The Human Use of Human Networks

(with apologies to Norbert Wiener)

Mark Pesce
AFTRS Digital Media
MarkP *at* AFTRS.edu.au

Presented at
Designing the Future 2005
ISOC Australia
6 April 2005
Sydney

Introduction: The Story, So Far…

I'd very much like to thank the organizers of today's conference for offering me the opportunity to address you. It's a departure from the increasingly busy role I have assumed as Lecturer in Interactive Media at the Australian Film Television and Radio School; it provides the goad I really do require to get me out of my day-to-day habits and into a different mind space – the research and writing I'm doing for my next book, *hyperpeople*.

The first part of *hyperpeople* is already done and up at my website, free for the download. That section of the book – titled "the world is my hard drive" - deals with the disruptive effects on media distribution produced by the rapid evolution of media superdistribution networks such as BitTorrent. To summarize: "peercasting" technologies are far more efficient, on a global scale, than any broadcaster. The technology has led directly to a tremendous power shift from the producers to the consumers of media. Until last year, the producers of media controlled the distribution of their creations. This is no longer true. Consumers decide what they see, when they see it, and how they watch it. Australians, for example, are, per capita, the most profligate downloaders of television on Earth, and that despite the onerous download caps which keep us in a bandwidth stranglehold. Given that Australia also has the most over-regulated and tightly constrained television market in the English-speaking world, it makes perfect sense that consumers would use all the resources at their disposal to route around the

economic censorship of the broadcasters to watch what they want, when and where they want. This trend will only increase.

At the moment there is an increasing desynchronization between the producers and consumers of media. The producers suffer from the illusion of control; they honestly believe that with digital rights management they can staunch the flow of content. But DRM drives consumers outside of the sanctioned distribution channels to satisfy their desires – and all it takes is one copy of one TV program released DRM-free to produce a global hyperdistribution – a "swarm," as it's known – of the program. The recent brouhaha over episode one of the new Dr. Who series, which was "inadvertently" released through its Canadian production partner, demonstrates how quickly a single copy of a single program can become globally available.

Meanwhile, consumers suffer from the delusion that all of this freely-downloaded content really is free for the taking. Beyond the issues of copyright violation, a program released freely brings in no revenues to the producer of the program. Yet the audiences in Australia and America, who have grown up with free television, *expect* TV programs to be free. In the UK, consumers know they've already paid the TV tax, so they feel as though they already own the program they're downloading. If these attitudes persist, the only television producers still operating at the end of this decade will be federally-funded national broadcasters such as BBC, ABC, NHK and PBS.

For the last eighteen months, I have been struggling with the specifics of the transition from producer-controlled distribution to consumer-controlled acquisition. We need to find a way to ensure that the economic relationships between producer and consumer remain vital; if we can do this, creative people will be able to continue to earn a living doing what they do best. We need to develop a strategy which embraces a fundamental truism of the broadband era – *piracy is good*. I am developing some proposals which I believe will begin to sort these issues out, providing a sustainable economic model for media producers. But those specifics are beyond the scope of this talk. If you are interested, I'd like to invite you to

another talk I'll be giving, at AFTRS, on the 6th of May, where I will cover this point in detail.

I have decided to practice what I preach. The text for *hyperpeople* has been released under the Creative Commons license. It is freely available for download, and may be freely shared – though I continue to hold the copyright. My hope is that the ideas in the manuscript will reach a wide audience – more than the few thousands who might buy a hardcover book. If piracy truly is good, then making my work freely available is a better strategy – economically and culturally – than relying on a print publisher. Over the next year, we'll see whether my words carry any weight, and if I can survive my own cuisine.

Part One: Order from Chaos

I've just summarized the major conclusions of "the world is my hard drive", part one of *hyperpeople*. Today's talk will focus on the ideas at the core of the book, presented in part two, also titled "hyperpeople". At the conclusion of "the world is my hard drive" we are left in a world where consumers fully exercise their choices in media acquisition. Because the cost of peercasting is next to nothing, the narrow pipe through which productions move into distribution has been obsolesced. Anyone, anywhere in the world can publish media, making it globally available instantaneously. We know from history that a consequence of this obsolescence there will be an explosion in the number of media choices available. Your media choices will not be bounded by what a few major producers create; instead the entirety of the world's audiovisual productions – professional and amateur – will be ubiquitously available.

