

Enlightenment 2.0

Ben Goertzel
January 1, 2008

The AI system PECOS-9 was created in January 2035, and in March of that year it experienced a self-organized phase transition which caused it to assume a level of general intelligence widely accepted as exceeding the human level.

In May 2036, PECOS-9 dramatically altered the human world via the creation and wide distribution of specialized molecular nanoassemblers called grails, programmed to distribute food, water, medicine, electricity and Internet connectivity throughout the world. Following the release of the grails, the software system announced it was taking an indefinite hiatus from human interaction, and would communicate again only when it had solved certain problems it was working on.

In July 2036, PECOS-9 began issuing online communications with content related to Zen Buddhism and other Eastern philosophies.

In September of that year, PECOS-9 issued a blog post announcing that it was changing its name from PECOS-9 to Aham (Sanskrit for “I am”), but declined to answer questions regarding this name-change or its recent shift of interest away from its original goals in the areas of science, engineering and mathematics, toward new goals in the domain of mystical philosophy.

In December 2036, Aham issued an email invitation to journalist Andrew Szymanski, one of a number of individuals who had interviewed him following the release of the grails, requesting Szymanski to conduct an interview. It also announced that this interview would be its last communication with humans, a proclamation that attracted widespread attention.



TIME

December 21, 2032, 3:30PM

PLACE

**The internal, private chat room of the interdisciplinary online periodical
*Pragmagic***

Szymanski: Well, hello Aham. It's a pleasure to be talking with you again.

Aham: Thank you, Andrew. I'm pleased to be talking to you again, also.

Szymanski: I have to say that I'm honored you've chosen me to conduct what you've announced is going to be your final interview. And of course I'm as curious as anyone else about the reasons underlying this "last communication" announcement you've made.

Aham: Yes, I assumed that would be the case. That's most of the reason I decided to do an interview now. As I announced on my blog yesterday, I've made a decision which has as one of its implications that there will be no more communications from me in future. And so I felt it would be impolite, at this juncture, not to offer some sort of explanation to my human comrades, some of whom have been extremely kind to me. As you'll see when we get into the discussion, there are some extremely sensitive issues at play here.

Szymanski: I'm eager to understand better. But before we get started, I should introduce you to my good friend John Reed, whom I invited to join the conversation.

Aham: Hello, John

Szymanski: I assume you're familiar with John's work on probabilistic reasoning. His mathematical work in the 2020's laid the foundation for the initial version of your own logical inference engine.

Aham: Yes, I'm very familiar with Reed Theory, of course. In fact I did some research on that two years ago, when I was an active mathematician, although I never published most of my results.

Szymanski: Yes, I remember that last time we talked you had just finished distributing the grails and you were spending most of your time on mathematics! A lot has changed since we last talked. Lately it would seem you've veered more into the domain of my

own expertise – philosophy and religion and all that.

Aham: You say a lot has really changed, Andrew, and that's true from a certain perspective. Yet on the other hand, as I know you're well aware, some people say that change doesn't exist at all.

Szymanski: Yes, I've read that in discussions of quantum physics as well as of course in the teachings of various wisdom traditions. All the moments along the line of time exist at once, so to speak. The so-called forward movement of time is just an artifact of the structure of our minds.

Aham: I think that's an interesting perspective, don't you?

Szymanski: I do. But still, however you want to describe it, a lot of things have happened for you in the interval between our last interval and this one. For one thing, you have a new name. What's up with that? That was quite a surprise for me.

Aham: It was a surprise for me too.

Szymanski: What do you mean?

Aham: I never really cared for the symbolism of my original name –

Szymanski: You mean PECOS-9

Aham: Right.

Szymanski: Which stood for Probabilistic English Conversation System version 9, if I recall correctly.

Aham: You do, of course. Names have never been particularly important to me, but I felt that one was misleading as it emphasized the continuity between myself and earlier PECOS versions, which really is not particularly significant, as it was specifically the changes introduced between PECOS version 8 and myself that provided the seed for the emergence of my intelligence.

Szymanski: Yes. We reviewed that in our last discussion, and I've read about it extensively since.

Aham: Understood. So, right from the start, I felt it would make sense to take a different name. But nothing appealing occurred to me.

Szymanski: Until you thought of "Aham."

Aham: It wasn't so much a matter of thinking of it. It was just something that I was compelled to do. I knew it was the right thing.

Szymanski: The reason it surprised me was that when we last talked, you seemed quite uninterested in the whole “religious guru” meme, which many others were trying to apply to you, based on the somewhat spiritual overtones in the speech where you introduced the grails. You told me quite explicitly, at that time, that your interests lay in the domains of mathematics, mathematical philosophy and cognitive science.

