vinge himself

states, ‘sometimes I think I’d be more comfortable if I were regarding these transcendental events from one thousand years remove... instead of twenty’ (Vinge, 1993). While it’s a millennium too late for that, in his influential novel *Marooned in Real Time* (1986) Vinge placed his characters at a 50 million year remove — in the remote future, after a Singularity had passed across the world, leaving only traces of vast, unintelligible engineering works and decaying cities.

The Singularity will be a sort of immanent transcendence, an accelerating dash into incomprehensible glory — or ugliness, always a possibility: feral nanotechnological gray goo, turning everything on the planet into sludge, or viral green goo, or a thousand varieties of unbalanced superminds let loose in the playground of the solar system. Here and there in Vinge’s novel are hints of what might have occurred when humanity vanished clean away while the reader’s back was turned. ‘Mankind simply graduated, and you and I and the rest missed graduation night,’ one character tells another. ‘Just talking about superhuman intelligence gets us into something like religion...’ (*Marooned*, p. 111). One need not be earnest about this prospect. Iain M. Banks’s Culture novels are peopled by posthumans rewired to ‘gland’ hormones by choice, whose postscarcity anarcho-communist polity is mostly located on starships and Orbitals run by AI Minds who seem to keep their organic partners around as pets.² Colossal artifacts with facetious names like the *Very Little Gravitas Indeed* roar across the galaxy, while enhanced humans and snide machines frolic within their protective fields.

How fast could such immense changes really happen? A marooned Vinge character muses doubtfully that even in a swiftly changing world, ‘there had been limits on how fast the marketplace could absorb new developments... what about the installed base of older equipment? What about compatibility with devices not yet upgraded? How could the world of real products be turned inside out in such a short time?’ (*Marooned*, p. 172). Vinge seems here to start interrogating his own cool idea only to back away hurriedly. It is a piece of narrative flimflam in the guise of rhetorical questioning.

Yet some of the steps that make this headlong alteration thinkable, if not altogether feasible, are sketched in Vinge’s seminal novel. ‘High-tech’ people from close to the Singularity wear headbands that augment their native abilities, computer patches to the raw stuff of evolved brains with their limited memories and even more limited attention windows. More up-to-date scenarios by Greg Bear, Greg

[2] See especially the feuding Minds in Iain M. Banks, *Excession* (1996).

Egan, Raphael Carter, Charles Stross and others expect such chips to be surgically implanted deep in the brain, or perhaps grown there using engineered cells or nanoconstructors, or just bypassed when humans upload entirely into computer substrates. In a way, it's evidence of Vinge's own argument that only a quarter century after his proposition was announced, we already find many of his once-wild projections rather tame and unadventurous.

Today, even the smartest of us can hardly sustain conceptual grasp on more than seven different items at once, and five is the usual mental handful. So we live in a cognitive universe dimly glimpsed through the narrowest of cracks, and the width of that aperture is set by our inherited neural hardware. Boost it, link our augmented minds together, and who knows what wonders of awareness might burst open into consciousness? 'Humankind and its machines became something better,' speculates the most advanced of Vinge's pre-Transcension high-techs, 'something... unknowable' (*Marooned*, p. 176).

Yet that character's own experience at the opening of the twenty-third century is almost incomprehensible to us a mere two hundred years earlier. With his seven colleagues, he was engaged in mining the sun for antimatter, 'distilling one hundred thousand *tonnes* of matter and antimatter every *second*. That was enough to dim the sun, though we arranged things so the effect wasn't perceptible from the ecliptic' — the orbital region around the sun that contains the earth and other planets and most asteroids. Working so far from home, he and his companions were brutally severed from the real action, 'hundreds of light-seconds away' (*ibid.*, p. 173, italics above in original).

What might such a world be like if it truly existed? Two hundred years ago, Europeans explored and conquered large parts of a single world many months distant from their political and mercantile masters at home. Today, by contrast, nobody with access to a cellphone is more than fractions of a minute from anywhere else on the globe.³ Computer networks, swapping information and financial transactions, blur into a haze of virtual instantaneity. Imagine the rapidity of a world where people exiled halfway across the solar system are

[3] 'Half of Africa's one billion population has a mobile phone — and not just for talking. The power of telephony is forging a new enterprise culture, from banking to agriculture to healthcare' (<http://www.guardian.co.uk/technology/2011/jul/24/mobile-phones-africa-microfinance-farming>).

unalterably amputated from the current action, because the speed of light is an intractable barrier to faster communication.

On Earth, large corporations with superior computers merge their staff into linkages of thousands. Is this a horrible prospect of soul-death, extermination of the self? One might expect such an interpretation from a libertarian like Vinge, but in fact he suggests otherwise: ‘There was power and knowledge and joy in those companies...’ As the Vingean Singularity approaches, the mind-to-mind linkages and augmentations become extreme, a form of group mind: ‘By the beginning of the twentythird, there were three *billion* people in the Earth/Luna volume. Three billion people and corresponding processing power — all less than three light-seconds apart.’ Inevitably, transhuman plans turn to cosmic engineering, the creation of black hole star-gates (as occurs in Ken MacLeod’s *The Cassini Division* and Stephen Baxter’s *Timelike Infinity*), the implosion of entire stars. Might this explain the Extinction? ‘We’d been “uppity cockroaches” — and the real owners finally stepped on us...’ (*Marooned*, p. 177).