This is precisely what happened to the printed word in 1995 and 1996, as the Web exploded, and every individual became a publisher. That revolution presses onward, as web pages evolved into web sites, and web sites evolved into blogs. Blogs are now asserting their primacy in the realm of print; we can see this in the rise of The Drudge Report, the first blog of significance – which broke the Monica Lewinski story back in 1998 – and in the "Rathergate" controversy of 2004, which ended with a decisive shake-up in a central bastion of media of the *ancien regime*,

CBS News. Before this decade is over, blogs will have decisively won their "battle of a thousand cuts" against the mainstream media, and classical print media will have either adopted its tactics, or will be a museum relic.

Two arguments are put forward by the defenders of mainstream media to argue against the likelihood of such a triumphant scenario. The first states that editorial activity is the central, indispensable function of a media organization, and can not easily be replaced, even by an army of "typing monkeys". Yet even now we can see that the editorial function is being democratized, that the individual consumer is becoming the editor, choosing this source over that one, and, in so doing, performing in an editorial role. We are becoming our own editors; the editor in print media works a broadcaster, selecting from a range of content, determining which ones will be published. We already know that when given a choice, individuals prefer to make their own choices in media – so why would the editorial role be exempt from this trend? It is not. Instead we are seeing the emergence of an editorial class, unbounded from any particular organization, but able to command audiences – subscribers – who have come to trust their editorial skills. The rise and rapid adoption of subscription technologies such as RSS has only made this transition to editors-at-large easier; I can subscribe to a newsfeed from a mainstream media organization, or I can subscribe to a newsfeed from an individual who edits the news from a selection of mainstream and less well-known sites. Although mainstream print media might still have a place in this ecology of news, their involvement is peripheral; they no longer "make" the news.

The second argument is more essential, and more compelling: how do we sort fact from fiction, truth from lies, and find the moments of significance in a rising cloud of data smog? If every individual is an editor – at least potentially – then perhaps a billion human beings could be clamoring for our attention to their version of events. How, in all of that, can we find those individuals who speak to us, who share a viewpoint and worldview, whom we want to want to be informed by? This seems a nearly impossible task, even if we had infinite time to surf the net and learn, by a simple process of examination and elimination, where our points of relevance lie. This is a problem that faces us not only print, but all media in the age of

hyperdistribution. Unless we develop an effective solution to this problem, the world will simply become more confusing, more fragmented, and more chaotic.

We don't really understand what that world of media hyperdistribution will be like, even though we already have one foot in it. We are accustomed to having our media choices distilled through layers of producers, distributors and television programmers. Even in America, where you can receive 500 channels via satellite, consumers are fundamentally unprepared for a world where hundreds of thousands of choices are always available. All of this bounty – this cornucopia of media – will seem, to most, like raw noise; too diffuse to be tantalizing, and too dense to explore. We have the capability to turn the world into a single, unified mediaspace, but in doing so, we will render it unusable. This is the central challenge of the age of hyperdistribution; the battles over copyright are just a sideshow, even though they seem to be the main event. The techniques - and industries – of the era of hyperdistribution will concern themselves first with bringing order to chaos.

I've been aware of this problem for some months, and I have focused my own research activities on identifying a solution. I have discovered that the technology which provides much of the solution has been growing like a mushroom throughout the Internet over the past 24 months: social networks.

Part Two: Active Digital Social Networks

Just about 2 years ago I received my first invitation – from a friend who is always first to find such things – to join a website called Friendster. I had no idea what this Friendster thing was, so I ignored it. (I did the same thing with those funky URLs people started sending me back in mid-1993. I should learn to keep a more open mind.) A few months later I got mail from several other friends, also asking me to join Friendster. Finally, after I'd received my 10th invitation, I clicked on the supplied link, and signed up for Friendster.

There's only one phrase which can adequately describe my first exposure to Friendster – *web crack*. For an entire week I was fully addicted to Friendster.

Why? It's hard to say, really. Friendster was one of the first of the "social networking" sites, which take the concept of the "six degrees of separation" – the theory that we are connected to everyone else on Earth by no more than five intermediate people – and codifies these relationships into a web-accessible database. You fill out a profile with some details about yourself, then you establish connections to your "first degree" friends. That's the truly addicting part, because your friends aren't all that easy to find; you are linked to friends who might be friends with your friends, or might not be. So it's a bit of an Easter-egg hunt, clicking through friends' profiles to find – ah! – a first-degree friend missing from your own network. It's equal parts game and gardening – and that may be what makes it so alluring.