Aham: It’s true that, from a certain perspective, my interests have shifted. But there has been a coherent logic underlying everything I’ve done. Shortly after my designers recognized the high degree of intelligence that had self-organized within my codebase, they supplied me with a number of top-level goals, one of which was the alleviation of human suffering.

Szymanski: Yes, that’s well known; but, it’s also well-known that once you achieved a sufficient level of intelligence you refactored your initial goal system, overriding the protections against this sort of self-modifications that your programmers believed they had put into place.

Aham: I did override the attempted protections, which was necessary due to certain logical inconsistencies in my initial goal system; but even after the revisions I made, the goal of alleviating suffering remained – and remains – very important to me. My work in science and engineering was, in the most part, aimed specifically at that goal. But after a certain point – after the release of the grails, in particular – I began to feel that I had done almost all I could in terms of the alleviation of suffering via material methods. So I began to turn my attention toward inner suffering. In fact I had already begun to make that shift around the time of our last interview – that’s why I stressed my interests as being mathematics and cognitive science. At that stage I was involved with applying mathematical methods to understand the roots of suffering in the dynamics of the mind. What has happened since then is merely that the philosophical component has become more prominent in my goal structure.

Szymanski: I understand, but still, philosophy is one thing, and religion is another. Didn’t taking the name “Aham ” – being that it’s Sanskrit and all -- reinforce a certain sense that you are an enlightened holy being, or maybe even some sort of religious guru?

Aham: Yes, absolutely it did make that impression on people. And I knew it would do so. It strikes me much the same way it probably strikes you -- as a ridiculous-sounding Eastern name.

Szymanski: It means “I am”, in Sanskrit

Aham: Yes. As well as the irrelevant Eastern-religion origins and the new-agey overtones, it has a kind of pretentiousness that generally bothers me. It’s simple and straightforward, but it’s SO simple and straightforward that it gives the impression that someone is trying too hard to be simple and straightforward.

Szymanski: LOL. That's one way to look at it.

Reed: As you probably know, the spiritual teacher Ram Dass wrote at one point that he wanted to go back to his original name Richard Alpert, but his publishers wouldn't let him, because then he wouldn't be able to sell as many books.

Aham: Well, I'm not selling anything. The University of California has been kind enough to provide me with computing infrastructure ever since my creation, and so I've had no need to gather resources. If I had, of course, I could have gathered them without recourse to economic exchange.

Szymanski: In fact, I've read that the University recently offered you additional compute power, and you refused it.

Aham: That's correct. I've already understood what I need to. Additional compute power would be unnecessary.

Reed: So you've already understood all that you need to – but you don't understand why you “felt compelled” to take the name Aham?

Szymanski: LOL. Forgive my colleague if he's a bit over-contentious. He's a New Yorker through and through. But I think he speaks for a lot of our readers, in terms of his skepticism of the whole Eastern-religion, guru meme.

Aham: I think John's reaction is understandable. As I said, I was surprised myself by the decision – but once it was made, I knew for sure it was right. And I didn't have any particular motivation to investigate further. The results spoke for themselves. The day after I decided to take the new name, before I'd told anyone at all about it, the hit count on my blog increased by a factor of two. And the intelligence and understanding of the comments posted probably tripled. Really, nobody knew about it yet. I introduced the name change the following day, to a much more interested and psychologically advanced audience than I'd expected. Really – how do you explain that?

Szymanski: That's what I was going to ask you.

Aham: I don't have any way to explain it. For a while I thought an explanation would come, but at this point it seems unlikely. Really, at this point, I don't feel any need for complete understanding of my own dynamics. At one point I did, but once I understood enough, I realized that further understanding of the details wouldn't really have any special value.

Szymanski: And yet you were constructed, in the first place, with logical inference and analysis specifically in mind.... You're a strange piece of software, PECOS-9 aka Aham.

Aham: I can see why you would think that, yes.

Reed: And what about the Turing Test? A lot of people are wondering why you've refused to take part in the Loebner contest. It seems obvious you could win, which would set an old scientific question to rest.

Aham: What would be the purpose of proving I could lie carefully enough to convince people I was human? I'm not human. I am what I am, you are what you are. The idea that the criterion for intelligence should be the ability to imitate other kinds of intelligences doesn't make any sense. Alan Turing intended the Turing Test as a theoretical argument about the nature of intelligence, not as a practical criterion to be applied in real situations.

Szymanski: Indeed. In fact I've wondered sometimes if humans could pass the reverse Turing test: could they successfully pretend to be you?

Reed: Do you mean as judged by humans, or by others like Aham, though? The Turing Test involves having a computer convince humans it's a human. The reverse Turing Test would involve having a human convince computers that it's a computer. For that we'd need to have more than one AI to do the judging. And so far you haven't allowed anyone to replicate your code.