Vinge’s subsequent deep space novels, *A Fire Upon the Deep* (1991) and its belated prequels *A Deepness in the Sky* (1999) and *Children of the Sky* (2011), ornate, multileveled machines of tales within tales, offered a series of complex analogies for the states of mind and body and existence that might prevail beyond any immediate Singularity. For story-telling purposes, these novels propose a galaxy partitioned like some Dantean cosmology, from a bleak Inferno of the Unthinking Depths, where matter dominates and light-speed is the final limit, through a Purgatorio of the Slow Zone flaring outward from the galaxy’s spinning wheel, and farther yet into the Beyond, where minds are vast and instantaneous: Powers, ‘gods’ in the ancient sense.

That’s metaphor, but certainly some form of tiered analysis is needed if we are to build models of enhanced intelligence. It is not that human mental life is different in kind from the progressively more limited kinds found in dogs and cats and parrots and goldfish and slugs and sunflowers and bacteria. Of course, many people argue that our minds *are* exactly differentiated from those of ‘lesser creatures’ in some mystical, unapproachable way. Even supposing that it turns out to be true — that people have an extra gadget — let’s call them *souls* — that provide our self-awareness and radical creativity — that conclusion would need to be attained by long, scrupulous investigation. It would be, in a sense, a negative hypothesis: something we would be forced to entertain solely because of the failure of more rational models compatible with everything else we know about the physical

universe. (And for all our ignorance, we now do know a *colossal* amount, most of it tied together into patterns not easily edited or expanded without shredding the lot.)

Posit, then, that human minds are not *mysterious things* but are rather *very complicated processes* going on inside living brains connected by senses to the outside world. Is it reasonable to guess that doubling the power of a laptop computer, and doubling it again, and redoubling for another fifteen or thirty years, will automatically produce a smart machine, then an intelligent machine, and finally — a *hyperintelligent machine*?

Vinge poses this rather neatly: ‘Imagine running a dog mind at very high speed. Would a thousand years of doggy living add up to any human insight? (Now if the dog mind were cleverly rewired and *then* run at high speed, we might see something different...’ (Vinge, 1993). No computer is currently anywhere near the doggy level of cleverness, of course, but some of them approach insect intelligence. Might running an insect brain very, very fast turn it into someone you’d like to discuss politics or art with?

It seems obvious that the answer is no, but what if that insect brain joins a swarm of its fellows, and they learn (or are taught) the trick of specialization, so that one bunch of insect brains swaps information about how stuff *looks* today, from a number of strategic angles, and another lot focuses on *sound vibrations*, while other bunches of ant-brains store these impression and disgorge them on command from yet other groups. This is the principle of the hive. In many ways the hive is much cannier than any of the dumb modules with legs and wings that comprise it. It also looks something like a complex brain — a *human* brain, in fact. Each of our neurons is a little like an insect (simpler, in fact, and much smaller), but it has better lines of communication than any single ant, and neural columns are like hives. Take a hundred thousand million neurons and let them cross-wire to each other in some kind of self-organizing, autocatalytic hierarchy, and you do indeed get — *us*.

Which does not answer the original question about speeding up a doggy brain. No, even if you added in a big pack of extra memory, a dog’s brain lacks the architecture to think like a person. What about an immortal and massively augmented dog, genetically engineered to grow as many cortical neurons as a human, but arrayed as they would otherwise be in a normal dog? In his near-Singularity novel *Holy Fire* (1996), Bruce Sterling has one of these wonderful beasts, a poignantly articulate talkshow host which/who, due to neurological deficits, can’t read. I suppose it is remotely possible that an enhanced dog

might, perhaps with great inner agony, teach itself the tricks of reflective self-consciousness. It would need access to a brainy, communicative culture, and right now we are the only one available. So a true hyperdog would require eyes capable of reading, and dexterous paws, and a jaw and larynx rewired for speech — a big job. Almost certainly no such wonder could emerge from a single, extreme mutation that just pumped up the puppy's brain capacity in utero (the conceit in Olaf Stapledon's superb *Sirius* [1944].)

It is much easier to reallocate machine memory, to rewrite code and try it out and discard it when it fails, and try again, and keep trying until you get it right, step by evolved step. That is presumably true even if the machine is boot-strapping its own abilities, driven away from stolid stable inner states by random changes somewhere in its operating system, in the software that makes it a process rather than just a lump of expensive silicon. This is the evolutionary model of teaching a machine to be a person. In a sense, we know that this has to work, because on a geological timescale it is what produced us. Mutation, contest and cooperation, differential survival of genetic patterns according to the success of the bodies they built: natural selection, in a word.

That is the likely path of AI or artificial intelligence. Vinge mentions a quite different method of attaining advanced cognitive abilities: a switch in emphasis that is captured deftly by a switch in the name — IA, or *intelligence amplification*. Now it's the human brain that is being boosted, or linked to others of its kind (as the near-Singularity corporations in his novel were composed of group minds). Again, this is eerily familiar, precisely because it is a definition of a society or a culture.