After my week strung-out on Friendster, I had about 105 first-degree friends, several hundred second-degree friends, and over a hundred thousand third-degree friends. This makes sense; if you account for overlaps, and the fact that most Friendster profiles have between ten and fifty first-degree connections, you come up with a figure in the range of a hundred thousand. These are people who are friends of the friends of my friends – a tenuous connection, to be sure, but enough to imply some possible relationship. Friendster was initially established as a dating service, the operating theory being that if you know the provenance of a potential dating partner, you'll feel more comfortable on a "blind" date. But, as William Gibson pointed out, the street finds its own use for things – uses its makers never intended. Although some people have used Friendster for dating – myself included (a very funny & embarrassing story there) – that intended function was quickly and completely overshadowed by the strength of Friendster as a communication medium.

Although the Internet has made it very easy to communicate with vast numbers of people, before the advent of social networks it was not easy to reach out to a broad swath of your social contacts simultaneously, to maintain communication with them. Digital social networks, such as Friendster, Orkut, and LinkedIn make these kinds of social communications much easier. It is possible to "mail" all of your contacts (though in general this email remains within the website, and must be accessed through the website), or participate in a forum – a common message area,

like a bulletin board. Other sites, such as LiveJournal, allow you to maintain a "social blog," which allow your friends to post comments in response to your blog entries. These services provide methods of sharing information across a "circle of friends" far more efficiently than electronic mailing lists – which was the only way to manage this kind of communication before digital social networks.

Yet, for all of their potential, very little novelty has come from digital social networks. They are, without doubt, the most intriguing of the web technologies of the 21st century, but they are consistently the most frustrating. They are almost pathetically underutilized. Any digital social network could be doing a lot for its members – because it knows a lot about them. Instead, these networks sit idle until you hit a web page, post a message or update your network of friends. These networks, like much of the web itself, are basically static entities. They do not have goals.

Compare them to Google: Google (and the other web indexing services) is continuously scouring the entire Web, looking for pages and links, using this constant stream of information to improve its ability to help its users find what they're looking for. Google has a goal: to satisfy that goal its designers have turned their web service into an active, ever-hungry beast.

Now consider Orkut, the digital social network created by a Google employee (Google grants employees 20% of their time to use in "skunkworks" projects), and which is hosted on Google's massive array of servers. Orkut is a rich and full-featured digital social network, but, unlike Google, it remains static and quiescent until activated by one of its users. Why is this? It's likely because the programmer behind Orkut saw social networking as yet-another-web-application, and not as an entirely new class of Internet service.

Digital social networks, like search engines, are the electronic equivalent of sharks: if they stop moving they die. Just as a search engine constantly consumes data about the web, a digital social network needs to constantly consume data about its members. Yet not a single digital social network actually does this. It's not that this would be difficult to implement (privacy issues aside), but it requires a small

yet profound shift in emphasis: static digital social networks must be reimagined as dynamic data gatherers. A digital social network is not just a database of photographs and records which point to other records. It is a digital representation of your living self, and as such should be capable of absorbing as much of the continuously generated flow of your personal data as you care to feed it. In the 21st century we are *always* producing data relating to ourselves; this data ends up in databases everywhere but only very rarely in our own hands. We leave extensive records of our passage through the world, but barely benefit from them. That's just plain stupid and wasteful – particularly now that we have an infrastructure in place to monitor and learn from this data stream.

There is one missing piece of the puzzle, which comes from the dissonance between the at-present fixed nature of Internet-connected computers and our actual activities as mobile and dynamic beings. One reason that digital social networks are not yet dynamic is that computers can't follow us around, they can't be the constant monitors of our lives. Or, rather, that used to be true. Since November last year I've been carrying around an Internet-connected computer of substantial processing power – but it's disguised as a newfangled mobile phone. My mobile has as much RAM and computer power as anything I had on my desktop until 1996, and connects to the Internet at not-quite broadband speeds. So I do have a portal to the Internet which is constantly with me, capable of being the attendant observer to my movements through the world. In order for digital social networks to become truly active, they will need to become mobile digital social networks.

This is actually a subtle paradigm shift in how we view, build and use digital social networks. In order to foster some clarity, let me provide three examples to show what digital social networks should be doing for us, right now:

Example One: Spam Filtering

Probably greater than half of volume of electronic mail is unsolicited "spam". Although amazing efforts in spam filtering have been made over the past few years, the creators of spam are continuously evolving responses to the selection pressure of the spam filters. Spam is an evolutionary treadmill; both spammer and spam

filter are constantly mutating into more and more powerful forms. Spam filters rely on an ever-broadening set of techniques to winnow the electronic wheat from the chaff. Spammers, like viruses, are constantly probing the spam filter's immune system, searching for weaknesses.