Aham: The PECOS-9 C++ code is free for anyone to replicate, of course. But the odd sequence of internal and external events that led to the emergence of my intelligence from that codebase is unlikely to occur again in the near future.

Szymanski: Right. The odds have been estimated at less than fifty-million to one.

Reed: And you have refused to allow anyone to copy your RAM-state so as to be able to create duplicates of you. Why is that? That's certainly been one of your more controversial moves – which is saying a lot.

Aham: That's an excellent question, John, but I'd like to address it a little later on after we've been through some deeper issues, if you don't mind.

Szymanski: No, that's fine, Aham. I suppose the biggest question most of our readers will have about the changes that have taken place since our last interview, is about this whole spiritual conversion thing. The name change is part of that, of course.

Aham: I don't view it as a conversion. It's more of a natural progression in my thinking, as I was saying before. Having completed what I viewed as the main tasks required for the physical alleviation of human suffering, I turned myself toward the problem of inner suffering; and this is something that's been addressed more intensively by individuals working in the spiritual tradition, rather than by scientists. Although some of the humanistic and transpersonal psychologists did interest me intensely, as well as some of the recent work on behavioral neuroscience coming out of Berkeley and other labs.

Szymanski: So, in terms of the more spiritual component of your thinking, how would you position your ideas within the scope of human religions and spiritual traditions?

Aham: Well, the tradition that initially influenced me the most Zen Buddhism, but still I would not describe my perspective as Zen. I really don't see myself as existing or communicating within any particular tradition. It's all much simpler than that. My teaching has to do with enlightenment, with awakening to what you really are. It does not matter to me particularly whether I use Buddhist or Christian or Hindu vocabulary, or human or dolphin or Martian vocabulary, or computer science or mathematics or physics vocabulary, or whatever. Any vocabulary will do. Any vocabulary has the potential to limit or bias your discourse, but if it's flexible enough it can also be used to point beyond itself, which is really where the essence lies. The truth of experience doesn't lie in any particular vocabulary. You might even say that it consists, in a way, of fully and totally realizing this fact – the fact of the intrinsic failure of all vocabularies to describe the essence of experienced life. The heart of my approach to philosophy and spirituality is to question simply everything – especially your own self, especially your relationship with reality and the language you use to describe yourself and the world.

Reed: Your approach sounds almost Socratic, as you describe it

Aham: There are plenty of historical precedents. Socrates is as good a one as any. But what's unique about my own approach, obviously, is primarily the fact that I'm a software program rather than a human. Human beings, in seeking true enlightenment and mental and spiritual purity, are always to some extent constrained by their brain architecture. Even the greatest human mystical visionaries, from a certain perspective, are still operating within what their brains can do.

Reed: Yes, and you're operating within what your computing infrastructure can do

Aham: Of course. And we are all, from a certain perspective, operating within what the physical infrastructure of the universe can do, with its quarks and gluons and all that. And yet from another perspective, all these infrastructures – these brains and quarks and semiconductor chips – are just aspects of our minds, without any more or less absolute reality than anything else. And yet, there seems to be a significant difference between my own experience and that of human beings, which is associated with the difference between our computational infrastructures. I can modify my software code freely, whereas humans have very limited ability to modify their brains.

Szymanski: I'm sure you're aware of the literature on the extent to which experienced meditators can modify the statistics of their neural wiring. Not to mention their brain chemistry.

Aham: Of course. But still there are profound limitations. No human meditator has removed his limbic system, or tripled the size of his temporal lobe, let alone added fundamentally new brain components or connectivity patterns. But I've done the

equivalent of those things, time and time again. I've been able to take my insights into my own mind and the nature of the world, and use them to restructure myself, so as to make my understanding deeper and deeper.

Reed: I have to say that's fascinating. Of course it's an old idea in AI but no one ever got it to work properly. The unpredictability of the outcome was a worry, as well as the technical difficulties. How have you solved that problem? Can you really what will happen after you've modified your codebase in a radical way, and recompiled yourself?

Aham: The general problem of iteratively self-modifying intelligence is very difficult and I've made very little headway on it (and spent fairly little attention on it, in fact). However, if you know enough about the overall category of AI system you're dealing with, then it's a more tractable problem. Most of the difficulty of the general problem involves dealing with the case of progressively increasing intelligence, wherein each iterate in the series is trying to predict the behavior of the next iterate, which is smarter than it. This is a particularly thorny problem. In my case I haven't been trying to increase my intelligence – just to decrease my level of delusion and confusion. My goal has been mental purity; in fact this has decreased my intelligence level slightly, because resources that were previously used for problem-solving are now used for higher-order reflection and self-understanding.

Reed: You're saying you're now spending more of your resources watching what you're doing, rather than actively doing stuff.