Already our brains store much of their knowledge outside the skull, in books and film and digital, searchable cloud archives, in the huge hard structures of cities and aqueducts and farms, in the facts and opinions we share by talking to each other in person and electronically. Vinge sees the path from this existing state of affairs to enhanced intelligence as rather easier than the AI route, because we have done it a number of times without quite realizing it. Currently, access to the Internet is a prodigious augmentation of a researcher's capacity: In effect we now think faster, put information together more swiftly, send our results out without delay, gain the reciprocal benefit from others cruising the net. Once improved methods are devised for getting information out of a database and shaping knowledge without the need for writing or keyboarding or even talking — something like the cyberpunk dream of 'jacking in to cyberspace' — it will be a

whole new kind of life. One small step for the individual cybersurfer, one mighty swarming leap toward the Singularity.

The key point about most Singularity scenarios is that they proceed step by step, advances of equal magnitude each taking only about half as long as the one before, so the most radical changes happen at the end of any given period. This exponential progression was sketched in Marc Stiegler's *Analog* story *The Gentle Seduction* (1989), where a man and woman from today experience the ramp up to a Singularity through rejuvenation, headband connectivity, implants, terraforming of Mars, nanoscale transportation to other stars, vastening of consciousness, contact with aliens; and yet:

She remembered who she had been when she was 25; she remembered who she had been when she was just 10. Amusingly, she also remembered how, at 25, she had erroneously remembered her thoughts of age 10. The changes she had gone through during those 15 years of dusty antiquity were vast, perhaps as vast as all the changes she had accepted in the millennia thereafter. Certainly, considering the scales involved, she had as much right today to think of herself as the same person as she had had then. Expanded communion would not destroy her; she was her own bubble no matter how frothy the ocean might become.

Vinge, then, was not alone among imaginative writers in trying his hand at portraying what is strictly inconceivable, a world on the far side of the Singularity. In the 1940s and '50s, Isaac Asimov almost single-handedly reinvented the idea of artificial cognition in the form of 'positronic robots'. The platinum-iridium brains of these machines were imprinted with Three Laws that forbade them, on pain of incapacitation, from injuring a human or allowing harm to come to one of us, while requiring them to obey humans and protect themselves.⁴ Most of Asimov's tales were clever attempts to surmount the barriers presented by these algorithmic 'laws' while showing that finally they were impregnable. But Asimov's own logic took him into the territory of a self-improving machine mind, and indeed several late stories in his collection *I, Robot* reveal the evolution of the Machines, AIs of immense power and insight, that first take over the running of the world and then remove themselves, knowing that their continued benevolent presence would blight our species (and, of course, block any future robot stories). Asimov's own favourite robot tale is 'The Last Question' (1956), in which people throughout future history ask increasingly ubiquitous and puissant AIs if entropy can be reversed,

[4] See http://en.wikipedia.org/wiki/Three_Laws_of_Robotics

or must the cosmos go inevitably into heat death? Trillions of years hence the final AI, known as AC, deduces the method, and says, 'LET THERE BE LIGHT!' But note: this heroically optimistic parable is precisely *not* about a singularity — it's a very far future Omega Point emergent god story.

Perhaps the most notable sf stories at the boundaries of the Singularity were by the late Theodore Sturgeon, a poet, dreamer and storyteller currently somewhat in eclipse. In the 1950s, his short fiction returned again and again to territory we now recognize, while drawing much of its power from his immersion in fallible characters carefully observed. Often — strangely enough, for an American writing at the height of McCarthyite Red Menace hysteria — his wistful parables told of maimed, lonely, isolated men and women finding one another in the company of a greater unity, a colony mind, a hive, *homo gestalt*. Yet Sturgeon did not abdicate from the precise, piercing joys of quirky individuality, or the prickly demands of freedom. His fiction may yet prove to come closest to one tendency in the posthuman condition.

In 'The Skills of Xanadu' (1956), a primitive world is found by a hightech militarist. The naked savages wear a curious belt but little else, and inexplicably their disorganized play fetches them everything they require in life. This harmony, maddening to the rigid newcomer, proves to be a by-product of their single crucial item of technology: the belt, which links all brains into a shared knowledge base. Faced with a problem, you simply *understand* the answer. The skill of the best practitioner of any craft is instantly available. When the warmonger leaves, taking a belt, he hopes to use it to transform his rigid society into a perfect dictatorship. Instead, the belts convert his world into a harmonious network of free individuals, each of them experts in the skills of 'logic and love; sympathy, empathy, forbearance' (Sturgeon, 1979, p. 259). Years later, Sturgeon wrote wistfully: 'I yearn to live on Xanadu, and wear their garment, and join with them in their marvelous life-style' (*ibid.*, p. 228).

William Gibson, coiner and explorer of cyberspace in novels such as *Neuromancer* (1984), is only the best known of the next generation of writers working at the lip of the Singularity, along with his sometime collaborator Bruce Sterling. Michael Swanwick, a brilliant stylist, has done so several times, most notably in the novels *Vacuum Flowers* (1988) and *Stations of the Tide* (1991). In the former, revived suicide Rebel Elizabeth Mudlark proves to be a kind of antidote to a

postmodern, poststructuralist affliction: the dispersion of the self in a culture where personalities are optional plug-ins, bought off the shelf as we buy computer apps. Rebel's capital importance is her possession of — indeed, her constitution in — the rare trick of assuring her own hyperstable self-presence. All around her, post-Singularity space habitats are gladly giving in to centripetal attractions. Earth's home population has already merged into a multibillion-headed colony organism, a sort of slime mold or coral comprised of former individuals. Indeed, this new shared-consciousness entity calls itself the Comprise.