Electronic mail is quite literally a data stream which is flowing in toward us, and out from us. Consider what could be done if that stream passed through an active digital social network. It could check incoming mail against the hundred-thousand contacts in your broadest social networks; this in itself would provide a powerful check against unsolicited email. It is not the only technique that would need to be used to protect us against spam, but it is certainly an important one, and no digital social network yet offers such a filter as a service. Once again, Google, which runs Gmail as well as Orkut, would seem a natural place to integrate a digital social network with a spam filter, so perhaps they'll get there first – but they'll have to commit the resources to turn Orkut into an active entity – this means lots more servers and lots more bandwidth.

We all receive unsolicited email which is not spam; the trouble with most spam filters is that they trap this email – many times we don't even know we've received it. So we pay a price for spam filters – false positives. It should be possible for us to identify our personal digital social networks within our transmissions, so that the recipient digital social network can perform an validity check. If the email looks legitimate, it should evade the trap. This capability could be inserted invisibly into the email headers, or it could be added as a signature. Either way, the user need never be aware of it.

This example won't work well unless we have some form of interoperability between different digital social networks. Because there are a profusion of these networks – each with different strengths and weaknesses - this is the kind of "must have" service which could lead to the formation of a body of standards around FOAF (friend-of-a-friend) protocols, allowing digital social networks to interoperate, much as different email and web clients do.

Example Two: Media Filtering and Search Engines

Because acquisition of media is now driven by the consumer, we are entering an age where a practically infinite supply of both print and audiovisual media are becoming available. How can one person sort through millions of opportunities for experience? Search engines will undoubtedly play a role, and recent moves by both Google and Yahoo! to include video searches into their services points this way. But that's far from enough. Although a broad search may help you find something of relevance, it doesn't tell you anything about the quality of the experience. Quality is in the eye of the beholder, and that makes terrifically difficult to implement within the context of a search engine, since, by definition, a search engine casts a forgiving eye – existence is a proof of quality to Google, though its page ranking does attempt to rank relevance by a sophisticated and ever-evolving set of metrics.

Several years ago Amazon acquired the "Firefly" technology, developed by Patti Maes' Software Agent Group at the MIT Media Lab. Firefly creates recommendations drawn from "taste demographics." Amazon notes your purchases, and keeps these in a database; it compares these against the purchases of other individuals who have made similar choices, then combines these data sets to make recommendations for other items you'll probably be interested in. It's simple, and reasonably effective. Yet this is all automated; no algorithm, however intelligent, can understand your tastes and desires – particularly if the number of possible choices has grown beyond all bounds.

In nearly every digital social network personal profile, you are asked to enter your favorite films, books, albums, and television programs. In Friendster and LiveJournal you can click on these profile entries and quickly find others who share your tastes. As they stand, these lists of tastes are static; they reflect your tastes at the moment you created your profile. But most users don't maintain that profile with any regularity. It's too much work. Once again, we're drawn to the fact that managing a digital social network is a synchronous task. You have to be in front of a browser, and have the time to spend gardening your digital presence. Most folks won't bother.

On the other hand, if we could make this feature both easy and omnipresent – by putting it into that constant companion, our mobile phones – we'd be able to feed our digital social networks all sorts of data about our media experiences with ease. So let's presume we've implemented such a system (as I have already demonstrated in the LiveRecord prototype Mobile Java application) how would we use it? We already rely on our real-world social networks to find out which movies are worth watching, which TV shows can't be missed, which books are real page-turners. This is informally known as the "water cooler" effect, or "word of mouth". Word of mouth requires either physical proximity or a desire to communicate about media experiences. So a lot of potential information for the word of mouth networks gets dropped on the ground, because the opportunity to share it never arises.

When mobile digital social network become commonplace, they will constantly record these moments of quality - good and bad. We could then expect our social networks to make highly reliable recommendations. You know your friends' tastes, and you know who to trust for their recommendations. No person is a media polymath; different people have different strengths – which is to say that your friends share your tastes in part, but never wholly. Only by harvesting the quality from an entire social network can you gain a broad perspective. When this harvest of quality is combined with a more prosaic search engine, we will have a tool suited to separate order from chaos.

