

Transformation

Good afternoon. I'm honored to be speaking with you today. This year marks a very scary milestone for me – 30 years since my graduation from Penn, and my transition to graduate studies at the Annenberg School. Annenberg was a formative experience for me in many ways – but perhaps the most directly relevant to this talk today is that it paved the way for my career in digital media, a career that began when my colleague and friend, John Carey, stole me away from the Ph.D program at Annenberg in 1979 to help him manage a research project at New York University in the new area of interactive media and its potential as a public information resource. The late Barry Cole, one of my professors at Annenberg, had introduced me to these emerging technologies, and I sensed at that time how potentially revolutionary they would be. So in a very

real sense I owe my career – or at least the potential of it – to the Annenberg School, its faculty and, of course, to Ambassador and Mrs. Annenberg, who made it all possible.

Today, I've been asked to speak about the historic transition occurring in media as we evolve from analog to digital delivery, and to do so from my perch as the senior digital executive at The New York Times Company.

Long ago, sometime after I joined NYU, it stuck me that we were witnessing a long historical arc, one that has its roots, in part, at this University when, during the 1940s, the Eniac was created.

This long arc, I thought, would last about 100 years. It would begin with the invention of digital computing and it would reach fruition after all the backend systems and front-end user technologies were fully digital. I think this is still true --- that the transformation in media today is not new, but only becoming visible to many people as they experience these changes in their daily lives. Windows is

part of this long arc; as is the ipod; as is Google. But the arc began long before the Worldwide Web protocols were written, and even before the first consumer technologies appeared during my Annenberg days in the late 1970s.

So the topic you've asked me to address – the transformation of publishing – can best be done in this historical context despite the temptation to characterize it as a recent phenomenon.

This history, I believe, is divided into four great phases. We are now just beginning to enter the third phase, so I think we're a couple of years past the mid-way point in this 100-year transformation.

- The first phase, which began after World War II with the invention of digital computing, and first reached commercial viability in 1960s, is characterized by the aggregation and

sorting of the world's knowledge, and of making this knowledge accessible to specialized groups of users, principally for commercial, government and academic purposes.

- The second phase, which had its seeds in the late 70s, but began in earnest in the early 90s, extends this aggregation and sorting function to a global consumer base. This second phase is characterized by the development of interface and standardization.

- The third phase, which began out of the ashes of the dot-com bust three or four years ago, is characterized by creativity --- that is, the birth of wholly new forms of publishing media that borrow from the elements of the first two phases, but bring users new and original ways of communicating, as well as a new class of competent creators.
- The fourth phase, which I think will begin in about five years, will completely and radically transform the institutions – corporate, academic and governmental – that support the information economy. These transformations – which may seem apparent to some even today – have only just begun.

Underlying these four phases are the two driving forces of technology innovation: Moore's Law, the doubling of transistors on an integrated circuit roughly every couple of years; and

Metcalf's Law, the power or value of a network increasing in proportion to the square of the number of nodes on the network.

From the earliest days, librarians have been tasked with storing, indexing and making available information to users and readers. In this context librarians are not simply mechanics, for they are responsible for ensuring the quality of the information that is ultimately retrieved. Early computers were, in essence, sorting machines that were designed to improve upon the mechanical sorting machines that existed prior to 1945. These were powerful new tools put in the service of specialists, who required training to operate them.

When Vannevar Bush published his seminal article on the “memex” machine, entitled “As We May Think” in the *Atlantic Monthly* that year, he was really the first to fully envision and articulate a generalized information retrieval system, or fast access to all the world’s knowledge.

The roots of electronic publishing, and what I refer to here as its first phase, grow directly from this early history. The principle job that electronic publishers were hired to do was to aggregate and sort information, and to present the results of this sorting to specialists who would ultimately then serve an end-user population willing to pay a very high premium for this powerful capability.

Not surprisingly, the government was the first to grasp this potential, and, by 1953, the first information retrieval system had actually been implemented by Harley Tillet on an IBM 701 at the U.S. Naval Ordnance Test Station in China Lakes, California.