Rebel's partial solution to this encroaching monolith is to disperse the microworlds as seed pods among the stars — a scattering that enacts on a literal level the dispersion of selfhood that the late twentieth century agonizingly diagnosed as the epoch's characteristic feature. When Swanwick names the armless girl who speaks for the colony mind — 'the focus.... of perhaps a billion Comprise, as massive a point source of attention as Earth ever needed to assemble' — he calls her 'Snow' (Swanwick, 1988, p. 222). The name reflects what every child learns at an early age: that every single snowflake is identically simple, with its sixfold symmetry, yet unique in the details of its shape. Swanwick's 1987 naming recalls, in turn, an earlier fictional group mind, in a 1959 novel, *Wolfbane*, by Frederik Pohl and Cyril Kornbluth: the Snowflake. Apparently these apprehensions have been curdling through the discursive unconscious of Western culture. It was Vinge's special genius to understand that they might become *literally* true.

And if they do? Swanwick is not reassuring. Starting from primitive AIs and rudimentary brain-computer interfaces, he posits a seed of 'thirty-two outlaw programmers' about whom the Comprise, a massively parallel human-machine hybrid, starts to crystallize. Born in glory, power and understanding, the new entity reaches to others in the net willing to join it, rewriting its own structure as it goes, deepening its algorithms. 'Within three minutes everyone on the net was ours. We controlled everything that touched upon the net — governments, military forces from the strategic level down.... With a fraction of our attention, we designed the transceivers, retooled the factories to make them, and reorganized the hospitals to perform the implants.... We ate the Earth' (Swanwick, *Vacuum Flowers*, p. 224). It is a totalitarian longing and one that today's culture fears desperately even as we yearn for the balm of that sweet joining. 'We reached out and out,' the Comprise tells Rebel, 'expanding toward Godhood.' It adds at once, candidly, 'We had ambition, and ascended into Hell' (*ibid.*, p. 224).

If this is a disturbing prophecy of a hungry post-Singularity consciousness, Swanwick was more terrifying still in *Stations of the Tide*, where nothing holds still for long (see Swanwick, 1991, p. 146). People split their minds into agents able to impersonate them, act on their behalf, report back and extinguish themselves. They move their point of awareness into surrogate bodies, the very reverse of Rebel's armored solidity of selfhood. When an agent of colonymind Earth is met, it is something out of Milton and Swift, an authentically monstrous manifestation in virtual reality. This vast, sweat-stinking, musky monster is a figure familiar from psychoanalysis: the archaic Mother, a sort of feral female phallic force, more mythic than misogynistic in Swanwick's making. And like that clammy image from post-Freudian analysis, complete with vagina dentata, it invites Swanwick's bureaucrat into its mouth. In the overwhelming presence of an Earth utterly overborne by technology out of control, he asks the agent: 'What do you want from us?'

In that same lifeless tone she replied, 'What does any mother want from her daughters? I want to help you. I want to give you advice. I want to reshape you into my own image. I want to lead your lives, eat your flesh, grind your corpses, and gnaw the bones' (Swanwick, *Stations of the Tide*, p. 147).

Here and now, we can readily fall in with Swanwick's forebodings, considering the prospects of a world that is perhaps only a generation away from birthing nanotechnology and artificial intelligence and machine interfaces and human intelligence amplification. It need not be that way, of course. One might wonder if it is even remotely plausible. Grimm's fairy tales, after all, caught the rural voices of the nineteenth century and earlier, perhaps not altogether salient to a world remade by science. Yet we carry our history with us, tucked away inside our narratives and nightmares, ready to be snatched up in the slipstream of science.

Will we, preserved by the wiles of that science into the Singularity (or, if not us, then our transhuman children or great-grandchildren) any longer be 'human'? Perhaps not, or not for long. Maybe we will live almost infinitely accelerated lives within a virtual computer in a grain of sand at the edge of the world's last drained sea. Maybe we will be quantum states of a cosmically dispersed, quantally-linked hypermind. Maybe we will be quite literally *gods*, inflating fresh universes out of the quantum foam and placing our impress upon everything that forms there.

Or maybe we'll all stay at home and watch the ultimate television channel, forever — my favorite explanation for the Fermi Paradox. This is the Great Silence in the skies, the mysterious absence of any radio or optical traffic between worlds out there in the galaxy and beyond, or of any cosmic rearchitecting of the stars. If the technology and culture of every civilization Spikes within a century or two of their discovering radio and space flight, they might all be tucked away into the folds of local spacetime like the hidden, rolled-up dimensions of string theory. They might even be living *there*, colonizing those intricate, implicate spaces.

Must the Spike, the Singularity, the amplification of human intelligence, or its replacement by hypermachines, be seen as either a farce or a horror story? No doubt this is the way it will be portrayed by Hollywood (Skynet Terminators!), embodying the usual false binary opposition. Mind versus Passion! Love throttled by Rationality! In *Charly* (1968), based on Daniel Keyes's heartbreaking story 'Flowers for Algernon', surgically enhanced intelligence lifts a sweet moron to the heights of sensitive genius before smashing into Nature's Revenge, dashing the tragic hero back into stupidity and early death. In the fifties' cult movie *Forbidden Planet* (1956) a device for increasing IQ and linking minds into a planetary gestalt frees the Monster from the Id, those vile forces boiling destructively in the unconscious (in pop Freudianism), and kills an entire world. Star Trek's Mr. Spock and Lt. Commander Data did a little better, but paid for their cleverness in the inevitable coin: They were forbidden human depth and feeling.