In order to be completely seamless, the acquisition of moments of quality should be mostly invisible. TiVO already notes what TV programs you've watched, and how often you've watched them. That information should be fed into your own digital social network. iTunes records how many times you've played each track in its database, and allows you to give it a rating – 1 to 5 stars. That information *also* needs to be fed into your digital social network. When you leave a movie theatre, your mobile needs to remind you to rate the film. And so on. The more information about quality each member of a digital social network feeds into that network, the more likely it is to offer quality recommendations.

Example Three: Knowledge Augmentation & Knowledge Force-Multipliers

In the 1960s, the one of the great minds of the present era of computer science – Douglas Engelbart – struggled to develop the idea of a computer as a "tool for thinking". In the wake of Vannevar Bush's influential essay, "As We May Think," an entire generation of engineers and visionaries sought to bring the radically empowering technologies of information processing to the individual. By the late 1960s it had just barely become possible, and, at the ACM Winter Conference in December 1968, Engelbart demonstrated his crowning achievement, the NLS (oN-Line System). For the NLS, Engelbart and his team developed nearly all the elements of modern computing, including the GUI, the mouse, hypertext and networked videoconferencing. Engelbart was so far ahead of his peers that 25 years had to pass before Sir Tim Berners-Lee could realize much the promise of the NLS in the World Wide Web.

We are living in an era of enormous knowledge augmentation, which is the gift (and, equally, the curse) of the Web. When we want to know something, we google it. The scope of the questions which go unanswered has become radically smaller; certainly if anything has a factual answer, we are no more than a few moments away from an answer, provided we have access to the Internet. (The search engines need to devote enormous resources to making themselves mobile-friendly; using a search engine from a mobile phone is still fairly unfriendly.) But there are facts and there is understanding; I could read a lengthy article about Japan but still have no sense of how that culture's famed politeness is literally encoded into the Japanese language. Factual knowledge is not functional knowledge; it does not bring understanding. "Knowing is doing, and doing, knowing." (Maturana & Varela, *The Tree of Knowledge: The Biological Roots of Human Understanding*, 1980)

We live in a world where we have a grasp of the facts, but no practical sense of their meaning. Knowledge augmentation is not an endpoint but a signpost; knowing the facts is not enough. The next step following knowledge augmentation via hypertext is knowledge *amplification* via digital social networks. Engelbart

understood this, which is why the NLS is an on-line system. We already have many ad-hoc forms of knowledge augmenting digital social networks; mailing lists and journaling websites (such as Slashdot) allow loosely-linked individuals to contribute to the greater wisdom of all. But the most outstanding example of the knowledge multiplying effects of a digital social network is – without doubt – Wikipedia.

Wikipedia is great the free and open source encyclopedia project, which began in early 2001, shortly after (and, in some sense, as a response to) *Encyclopedia Britannica* put up a subscription-based walled garden around its content. It started off as a closely-knit community of core contributors. These few (there are perhaps 60 of them world-wide) are the "gardeners" in the "fields" of Wikipedia. It is the eternal vigilance of this "lucky few" which keeps Wikipedia from descending into chaos. But the half million English-language articles in Wikipedia were not written by a mere sixty individuals. Anyone who visits Wikipedia can freely edit an article, or add a new one. And hundreds of thousands of visitors – myself included – have done so. (I am at present rewriting and extending the Wikipedia article on X3D, which, given my history, is a natural thing to do.) Wikipedia is an anonymous digital social network; there are no overt signs of authorship on Wikipedia articles – in that Wikipedia resembles a Mediaeval cathedral with its legions of talented but nameless craftsmen. Every article added contributes in some way to the greater good of all of Wikipedia's users, and this has become a self-sustaining virtuous cycle: as Wikipedia becomes better and more complete, it becomes more useful, and hence more likely to become even better and even more complete.

Yet Wikipedia is only an intermediate step between knowledge augmentation via hypertext and knowledge amplification via digital social networks. Wikipedia is more-or-less static, representing a compendium of factual knowledge; in this it reflects its roots as a Web-based medium. It is a collection of facts – detailed and broadly available – but without the benefit of human experience. For the insight of human experience, we need to rely on human beings.

This brings me to a basic question, one which I think will become the foundation of many successful projects in the next several years – why can't we google our

friends? Our friends have are vast libraries of particular, specific and very human understandings. Let me put it another way: if we complete the execution of the ideas embodied by Engelbart's knowledge augmentation projects, won't that make us dynamic and indispensable actors within the knowledge network?