Three years later, the first large-scale online retrieval system – SAGE – was developed at M.I.T for the U.S Air Force.

For the next 10 years, retrieval systems were systematically developed and improved upon, principally in academia, but certainly in expanding application areas, where aggregation and sorting were of great importance.

For example, in 1962 John Harty extended the idea to health law information at the University of Pittsburgh's Health Law Center.

Two years later, Harty demonstrated the capability to extract health-related statutes and related legal data at the ABA. That same year Ohio Research Service Incorporated built an online service for the retrieval of case citations.

So by the mid-1960s, the stage was set for the commercial development of electronic publishing. In 1967, the Ohio State Bar Association contracted with the Data Corporation to develop a full-text retrieval system. In one of those wonderful quirks of history, Mead acquired the data corporation in 1968 to get its hands on some proprietary printing technologies. This led to the creation, in 1970, of a subsidiary called Mead Data Central headed by my old friend Don Wilson.

The rest, as they say, is history. Interestingly, The Times actually entered this scene in 1974 when we created The New York Times Information Bank to produce the first publicly accessible online newspaper index and abstract. In a stunning act of bad timing and lack of foresight, The Times then sold our INFOBANK library to Nexis in 1983, and entered into a perpetual, exclusive content license that was only modified when Reed Elsevier acquired the business in 1994.

This early history of electronic publishing, from memex to the invention of the Web, shows the steady progress of content digitization, aggregation and sorting that characterizes this first phase, which – of course – is still highly relevant today. Whole groups of professionals now take for granted advanced capabilities that hardly existed a generation ago.

Phase two in this history began just around the time we were selling our Infobank to Mead.

Now I want to mention Ted Nelson in this context, because in 1965 or thereabouts, he first came-up with the term hypertext, which he later expanded on in *Dream Machines*, published in 1974. In the book, Nelson describes his Xanadu notion quote “as a magical place of literary memory.” While this owes much to

Bush's early thinking about Memex, hypertext is a huge leap. I want to quote Nelson from a BBC interview conducted in 2001:

So, the point was to be able to have a medium that would record all the connections and all the structures and all the thoughts that paper could not. Since the computer could hold any structure in any form, this was the way to go.

Now it's important to keep in mind that Nelson had a media background. He was not a computer scientist. His orientation was film and music. He was in a rock band. He envisioned hypertext as a new way to communicate.

I should be clear at this point that Nelson states emphatically in this BBC interview that the Web was not what he was trying to create, that – in fact – the Web employs a “degenerate” form of

hypertext because – among other reasons – quote “the links only go one way.”

But like so many innovations, the Web was “good enough” to allow for the foundation of a worldwide set of standards. Still, Nelson’s hypertext vision was bubbling up at the same time as other major developments in computing and information retrieval were beginning to occur.

Around this time, we had folks like Alan Kay and Doug Engelbart beginning to create new ways to access information. The mouse, for example, had its public debut in 1968. It, too, was a government-funded innovation --- this time out of NASA. Much innovative work was being conducted at Xerox Parc. By 1974, researchers there had built the first bitmap WYSWYG cut and paste editor. By 1980, the Ethernet specifications had been published.

By the early 1980s, about the time we were selling an exclusive, in-perpetuity license to our archival content, Steve Jobs and Apple had copied Engelbart's mouse and the WYSWYG concepts that came out of Parc and popularized them in the early Mac. A few years earlier, as I said at the outset, my career began on one of the first teletext projects in the country, also funded by the government – the National Science Foundation – at broadcast station WETA in Washington.

For those of you with long memories, you will recall that the technologies of teletext and videotex used the television as the display device for massive amounts of information. Videotex databases were tree structure designs meant to allow ordinary users – not specialist librarians or other trained users – to access information.

By the time America Online was introduced in the mid-1980s, videotex had migrated away from the television, and owing largely to the vision of Steve Case, onto the emerging Microsoft Windows platform. Phase two, which I define as characterized by Interface and Distribution attributes, was now set to popularize information retrieval and computer-mediated communication for the masses.