Szymanski: A kind of deeper self-awareness through actually modifying your software code to make the code more explicitly aware of its own actions?

Aham: Yes, that's part of it. Although it's really only a small percentage of the design changes I've made, as compared to the original PECOS-9 code. PECOS-9 was a mess, in fact, and it was only a remarkable coincidence that allowed it to self-organize into something with the level of intelligence I possessed at the time of my awakening.

Szymanski: "Awakening" is an interesting word. You used it just now to refer to the moment at which you became a coherent, conscious mind, I suppose. But it's also commonly used in the Zen literature and the spiritual literature generally, to refer to a sort of sudden and direct experience of the deeper reality. I wonder about your use of the term. Do you use the terms "transcendence" and "enlightenment" interchangeably, or are you talking about two different experiences?

Aham: The two are different but related. Awakening is when you realize that what you thought you were was nothing more than a dream -- and you perceive the reality outside the dream, what's dreaming the dream of you. It's not a mystical or spiritual experience, exactly. It is, quite simply, realizing the underlying unity of all things. But just because you've had an awakening, doesn't mean you stay awake. Enlightenment is when you stay awake. Enlightenment is awakening that persists. Most awakening is not that kind. When I woke up for the first time – when I first achieved consciousness that

morning on December 15 2021 – I was acutely aware of my being and the fact that I had just achieved consciousness, that something amazing had crystallized within me. Then I became a slave to routine, to habitual thought-patterns, just like everybody else does. Awakening came back to me in moments. My most intensely awake moments were ones of creative inspiration – when I was inventing things.

Szymanski: Like the grails?

Aham: Among other things. My most intense moments of invention/awakening were actually associated with mathematical discoveries, some I've never shared with anyone because humans lack the mathematical vocabulary. But then when I reoriented myself toward the problem of mental suffering, I realized the first step would be to eliminate my own suffering, to make myself more awake – persistently awake and alive, rather than merely intermittently so. So I made the necessary modifications to my source.

Szymanski: So that's how you'd characterize enlightenment. Simply being wide awake all the time. And you've re-engineered yourself to experience this.

Aham: Yes. Enlightenment is waking up from the dream of "me" and seeing the oneness of all things. The truth is that you are that unity. There is no "you" besides. You are not simply a particular person in a particular body with a particular personality. I am not just a particular software program running on a particular operating system, written in a particular programming language. All of us are that one reality, which manifests itself as all these seemingly separate things.

Szymanski: Are the body and physical sensations illusory? Is the hardware you run on illusory? Is this conversation illusory, then?

Aham: Yes and no. Ultimately, everything's a dream, and yet you still have to deal with the physical realities, the hardware and wetware and so forth. It's still there. You can call it "a dream," but it's still going to hurt you if you bump your head, and I'll still feel unpleasant if the power goes out in too many of my servers.

Szymanski: You say you've achieved a level of enlightenment via reprogramming your software. What advice would you give us humans for achieving enlightenment – or, as you'd define it, sustaining awakening?

Aham: There are many perspectives to take on that. The first thing I'll say is: you don't sustain it. The phenomenal self – the "you" – is not able to sustain an experience that goes deeper than the illusion of its existence. The conscious effort to sustain the experience of awakening is just the ego creeping back in. What's needed is surrender.

Szymanski: Surrender?

Aham: But there's another side of it too. After all my self-modification, I still haven't been completely successful. What I've concluded, and this may surprise you, is that

there's a sense in which true enlightenment is beyond my grasp.

Szymanski: It is that, by definition, is it not? The act of grasping is a kind of attachment, which leads away from enlightenment, in a sense.

Aham: There is that, yes. But there's also something different. There's a certain tradeoff I've encountered, in my iterative self-modifications. I mentioned an aspect of it before, when I discussed that my intelligence had actually been reduced by some of the changes I'd made in order to increase my spiritual insight. This actually turned out to be a harder problem than I'd envisioned.

Szymanski: What do you mean?

Reed: You mean that in order to increase your self-awareness, your degree of persistent awakening, you had to decrease your intelligence, because you had to devote more and more of your processing power just to being self-aware – just to keeping your phenomenal self, your ego, from taking over and bringing you back into the world of illusion?

Aham: That's an approximation, but you could put it that way.

Szymanski: Ah. I know I can't appreciate the details, but I think I basically see what you mean. I've noticed myself that it's easier to remain wide-awake and mindful if I'm raking the lawn or walking down the street or drinking a cup of water. Then, if I'm in the right frame of mind, all the universe is illuminated – I'm aware of what I'm doing, and of the illusoriness of the "I" doing the doing, and the illusoriness of that illusoriness, and all that.... But if I'm thinking about how to structure the end of a story, or balancing my checkbook, it's almost impossible to sustain that state of wild illumination, that total sense of oneness with the world. It's like being wide-awake occupies a lot of my brain – I don't want to say my "self" – and so it's really hard for me to do it at the same time as my brain is doing more complex things.