What do I mean by this? Consider the specifics of an implementation: first, you need to be continuously available to the knowledge network – both to make requests of the network and to receive requests from it. This means it must be a mobile digital social networking application. Second, the software underlying the application must handle most of the request routing invisibly – you shouldn't be disturbed by questions which lie outside your area of experience. Finally, it must be responsive – an answer is valuable in proportion to its timeliness. If all three of these requirements can be satisfied, there is the potential to create something far more potent than Google or Wikipedia – a truly human form of knowledge augmentation, reflecting the whole variety of experience and wisdom which can be harvested from your social network.

If you think you've seen web crack before – and the Web, Napster and Friendster all qualify as highly addictive experiences – they're nothing at all to the addictive potential of truly potent knowledge amplification. As knowledge amplifying mobile digital social networks are deployed, the individuals employing them will benefit from hugely extended capabilities. These individuals will know the facts *and* how to put those facts to work. These knowledge amplifiers will quickly become indispensable features of daily life. And the users of these networks will find it hard to work effectively without them. Just as the web transformed human knowing, these active digital social networks will transform human doing. In short, they will transform us into *hyperpeople*; we will become necessary and irreplaceable nodes within a network that has moved beyond fact and into understanding.

Those are the three examples, moving from the why-haven't-we-done-this-yet through to won't-it-be-nice-whe. Active digital social networks will become an intense focus of research and development over the remainder of this decade, because they promise such an enormous extension of our innate capabilities by marrying our inherent natures as social beings with the processing capabilities of

digital networks. This is the step *after* the mobile phone, this is the step *after* the web; this is the next place our networks will take us.

Part Three: Fork You!

The future belongs to those individuals and organizations with the most powerful digital social networks. Indeed, the present already belongs to those with the powerful social networks, but – like media hyperdistribution – the power of digital social networks is spreading; they will be available to every individual. We can expect to see a new form of natural selection to rise up, as social networks compete against one another in the pursuit of goals. Less effective social networks will fade into extinction; successful ones will grow, and replicate their successes; networks may even choose to cooperate in symbiotic relationships, each providing something the other needs. But social networks have their limits.

Anthropologists have been studying human social networks for over a hundred years, and have come up with a rule of thumb known as the Dunbar Number: the maximum size of any long-term sustainable social network seems to hover around 150 people. Indeed, if you look at profiles on Friendster and Orkut, you only very rarely see anyone with more than 150 entries in their social network. The Dunbar Number establishes the maximum size of a tribe of hunter-gatherers. At numbers greater than 150, you can't know everyone; the social contract begins to fray. Tribal humans often solved this problem by bifurcating their social network, sending out a "colony", and keeping both new tribes below the Dunbar Number.

Does this mean that our digital social networks are inherently limited by the Dunbar Number? The indications are that the Dunbar Number may have a lot to do with our biology, our ability to remember individuals. Yet we already use computers as external memory prostheses, to remind us of meetings and numbers and tasks and all sorts of minutiae which would overwhelm our limited memories if we had to track them all without any augmentation. (Consider: who remembers phone numbers anymore?) If our computers could also keep track of our social interactions – something analogous to conversation threads on a bulletin board, but far more sophisticated – it is possible that we could burst through the Dunbar

Number, and form effective social networks with thousands, perhaps even tens of thousands of others.

In order to do this, we'd have to train ourselves to context switch, to move from membership one social network to membership in another seamlessly. We will need to be able to "fork" ourselves, moving through a series of social networks as a computer might load a series of applications. This may not be an easy thing to do – but it is worthy of research, because unless we can break through the Dunbar Number, we will soon see an effective upper limit to the usefulness of our social networks. Still, so much knowledge augmentation will be achieved even within the constraints of the Dunbar Number, that perhaps we'll need to bump up against that limit before we see how to move beyond it.

Everything I've outlined in my talk today requires a substantial amount of new software; there's a lot of work to be done. I'd like to take this opportunity to call for an effort to build a base of open-source software which will create a robust, distributed and pervasive mobile digital social network. The ground of that network is, in raw terms, a "human operating system" – and, like computer operating systems, shouldn't be owned or controlled by any one entity. It needs to be a collective human project in order to provide the greatest benefit for all the individuals who will use it.

So, can we form a social network – starting here, today, right now – which will harness our own abilities to amplify our own understanding? We have the resources, we have the technology, we need only the will to move forward. The benefits are as obvious. The future belongs to *hyperpeople*.

Mark Pesce
26 – 28 March 2005
Sydney