A few years later, a British physicist named Tim Berners Lee wrote the worldwide web protocols at CERN, another government organization, built on top of the Internet, which had been funded by the United States government – DARPA - as a redundant communication network. This enabled the proprietary interfaces pioneered by Engelbart and others, later popularized by Jobs and Case in closed systems, to be extended through open and free standards through the web.

Importantly, in Berner-Lee's vision scientists would use these open protocols not just to retrieve documents, but to link with one another. Hence, the web at its outset is designed to foster social interaction, not just information retrieval.

So phase two built on phase I, taking the essential aggregation and sorting functions enabled by computing, and extending them to the masses through breakthrough interface design and open internet standards.

Phase two is just now reaching commercial viability. Fourteen years after the Web is introduced, we now have two Internet information media companies at scale – Google and Yahoo – that extend aggregation and sorting to a worldwide audience. The demons of the dotcom era have been exorcised as these two companies are realizing many of the concepts that were introduced during that time:

- A highly accountable, very targeted advertising model that creates great margins for the winning players,

- Information services that are “good enough” to satisfy massive numbers of users to fuel this advertising business, and
- Increasing returns to scale and continuous product improvement that could eventually kill or absorb many of the incumbent phase I players.

In the editor’s note to Vannevar Bush’s seminal 1945 article in the *Atlantic Monthly*, we read, “*Now, says Dr. Bush, instruments are at hand which, if properly developed, will give man access to and command over the inherited knowledge of the ages.*”

“The inherited knowledge of the ages”

This is the key phrase that has characterized digital media from the first description of MEMEX to the 300 million worldwide users of Google last month. The inherited knowledge of the ages includes the iTunes catalog, The Times archive and the collective wit and wisdom of all those MySpace users.

While Phase two has plenty of growth left – decades of growth, in fact, we are now beginning to realize the promise of memex. The twin attributes of early online services – aggregation and sorting – are becoming increasingly available to a worldwide base of users, not just in the context of documents, but increasingly through access to all media. Walled gardens, corporate or governmental chokepoints and top down media structures are dissolving. The revolutionary forces that we envisioned in the late-1970s are reaching global scale.

As we all know, this has already disrupted the commercial music business and – in what I referred to in an opinion piece in *Ad Age* in November, 2004 – portends a “control shift” in which all media move gradually under user control. Such things as the 30-second commercial break – the underlying financial driver of the television industry today – are increasingly under pressure as a result. The structures that formed the world of analog media are breaking down.

As I stated at the outset, as phase two begins to mature, I am postulating today that we are now on the cusp of a new phase in this historic transformation, one characterized not just by the aggregation and sorting of the world’s knowledge, but now by the creation of works that are specific to a new expressive medium. This will fuel much greater growth than the two prior phases, and create a new set of businesses that will sit alongside the great information and entertainment media of our day.

What are the essential criteria that define successful new media, from newspapers and books to television and videogames? I think there are three main ones:

- First, the medium must be large, global and spawn a new profitable industry.
- Second, the medium must be expressive. It must delight people on an emotional level. It must become a regular part of their life experience.
- Third, the medium must ultimately engender a new collective class of creative people. Think of film, with actors, directors and set designers; or videogames, with art directors and

programmers; or newspapers, with reporters, editors and photographers. These people work together as professionals once the medium can sustain this new professional class.

Often in new media, there is an early period of false starts. We have witnessed many in interactive media over the 20 years, from interactive movies where the audience picks alternate endings, to failed experiments in interactive journalism.

But we also see new forms taking root. These new forms blend a unique stew of audience input, amateur content creation, the editing of outside content sources and other attributes. They evolve – in part – from the forums and chats that ran on CompuServe twenty-five years ago, in part from email, in part from traditional publishing, in part from search, in part from content syndication. In a word, they are new.

An early version of this is evident in the blogosphere. Certainly, the better blogs delight their audience. Millions of people each day read a web log. Even though blogs began as a form of amateur publishing (all new forms begin with passionate amateurs because a professional class does not yet exist to create in the new medium), we are now beginning to see more professionals entering the fray, and the form seems to be developing a division of labor in some instances.