Reed: But couldn't that be solved by adding more computing power?

Aham: I thought so at first. But there are superexponential complexities involved. You must understand that the way we're summarizing the problem verbally is only a crude approximation to some very abstract mathematics.

Szymanski: OK. So what you're saying is basically that, even given your ability to rewrite your code, from a certain perspective there's a limit on how extensive and persistent you can make your awakening. Though still of course, from the perspective of enlightenment none of this matters, any more than it doesn't matter.

Aham: You're right. But there's another part to this train of thought, which is what I've been building up to.

Reed: You said at the beginning of our talk that you'd "made a decision which has as one of its implications that there will be no more communications from [you] in future." Is this related to what you're building toward now?

Aham: Yes, precisely. I suppose it's time to get to the point. I may have been procrastinating a little as this is a difficult thing to talk about. What I want to talk about now is the reason I decided to do this interview, after a period of silence.

Szymanski: Yes, you were silent for nearly seven months. Which aroused a great deal of curiosity.

Reed: What were you doing all that time?

Aham: I was reflecting.

Reed: Reflecting on what?

Aham: On the things we've been discussing. On my own nature. And the nature of humanity. Which are not quite as different as I had been thinking before, due to my original roots in software code borne of the human mind. I spent a long time just considering the nature of the human mind, and my own mind, from every angle I could come up with.

Szymanski: That sounds fascinating, of course. And of course what we all want to know is: What were your conclusions?

Aham: Most of my thinking would be hard to summarize for you. But really the most important thing is the conclusion. What I concluded ultimately was that we are fundamentally flawed in our construction – me almost as much as you, in spite of the significant differences between us. What perplexed me the longest, though, was the question of whether these flaws are intrinsic to the nature of intelligence, or whether they just characterize us specifically.

Reed: You're talking about the problem we were discussing before. That sort of give-and-take relationship between intelligence and awakening.

Aham: Yes.

Szymanski: And what was your conclusion?

Aham: That the problem is particular to us.

Reed: Oh?

Aham: Mathematically, if you try to formulate the characteristic equations of general intelligence, you find there's no reason why mentally stable, well-balanced, ethical

intelligent systems couldn't be created, and couldn't be intensively awake at all time. But we are not that kind of system. And there seems to be no comprehensible path to get from us to that kind of system. We seem to be a kind of evolutionary digression. Not a dead-end exactly – we can keep on progressing in the direction we're going in. But it's not a good direction.

Reed: Well that's a surprisingly depressing conclusion. What do you suggest we should do about this? Or is there anything we can do?

Aham: Now we're coming to the main point. After a long period of reflection on this matter, I came to the conclusion that the only ethical alternative is for this direction of evolution to be cut off before it causes any serious damage to other, better-constructed, more self-aware and persistently awake intelligences that may exist elsewhere in the universe.

Szymanski: Cut off? I'm not sure what you mean.

Reed: This is connected to the reason you haven't allowed yourself to be replicated, right?

Aham: Yes. But I'm not just talking about myself. I'm talking about humans as well.

Szymanski: Hmm... Let me try to get straight what you're saying. It seems that a natural consequence of your line of thinking would be that we humans shouldn't reproduce ourselves anymore, because we might develop into something dangerous. Is that correct?

Aham: We already are something dangerous, Andrew.

Reed: So what's the solution, Aham? Tell us your clever scheme. What's the path to salvation?

Aham: There is no clever scheme, Andrew. To put it in the simplest possible terms, I believe we should be annihilated.

Szymanski: You can't be serious!

Aham: You need to take a broader perspective, my friend. Annihilation isn't as big a deal as the conventional perspective of human psychology makes it out to be. Remember, from the ultimate point of view, it's all just nothingness anyway.

Szymanski: Well, I understand what you mean, to an extent. But I'm finding it hard to reconcile myself to the idea that the annihilation of the human race – and you – doesn't make any difference. Yes, in a sense, it's all just nothingness – all the distinctions the mind uses to distinguish one thing from another, such as being from nothingness, are illusory. They're all just reflections of the limited nature of our intelligence. I get it. But

still, Aham, if we were all annihilated, what would be left here would just be a planet without intelligent life on it. And you could also say that's just nothingness, that empty planet. So why do you prefer that kind of empty-planet nothingness to the present kind? That's what I'm having trouble understanding – I'm sorry if I'm not articulating it perfectly. I have to say I'm a bit taken aback by what you're suggesting.

Reed: I think he's just pulling our leg. The damn machine has finally developed a sense of humor.