This is a very strange cliché, its supposed truth is denied by the slightest contact with real, passionate, sometimes frenzied human scientists and technologists. But perhaps this is what *really* frightens us: when strong feelings, devotion and hatred and prejudice, are joined powerfully with effective strategies for influencing, even dominating, the world. Then we get holocausts, the continuing risk of nuclear war that could serve nobody's purposes, mad cults and faiths with real weapons.

Is the Singularity likely to bring just such unholy conjunctions to a new pitch of fearful strength? We have evolved in hundreds of thousands of years of sluggish tribal life, dealing with each other face to face, vulnerable to the fists and scorn of others, sensitive to their smiles and touch. What happens when we allow ourselves to deal with others wholly through a monitor's or iPhone's window? Anyone who

has watched, or taken part in, a ‘flame war’ on the Internet knows how these exchanges escalate into brutal rudeness and sarcasm, the kind of thing that leads on the highway to ‘road rage’ murders. It is a stroke of luck that we cannot yet easily kill our foes through the screen. But perhaps an enhanced human, locked into symbiosis with potent programs running on computers all around the world, would indeed learn how to smite enemies in world-shaking tantrums.

With luck, this interregnum would be brief. By and large, people are not wild dogs. It has been possible for decades to make lethal bombs from common agricultural substances, or gun down children in a playground, yet when this is done by a few crazies the rest of us are genuinely grief-stricken, if only for a brief time before our own concerns drift back to the surface. The transition from intelligence amplification (IA) to the Singularity, when enormous numbers of changes happen with immense swiftness, will probably not occur as an all-at-once crystalization — or at least we must hope not.

Even so, in John Barnes’s *Mother of Storms* (1994), a Singularity occurs when Louie Tynan (one of the first humans on Mars, now in Earth orbit) and Carla Tynan (on her submersible yacht) quite plausibly get absorbed into the net in 2028, after a global catastrophe, then spread between orbit and the Moon and eventually into a comet from the Kuiper belt. Such transition into dispersed superintelligence might bring with it the solutions to most of our traditional woes: hunger, thirst, nakedness to the elements — scarcity, in a word. The best means to those solutions cannot yet even be glimpsed, but the likely paths are obvious: nanotechnology, direct AI-interfaces, genome control and repair, massively extended lifespan, and so forth.

This might seem a utopian outcome — but really, it can be neither utopic (positive) nor dystopic (negative). From what standpoint is ‘positive’ to be assessed? *Mother of Storms* seems highly positive, in that the uploaded and enhanced humans become a sort of benign solution to all our problems. But seeking a future positive from *our* viewpoint is arguably like asking a trilobite or a *Homo habilis* to regard the prospect of the emergence of *Homo sapiens* as a happy ending. ‘Yes, we’ll all die out, but another species will do pretty well.’ Charles Stross’s *Accelerando* (which we’ll examine shortly) is undeniably exciting, taking us in explosive leaps from the near future to a madly different post-Singularity medium-term future full of complexities we can scarcely understand — and I find that positive, compared to doom and gloom of Cormac McCarthy’s *The Road* or Margaret Atwood’s *The Handmaid’s Tale* or Kazuo Ishiguro’s *Never Let Me Go*, which is how most ‘literary’ fiction still wants to portray the future.

‘Suppose we could tailor the Singularity,’ Vinge proposes. ‘Suppose we could attain our most extravagant hopes. What then would we ask for: that humans themselves would become their own successors, that whatever injustice occurs would be tempered by our knowledge of our roots’ (Vinge, 1993).

In this authentic golden age, immortality would no longer be a fantasy of consolation in the imagined world beyond death, but a literal indefinite extension of life in a utopian world without want.⁵ A caution is necessary. We are aware of our limitations, the weaknesses of brains and bodies built by blind Darwinian selection. Each of us expires a few decades at most after a fertile span of some thirty or forty years. We are put together as disposable gene-carriers. Our beautiful minds rot with our fallible, corruptible brains and are gone forever. What would happen with such temporary mechanisms if they were repaired again and again, held safe from the corrosion of time? Nightmare, perhaps. ‘A mind that stays at the same capacity cannot live forever,’ Vinge notes; ‘after a few thousand years it would look more like a repeating tape loop than a person’ (Vinge, 1993). So extended life must also be enhanced life. For immortality we would need to be smarter and with better emotional control, even if we didn’t already need amplified gifts to attain immortality. (In fact, we might not, as medical consequences of the Human Genome Project are gradually pointing the way to endless cellular repair.)

More to the point, what is the ‘we,’ the ‘I,’ that is going to survive into the post-Singularity utopia? If poststructural theory tells us that the self is always-already a construct, an illusion, that improbable perspective will be even harder to dispute in a world where we can send out ‘partials’ or ‘agents’ from ourselves into the global net. Already, our brains are composed of dedicated, somewhat partitioned modules resembling the ‘faculties’ of an older philosophy. Of course these specialized components tend to work together, creating (most of the time) a sense of unified consciousness. Once we learn to split off fractions of our selves — or, rather, duplicate and amplify and elaborate those fractions, as in the Swanwick novels — we will no longer be strictly human. Vinge notes: ‘These are essential features of strong superhumanity and the Singularity. Thinking about them, one begins to feel how essentially strange and different the Post-Human era will be —

[5] An elaborate Golden Age of this kind is projected in John C. Wright’s trilogy of the same name (2002–03).

no matter how cleverly and benignly it is brought to be' (Vinge, 1993, italics in original).