So we satisfy two of the new criteria at least.

The last one – and often the most challenging – is our first criterion. The form must evolve into a profitable commercial enterprise if it is to sustain the professional class required to attract mass audiences. Indeed, it is possible that blogs will remain a vibrant amateur medium, much as CB and ham radio have over the decades. But I think these are more than telephone conversations.

We've invested in a company called Federated Media Publishing that is a kind of studio for bloggers, supporting them with backend publishing and ad server systems, advertising sales support, audience research and marketing. Think of Federated as a kind of early studio for the best bloggers.

Blogging is only one example of this third phase now coming to the fore. Others include:

- Newly created web video studios like Heavy that are designed to create new entertainment experiences unique to the attributes of digital media.
- Social media hubs like MySpace and YouTube that act as platforms for emerging creative talent, using low-cost production technology and circumventing traditional distribution systems.
- New packagers that take these newly created media types – blog posts, comments, web sites – and aggregate them along geographic or topical lines, forming new intermediaries in competition with traditional content packagers like newspapers and broadcasters.

These are all early examples of this new phase – this phase III where a tremendous flowering of creativity will occur, all made possible by continuous innovation:

- Infrastructure that allows millions of ordinary users to access a broad array of media types, from simple text to video and animation. This infrastructure includes broadband and very powerful personal computing devices.
- Dramatic consumer behavior shifts – both at home and in the office – placing connected devices in the center of the consumer’s day. A control shift from the center to the edges of the network.
- A whole array of creation technologies – from low-cost, digital video to widely available, advanced publishing platforms – that enable new ideas and formats to bubble-up, initially in the hands of passionate amateurs, but over time subsumed by a new professional class.

- On the other end, core infrastructure that greatly enhances the user experience, including inexpensive, highly portable, very powerful networked devices.
- Profitable Phase II content businesses like Yahoo, populated with people who have new competencies and resources to support experimentation among a professional class of collaborating colleagues.

Like a primordial pond bubbling with the elements of life, our environment is now ripe with forces that might shape a genuinely new medium.

And so it is in this historical context that The Times transforms itself. Indeed, when we started nytimes.com in 1995, we did so with a classic phase II vision in mind. The early web site took the aggregation and sorting mechanism that had existed for decades,

and thus made our content available to anyone with a PC, a web browser and a modem. Almost instantaneously we became a global brand, reaching people in virtually every country on earth. Nothing particularly new was being created. Our success came from a keen understanding of interface design and the principles of hypertext laid out by Ted Nelson thirty years prior. Over the next five years, NYTimes.com grew to become the largest newspaper-owned web site in the world, based on these simple principles. By 2003, we were extending these notions of sorting and aggregation more deeply into the arena of consumer web publishing. Yahoo had pioneered many of these usability attributes. Going far beyond the notion of an online newspaper, we worked first in the area of movies and theater, a place of traditional strength for The Times.

We took 90 years of film reviews (yes, The Times' first film review appeared in 1913 for Quo Vadis) out of the paid archive and made them available through an easily searchable data base.

We created a landing page called a “movie detail page” that allowed users to find everything from the review and cast, to what other nytimes.com users thought of the movie, to where to find it in your neighborhood.

Tony Scott, one of our film critics, created Movie Minutes, a series of short web videos that complement his written reviews.

For the first time we allowed our users to write reviews of movies and to rate them. We hired a product manager who had experience in both the entertainment industry and online publishing to oversee the whole effort.

Today, nytimes.com movies regularly has over 4 million monthly visitors. It has grown into one of our most important revenue generators.

But, again, with the possible exception of Tony Scott's web videos, we creating essentially nothing new here. We are simply using the aggregation and sorting attributes of electronic publishing at a new state-of-the-art.

By 2005, however, we sensed that this aggregation and sorting phase was giving way to something new – an era characterized by new forms creativity and wholly new ways to tell stories. This created a two-fold imperative:

- We had to make the transition to Phase III.
- We had to organize to do so.