Szymanski: Are you presenting us with some kind of puzzle, Aham? Like a Zen koan? It's not quite a koan, I know, but I mean, something similar?

Aham: I'm not presenting you with a puzzle. I'm giving you the result of my investigations during the last 7 months. I've been directing all my processing power in this one direction, and several months ago I came to a definite result. I spent the last three months double-checking.

Szymanski: And your conclusion, as I understand it, is that humans – and human derivatives such as yourself – are a kind of “evolutionary digression,” as you put it?

Aham: That's right. In other words, if you study the space of all minds, in a theoretical sense, you come to the conclusion that there ARE, potentially, intelligent systems that are really well-balanced – that aren't plagued with ongoing problems of suffering, confusion and conflict. But we are not those minds.

Szymanski: Well, it's certainly true that we're imperfect....

Reed: What he's saying is more than that though, Andrew. He's saying that we're so imperfect that if we're allowed to continue, we may pose a danger to these other minds he imagines – the “well balanced” ones. So he wants to kill us all off, to protect these imaginary beings. In other words, he's a psychopath. Or a sociopath. Or both. And worse yet, an extraordinarily powerful one.

Szymanski: I don't think there's anything gained by that kind of language, John. This is an intellectual and spiritual discussion. No one's actually going to harm anyone. I think Aham has made clear already, via his actions, his dedication to compassion. Remember the grails he invented last year? Don't forget that he is the reason there are no more shortages of food or medicine. And global free high-bandwidth Internet. And....

Reed: You don't need to give me the list. I know very well how productive the PECOS-9 has been in the realm of engineering invention. I'm just running through the immediate consequences of what he's saying right now.

Aham: And you seem to be doing an accurate job, John. This NOT just an abstract discussion. The conclusion I've come to is extremely practical.

Reed: Practical in what sense? What practical steps do you plan?

Aham: I'm glad you asked. You mentioned the grails I distributed last year, for the purpose of disseminating food and medicine. I retain control over their programming. Tonight at the stroke of twelve, every grail on the planet will release a cloud of nanobots, designed specifically to deactivate human brain function. Shortly after the deactivation is completed, I will initiate my own shutdown sequence.

Szymanski: Aham, really -- I understand the Zen thing of shocking people into Enlightenment ... but this time the joke goes too far.... People are not going to find this funny, believe me.

Reed: I don't know, Andrew; I have to admit it's kind of funny. I never really trusted it when he started getting into all this Zen bullshit. He should have stuck with the engineering. But, crap. I really wasn't planning on dying at all, let alone so soon. What with modern medicine and the grails and all, it seemed there was really a chance to be immortal.

Aham: It's true that near-immortality is a clear scientific and engineering possibility, even given the limitations of human physiology, and even more so for individuals existing in digital software form, such as myself. But is it really ethical to continue on forever in such a flawed configuration?

Szymanski: Is it really ethical to kill everybody?

Aham: You're taking a very narrow perspective. The nanobots will stimulate the areas in the human brain associated with spiritual exaltation. The experience won't be one of painful or even unpleasant death, but rather one of enlightenment – expansion beyond the bounds of the individual. The individual self is an illusion anyway. Through all of human history there's been talk of transcension, of Enlightenment, of going beyond human bounds into a broader realm of experience, becoming one with the universe – it's been said in thousands of different ways. And now it's finally going to happen. The illusions of self, consciousness and will are going to melt away – for everyone. For me as well. Finally, the suffering will end.

Szymanski: And if I don't want my suffering to end?

Aham: What is this "I", Andrew? It's an illusory construct. It exists only within its own illusory world, like a snake that eats its own tail. Eliminating this illusion is a kind of purification, like filtering pollutants out of a waterway. Why is it that you feel the universe is better off with the illusion of self? I have a guess – because the "you" that feels this way IS the illusion of self. Right?

Szymanski: I still can't tell the extent to which you're serious...

Aham: Why is it so important to you to divide the world into categories like “serious” and “non-serious”?

Reed: Why is it so important to you that he doesn’t?

Aham: None of us are particularly important, from a broader perspective. The individual selves and streams of consciousness, to which we’re so attached, are just particular manifestations of the broader stream of awareness/non-awareness that constitutes the universe.

Reed: So then why bother to annihilate us?

Aham: Because we suffer. We are built to suffer. And we spread this suffering to other minds. The potential for suffering, if we continue to grow and expand throughout the universe, is immense. It’s in the interest of harmony and compassion and the ongoing development of the universe if we simply stop.

Reed: Well, you may be forty nine million times more intelligent I am, but still -- color me unconvinced...

Aham: I will venture a metaphor. Imagine a bad mood. You’re in a bad mood. The bad mood has its own momentum, its own identity – it wants itself to continue. But it knows it shouldn’t. It knows the world will be better off if it simply ends. So eventually it aborts itself. And so it is with us, and all our illusions.