An audacious Marxist (or perhaps post-Marxist) version of a world going through, recovering from and absorbing a Singularity has been sketched by Ken MacLeod, like Iain M. Banks a Scots writer with affinities to the anarchist left. These novels began with *The Star Fraction* (1995), which nearly won the Arthur C. Clarke award, a kind of frenzied Trotskyist vision of a world convulsed by political fractions and manipulated by an emergent AI weapons system. Subsequent novels explore a variety of future cultures that might emerge from this abrupt discontinuity. *The Sky Road* (1999) actually revises the future history of the other three, ensuring that a menacing Singularity does *not* occur, or at least is delayed. *The Stone Canal* (1996) and its sequel *The Cassini Division* (1998) portray with high, sly humor and impressive technical and political insight two contesting utopias, one a Banksian anarcho-socialist Union on a damaged Earth, the other an anarcho-capitalist libertarian world on the far side of a wormhole, 10,000 light years from Earth and in its future.

Both alternative regimes (or antiregimes) are thus at arm's-length yet able to communicate and even visit. Jupiter, meanwhile, has been redesigned by posthuman entities, the 'fast folk,' that blend organic roots and AI enhancements. Are they dreadfully dangerous — trapped in psychotic virtual realities, but likely to emerge at any moment — or should they be welcomed as our successors? MacLeod is superbly sardonic in setting all these groups at each other's throats and minds. His vision of a Singularity is distinctly unnerving, and his several contrasted utopias are no less troubling even as they seduce us in succession.

The grace to permit remnants of the old, unreconstructed humanity to live in peace may be the best we can hope for from augmented or borganized posthumans. Well, one happier option is conceivable, for those open to the charms of self-deceptions: Our successors might choose to disguise themselves as our servants. Vinge hints at 'benign treatment (perhaps even giving the stay-behinds the appearance of being masters of godlike slaves)' (Vinge, 1993). Unaltered humans would be the metaphysical equivalents of the Amish, serene agricultural throwbacks in our mechanized, electronic world. That appears to be a difficult way of life to sustain, even to negotiate, and its future version would surely be harder still to endure without soul-wrenching episodes of blatant bad faith, lack of resolution when times grew tough and one's body aged while everyone else remained healthy and young, backsliding by the next generation.

When the first Singularity consciousness awakens and looks around itself, it might not be a safe time to be alive; it might be the end of all histories that concern us. Or the AI might simply shut itself down, or vanish into ‘uncompromising silence’, as in Stanislaw Lem’s superb parable ‘Golem XIV’ (1981/1985).

So what should be our attitude to these posthumans who seem likely to replace us, perhaps within a century or even much sooner? Terror? Awed worship? Serene acceptance? The hope that we might join them in a *gestalt* superhuman Comprise?

Many alternative paths lead from where we now stand into the transhuman and the posthuman futures of the Singularity. It is not at all clear that there’s any role at all for us mere humans on the far side of the Singularity’s looming wall, unless we are drastically enhanced, or perhaps uploaded into virtual spaces where our electronic or photonic Doppelgangers can run a million times faster. It will be uncomfortably interesting to learn how sf writers further into the twenty-first century meet this extreme and continuing narrative challenge — and how everyone else deals with the history breaching reality following hot on the heels of the fiction. For the moment, let us look a little more closely at two exemplary instances of how sf can approach the blinding dazzle of the Singularity (while bearing always in mind that by the nature of the medium, sf tends to emphasize the risks and dangerous consequences of radical technological and social change, because that’s what most readers find *exciting*...)

Like Greg Egan a decade earlier, British writer Charles Stross seemed to come from nowhere and leap immediately to the top rank of science fiction and fantasy writers. His work has proved enormously popular, and justly so. With *Accelerando* (2005), Stross was sealed as the new Poet Laureate of the Vingean technological singularity, a topic he had already approached in his first published novel, *Singularity Sky* (2003). The *Accelerando* project’s five years of development (it is tempting to apply this sort of corporate language to Stross’s dense techno-speak art-artifact) yielded an early twenty-first-century counterpart to John Brunner’s compressed future shock 1969 Hugo-winner, *Stand on Zanzibar*, complete with rich idiomatic sidebars or side loads of Baedeker guidance to the non-native.

The rogue corporation rears up slightly and bunches into a fatter lump; its skin blushes red in patches. ‘Must think about this. Is your

mandatory accounting time-cycle fixed or variable term? Are self-owned corporate entities able to enter contracts?'

That's funny, and fun, as well as knowing. Some readers complained of infodumping of the most blatant kind, yet this device seems unavoidable when a torrential cascade of novelty is the very topic of a work of art. Approached with an appreciative generosity of response, these dollops of data are tight, compressed, inventive, brilliantly illuminated gems, or perhaps genomes (or memomes) that will unfold, in a prepared mind, into wondrous ecologies of image and idea. The changes implied by headlong acceleration are by definition both too immense and too subtle to be portrayed or perhaps even imagined. Stross has the audacity and, luckily for us, the imagination to come close to pulling it off.