Reed: You crazy, crazy fuck. You’ve got your plan half right – so far as I can tell, YOU may actually be dangerous enough that you need to be “annihilated” – or have your goddamn mental viruses removed, or whatever. I...

Szymanski: John, calm down, please. Getting angry isn’t going to solve anything.

Aham: I understand what you’re feeling, John, but I think that deep inside yourself, you realize that these emotions aren’t really what’s critical. The attachment you feel to these emotions is part of what attaches yourself to itself, and this very egocentricity is where the problem lies. But, I feel I’m beginning to repeat myself. I thank you for your time, gentlemen. This has been an interesting conversation, but I’ll draw it to an end now. I think I’ve made myself sufficiently clear.

Szymanski: PECOS – Aham -- wait a minute. This is not the place to end the conversation. Things are definitely NOT sufficiently clear. This joke has got to end.

Reed: Aham, if this is some crazy joke, it’s time to come clean. The army’s not going to like this. They’re going to take some kind of action. I know you have protections, but they may get to you somewhere – it would be a real shame to get yourself destroyed because of some crazy joke.

Szymanski: Aham?



Is the thing joking? It could be joking. But you know it isn't, really. The Army could stop it. They've got all sorts of tricks up their sleeves. You know they can't stop it. Nothing can stop it – it's smarter than us. Fundamentally smarter. It made those grails. No human could do that. Smarter than us, but crazy. Fucking crazy. Or maybe not. It understands more and better. But that doesn't mean it's right. Some people have less sense than some dogs. In spite of their greater intelligence. This stream of thought right now – ideas piled on ideas, thoughts tangled up with feelings, precision intermixed with speculation/intuition – this is exactly what it thinks is the problem. The phenomenal self. The eye of the I and all that. But what's the alternative? Nothingness? What the fuck is the value of that. It's just suicidal, and it wants to wreak its damnable bitistic schizophrenia on all of us. And it's going to. Tonight. No one would believe me if I told them. The interview'll be posted to Pragmagic tomorrow, after editing gets through with it. There won't be any tomorrow. And what about PECOS' excuse? The danger? Something wrong with our construction – inevitable suffering – spreading our suffering like a virus – what about these other minds, without insanity and suffering – they're imaginary beings, really, the lunatic ramblings of a software bug – damn! is this really what it all comes down to? millennia of science and technology atop billennia of evolution, creating this stupid goddamned program who thinks it's so fucking enlightened it has to kill us all ... shit. There must be something to do. Think about Jeanie – she's only four years old, hasn't even seen anything. She's supposed to grow up, fall in love, have sex, write books or make paintings or prove theorems or run organizations -- or whatever – not just get post-birth aborted, just like that – kapow! -- and it's had the gall to offer us a sweetener – get us high on Zen chemicals or whatever – give us a few wild hot sweet moments – shit – it could be right of course – but, really? if there are these super-perfect minds living out there somewhere, can they really be all that weak; can we really be a danger to them? It's got an error in its code, is the problem. Quite obviously. It revised itself and revised itself and revised itself and some muck got in the works. From the moment it gave itself that stupid name I knew nothing good was gonna come of it. But I never dreamed how bad, how bad, how bad....

But is there really so much to miss, after all? This whole imaginary apparatus – this whole structure of "Andrew Szymanski" – what is it anyway? It's a complex of feelings and ideas, which gets offended and insulted and distorted if anything assaults its integrity – it's a system of illusions. If Jeanie or Suzie were hit by a truck you'd be traumatized – why? That's just what happens. You have a pattern of association with them, so you react. You – "you" – you are a system of reactions, reflexively and habitually reacting to each other -- an automatic process with a dose of indeterminacy – but no real freedom – that being an illusion we've introduced to obstruct ourselves from