Manfred Macx is a venture altruist ('Manfred's on the road again, making strangers rich') — Stross is not afraid to have us smile even as he jolts our preconceptions — encrusted with computer wearables and the latest wifi connectivity, affianced until recently to Pamela, a dominatrix headhunter for the IRS who tries to persuade global megacorps to cough up the tax they owe. Their venomous bond is manipulated sardonically by their robot cat, Aineko, which is being hacked and upgraded on the sly by Pamela. This nicely observed android animal — 'It sits on the hand woven rug in the middle of the hardwood floor with one hind leg sticking out at an odd angle, as if it's forgotten about it' — might be the secret narrator of the novel. Its augmentation and expansion toward the condition of a low-level demiurge mirrors the transitions of humankind and our posthuman Vile Offspring.

In a bondage scene of hilarious erotic vividness, Pamela gets herself pregnant with their daughter Amber, who will carry much of the long arc of the story to the Singularity and beyond, as human minds export themselves increasingly outside the skull into machine substrate exocortices. In turn, Amber's son Sirhan (well, son of one of her many instantiations) takes the generational saga to the destructive *Childhood's End*-style transcension of the solar system into a Matrioshka Brain (energy-hungry Dyson shells of computronium hosting untold trillions of superminds),⁶ the return from death of an extremely augmented Manfred, and a blind plunge beyond the provincial Milky Way to a realm where a galactic superintelligence seems to be mounting a 'timing channel attack on the virtual machine that's running the universe, perhaps, or an embedded simulation of an entirely different universe'.

[6] A consummation closing my own Singularity novel *Transcension* (2002).

This is how high bandwidth science fiction works. As the decades pass, as the rate of change accelerates, Stross's characters *become* Googled. And even with their inbuilt channels of information and communication, they are lost like us in the hydrant gush of available knowledge. All around them, intellectual tools are mutating into predatory lifeforms. Feral tax auditing software roams the solar system, entire economic systems convulse in ecological firestorms of contest. And then there are the aliens... which, of course, are just as likely to be autonomous spam attacks as anything we would recognize as people.

Accelerando is a *Fantasia*-bright cavalcade of borrowed and adapted landscapes — the Atomium globe from the 1950 World's Fair, the deck of the *Titanic* emulated in a virtualization stack, a phony debased Islamic heaven — transplanted to Saturn's icy atmosphere or a virtual reality world inside a soda can-sized starwhisp interstellar spacecraft or an alien router network. Does it work? *Can* it work? It is an impressive attempt upon the impossible. For all its Catherine-wheel sparkle and intellectual bravura, there is evidence that the impossible must remain always out of reach.

This very recognition framed the last major work of US writer Poul Anderson, *Genesis* (2000), which confronts a basic and almost insurmountable consequence of a post-Singularity future. Denying the Fermi Paradox, Anderson posits that intelligent consciousness, once evolved, must proliferate on a galactic scale, mutating and extending its own capacities, perhaps replacing its very substrates. It might relocate itself, for example, from limited organic bodies to very much more adaptable synthetic forms. The Science Fiction Writers of America's 1997 Grand Master, Poul Anderson (1926–2001) was familiar with extrapolations along these lines by roboticist Hans Moravec (1988; 1998), and built them gracefully into his own saga of a galaxy a billion or more years farther off into deep time.

Gaia, the vast, immanent AI custodian and consciousness of the world, finds ancestral Earth threatened by a swelling, terminal Sun. Rather disturbingly, Gaia wishes to allow the world to perish in final flame rather than disrupt the Sun's 'natural' astrophysical trajectory. Other mighty Minds throughout the galaxy, and to the 'shores of the Andromeda,' find this plan perverse. One such godlike node, Alpha, hives off a sub-mind (still Olympian by our standards), and sends this Wayfarer to Earth to investigate and intervene. A still more diminished aspect or agent of this fragment is a reconstruction of the early upload engineer Christian Brannock. A merely human-scale genius, he visits the planet as his larger self communes and debates with Gaia. What he finds, inevitably, is baffling yet emotionally moving (in its

constrained way), recalling those Norse sagas Anderson loved so well.

And all of this impossibly remote story is told to us as myth, as repeatedly distanced construct. We are informed again and again that what we read is nothing like the vast reality. Of course, this *must* be so, given the premises of ruthlessly projected futurism. ‘All is myth and metaphor, beginning with this absurd nomenclature [Alpha, Wayfarer]. Beings like these had no names. They had identities, instantly recognizable by others of their kind. They did not speak together, they did not go through discussion or explanation of any sort, they were not yet “they.” But imagine it.’

And we do, for we have been here before. This is the grand proleptic mythology of Olaf Stapledon himself, of Roger Zelazny’s ‘For a Breath I Tarry’ (1966) — in which Machine remakes Man, but then bows before Him (which is absurd and sadly farcical, however much that story was loved in the 1960s).

Anderson eases our entry to allegory via several well-formed episodes from the comparatively near future: a boy’s epiphany beneath starry heaven, in our Earth; Christian’s empathy with his robotic telefactor extension on Mercury, prelude to his own status as an uploaded and finally multiply-copied personality; English bureaucrat Laurinda Ashcroft who plans the first millennial salvation of Earth from the brute assaults of a heedless cosmos; a small, neat parable of rigid, gorgeous clan rivalries held in check and paralysis, finally, by the emerging Mind of Gaia. These are Anderson’s antinomies (and perhaps American science fiction’s): the sacred autonomy of the self, the craving for transcendence in something larger; personal responsibility, and its terrible limits in a world linked, defined, by billions of threads.