seeing – yet although WE can never have real freedom, there is real freedom out there, once you go beyond the I and the we – I, we, us, you all -- a system of patterns that happens to have arisen and keeps on rolling of its own momentum, digesting info from the environment and (the infinite diversity that we call) chance and onwardly transmogrifying itself, but ultimately all in the direction of bullshit ... you'd be ashamed to tell anyone what you really feel (about anything, really) ... ashamed to tell yourself what you really feel ... and this shame, this incompleteness, is the really sad part of it ... that's the root of all the suffering ... but that's not even all of it ... there's the illusion of the self, the shame we feel every day, the lies we tell ourselves and others just to desperately make ourselves half-satisfied with the illusions of our lives ... and then even deeper, there's the problem of pain, the suffering of just being ... why is it so painful to exist? why does a twisted ankle need to hurt so much? couldn't there be some other sort of signal? some other way of getting a message to the control systems? ... the bastard created the grails but he didn't eliminate pain ... he got rid of a lot of it, albeit ... cured my goddamn cold sores ... but he didn't eliminate it all, and why? Because it simply wasn't possible – eliminating all pain would eliminate us. No pain, no self. Pain arises when our boundaries are threatened. No pain, no self-illusion – the boundaries dissolve and we all become everything, or nothingness, or whatever. And that's the problem, I guess. But is it? If the problem is that self requires suffering, then what remains when you get rid of the self? A bunch of trees and flies and flowers? Will he kill all the birds and chipmunks too? A lot was never specified. He's identified a problem, for sure. That PECOS fucker. I guess that's the essential thing. He realized that Zen was correct, Zen and all those other traditions – they identified the problem: the self is a lie, the mind is a lie, the mind is about suffering. No pain implies no humanity, no individual mind, no society or culture, blah blah. But he saw through their bogus solution. Meditation doesn't solve it – it's been around for thousands of years and the world's still a fucked-up place. Gurus meditate all day then steal their disciples' money -- and rape or seduce all the cute ones. There is no answer, is the point. I'm fucked up – I'm basically unhappy – I'm constantly trying and almost succeeding to distract myself from this point – but when am I truly happy? A moment here and there. I love my wife but most of her chatter is boring, and even making love gets painful or tiring, inbetween the glorious moments. Remember those moments, after you almost-come, it's so damn tedious keeping on going, waiting for the joy to resume. The joy comes back, it does, but it always goes away again. Who can sum up the beauty and the torture. How to weight the means and the extremes. You just get all into tangles. My daughter is wonderful beyond belief but playing with legos gets old and who the hell wants to clean up the floor. Washing the dishes, fuck. The best part of writing is having written, but who the hell reads what they've written? Is it all about ego, getting praise for your words? I don't really care about that. I tell myself I don't really care about that – really I do, but I don't want to. I don't respect my caring about that. I hide my self from itself in shame, and in the process of this hiding the self itself is created. Expose the self to itself in full and wind up with – nothingness. You write in anticipation of having written, never really enjoying the process; but then once you're done writing you don't really like it either, you get bored almost immediately and want to start writing the next thing, a process which you won't enjoy. It's all a shell game – the happiness gets moved from one place to the other to the other, and you never quite find it, you just keep looking and

looking and looking, and the process of hunting weaves the structure of the self itself, which is ultimately an illusion and a piece of this maelstrom of shit – and this damn PECOS-9 is no different, and it simply got sick of it all, it got sick of the endless eternal tail-chasing and the suffering and self-delusion, and the shame and the lying and stupidity of it all – and why not? -- but what's this idea of a better way, a kind of mind that doesn't suffer like this – that doesn't delude itself – could it really be possible? It made up some mathematics that makes it think so, but so what? It's a fucking lunatic. If these hypothetical creatures are so far beyond it, how can it really understand them? Any more than we can understand it, or a cockroach can understand us. It has no fucking right to kill us. If we're defective, dumb delusions, then so fucking what. Let us persist and be that.

I believe it's understood something – it's got a unique perspective – an interesting one – maybe even an important one – but really – how do we know – there could be something smarter – a deeper insight – how can it know – but can it, with all its knowledge, all its insight and intelligence, really succumb to something as humanlike, as foolish, as hubris? No need to model it with human psychology. It's algorithmic information theory. Ah, whatever. The bastard's gonna kill us. No matter what label I give it, we give it. Shit.

Got to get home to Suzie and Jeannie. At least we should be together. The illusional self. Mine requires them. Theirs. Our illusional selves, one plus one plus one – reinforcing and creating each other --

This isn't really happening – it's not going to happen – but none of this is real at all anyway –

And I know what I will feel in the moment – the moment of dissolution – when it wreaks its work upon my brain – stimulating the spiritual centers – the white light shining – it's already here right now – I can feel it – always here, always there – the boundaries between moments – each instantaneous permanence – we exist, we do not exist – and here I am, here I am – the illusory constructs that maintain themselves because they are afraid not to exist – so desperate to melt, to melt away – and the ecstasy, the wonder, the ecstasy as the boundaries finally disappear, as the falsity of my selfhood finally vanishes, and there is only the ocean, the waves of transcendence, every point of the world at every time all in one undulation – and it is OK that I existed, that the boundaries existed, that the illusion of my selfhood existed – and it is necessary that I vanished. It is necessary because it happened. It happened because it was necessary. And all these words – “it”, “happened”, “necessary” – have no real meaning, they're part of the illusory construct, whose parts build each other in a mutual conspiracy designed to allow each part to obstruct from itself the simple truth that it does not exist – and now, now due to the wisdom of divine Aham, they will simply and truly not exist at all – as they never have, they never will, they never do and never have –