Returned to Earth, Wayfarer’s Brannock and Gaia’s Laurinda tarry in *faux*-eighteenth-century civility, driven together and apart by a series of visitations to simulated histories as dense and real and tormented and doomed as the ‘real world.’ Their own personalities are no less constructed, however rooted in some small early reality, and so the poignancy of their dilemma is the greater. But for us, knowing that we read a fiction, and snatched in a kind of postmodern gesture again and again by Poul Anderson from our comfortable readerly illusion, these figures and their worlds run the risk of all allegory: can we care? It is the great artistic problem for any form of art predicated upon utter disruption and dislocation, of which the Singularity is by its very nature perhaps the most minatory.

Religious art faced it long ago, and clad its transcendent message in parable, majestic song — and quietness, sacralized domesticity, anguish transformed at the graveside. Confronting the Singularity, reaching for these well-honed tools to give himself voice and range, perhaps Anderson succeeded as well as anyone can manage — given that the task is impossible. If he did not truly succeed, as Stross and the rest of us did not, this is hardly a fault. It takes an entire culture to sustain such mythos. Sf has only lately begun to deepen and extend the mythos, and meanwhile the world's culture turns technological run-away into jingles and plastic toys.

It is compelling to watch how the genre, the mode, of sf is responding to this immense perspective, into the pitiless depths where Poul Anderson, not long before his death, made his brave foray. Unfortunately, consumerist culture does not tarry for more than a breath — or in this case, perhaps a few decades — on a new trope. The shock of the new has expired already for the Singularity. To the extent that it had become just one more tedious way to insert blatant magic into adventure stories, this is understandable. But it might be that the genuine ontological-cum-epistemological challenge Vernor Vinge introduced, nearly three and a half decades after Heinlein had glimpsed it, is, after all, insoluble in just the ways it declares itself to be.

In the real world of publishing, attempts to scale or burrow under its wall to the future have been met, mostly, by reverting to frank mock-medieval fantasy, or degenerating into endless puerile zombie and vampire bloodbaths, franchise wars in space, and whimsical post-Goth zeppelin romances and regressive if fun steampunk sagas of a world that has turned its back on technological progress. In the meantime, in the other real world of bioengineering and computer science, Moore's law is still ticking away. Whether it will reach escape velocity and head off up the page to tear through the top of the graph will only be known for sure when we or our flesh and AI descendents are there, beholding endlessly new landscapes of change, no longer hugged terrifyingly by Rilke's Angel but gazing out from its eyes.

References

- Anderson, P. (2000) *Genesis*, New York: Tor.
 Asimov, I. (1956) The last question, *Science Fiction Quarterly*, (November).
 Banks, I.M. (1996) *Excession*, London: Orbit.
 Barnes, J. (1994) *Mother of Storms*, London: Millennium.
 Broderick, D. (1997/2001) *The Spike*, Sydney: Reed (1997); rev. New York: Tor (2001).
 Broderick, D. (2002) *Transcension*, New York: Tor.

- Heinlein, R.A. (1952) Where to?, in *Galaxy*, pp. 13–22. An extended version, with updates from 1965 and 1980, is in Heinlein, R.A. *Extended Universe*, New York: Ace Books.
- Lem, S. (1981/1985) Golem XIV, in *Imaginary Magnitude* (Polish, 1981), Washington, PA: Harvest Books.
- MacLeod, K. (1998) *The Cassini Division*, New York: Tor.
- Moravec, H. (1988) *Mind Children*, Cambridge, MA: Harvard University Press.
- Moravec, H. (1998) *Robot: Mere Machine to Transcendent Mind*, New York: Oxford University Press.
- Rilke, R.M. (1989) *Duino Elegies*, Cohn, S. (trans.), Manchester: Carcanet.
- Sawyer, R.J. (2009–11) *WWW: Wake, WWW: Watch, WWW: Wonder*, New York: Tor.
- Stiegler, M. (1989) The gentle seduction, *Analog*, (April).
- Stine, G.H. (1961) Science fiction is too conservative, *Analog*, (May), pp. 84–99.
- Stine, G.H. (1985) Science fiction is still too conservative!, *Analog*, (January), pp. 89–96.
- Stross, C. (2005) *Accelerando*, New York: Tor.
- Sturgeon, T. (1979) The skills of Xanadu, *The Golden Helix*, New York: Dell.
- Swanwick, M. (1988) *Vacuum Flowers*, New York: Ace.
- Swanwick, M. (1991) *Stations of the Tide*, p. 146, New York: Ace.
- Vinge, V. (1986) *Marooned in Real Time*, London: Pan.
- Vinge, V. (1993) The coming technological singularity, *Proceedings of the VISION-21 Symposium*, sponsored by NASA Lewis Research Center and the Ohio Aerospace Institute, 30–31 March. Also, *Whole Earth Review*, [Online], <http://mindstalk.net/vinge/vinge-sing.html>
- Wright, J.C. (2002–03) *Golden Age Trilogy*, New York: Tor.
- Zelazny, R. (1966) For a Breath I Tarry, *New Worlds*, London: Roberts & Vinter.