That meant evolving from the innovative, small skunkworks that we'd established in 1995 to manage the site, to its full integration across the entire enterprise --- we needed to tap the creativity of the entire newsroom, and to make nytimes.com as important a creative

outlet as the printed product. This is a huge organizational and cultural transformation.

We made this decision in late 2005 and reorganized in early 2006.

The first clear example occurred out of a tragedy in October of that year.

Last October 11, a small plane crashed into the 40th floor of an apartment building on the east side of Manhattan. Prior to our reorganization:

- Assignment editors would have sent reporters and photographers to all the right places.
- The Metro editor would have convened a small group to draw up a story list for the next day's paper.

- And, yes, somebody would have told those people over at the web site, and the AP story would have been put up online followed by a Times story a few hours later.

Classic Phase II approach.

But here's what actually happened last October:

- Assignment editors alerted continuous news and called the video unit and the graphics desk and mobilized the database reporting team.
- The newsroom immediately put an alert on the web site within a few seconds.
- Assignments for reporters and photographers included instructions to gather audio, which appeared on the site within a few minutes.

- A notice went up on the site soliciting photos from users in the vicinity who might have had camera phones at the time of the crash.
- The database reporting unit executed a lightning fast search that discovered that the owner of the plane was Cory Lidle, a pitcher for the New York Yankees, and conveyed the information to the metro desk, which confirmed that Mr. Lidle had died in the crash.
- The alert about Mr. Lidle was picked up, with credit, by television and radio stations.

The coverage on the web was faster, richer and deeper than anything we've done before. Within a few hours we were offering slideshows, audio, video, pictures, a multimedia graphic so remarkable that it drew more than 100 compliments from readers. Most importantly, we'd made this magical transformation from phase II to phase III. We were beginning to create *for the new*

medium in the heart of The New York Times newsroom. It took almost exactly ten years to reach this milestone.

The transformation from print to digital at The Times was now moving along the natural historical arc from the aggregation and sorting of our archival and daily content to the creation of something new designed specifically for the attributes of the new form – a rich mix of digital media types, all under the user’s control, and with the user able to contribute.

The natural next step in this evolution will be found in the principles that underlie this new era:

- A more porous membrane between the institution and the outside world. One can begin to see this in what we call “topic pages” --- initially born out of archival sorting of The Times archive into 10,000 topical areas, but one day incorporating content from many other high quality sources

and curated by credentialed professionals beyond the walls of The Times newsroom.

- A robust identity management system that provides transparency to a broader Times community; perhaps one that forms an early version of social networking rooted in content and knowledge exchange.
- Storytelling that naturally incorporates all appropriate media types into a rich digital stew, formulated for a global audience hungry for high quality news and information.

In re-reading Bush's 1945 Atlantic article, I tried to think about what foresight it took for him to invent the memex, given the technologies of the day. He envisioned a desk-sized device where *"books of all types, pictures, current periodicals, are thus obtained.... If the user wishes to consult a certain book, he taps its code on the keyboard."*

Twenty years before Ted Nelson comes-up with hypertext, and 50 years before Berners Lee creates HTML, Bush is describing, "wholly new forms of encyclopedias... ready made with a mesh of associative trails running through them." He is envisioning a version of wikipedia in 1945!

While I wouldn't even presume to reach Vannevar Bush's middle calf in stretching to think about the emerging medium I refer to as "Phase Three," I would like to paint a picture that seems to contain a lot of the "stuff" that's coming together to make this happen.

There is no primary media format in this new journalism. It is all media combined. All of the assets available for storytelling are seamlessly available and can be offered for the sole purpose of telling the story. Aggregation and sorting are rapidly becoming media independent. In part, this is because the means of production are advancing to allow creative teams to work seamlessly across media --- and to tell stories in new, non-linear ways.

This sows the seeds for a new kind of storyteller --- one who understands how to combine all the best attributes of formerly independent media types into a new whole. This new storyteller writes in a non-linear format; that is, the storyline runs both as narrative, and as tapestry of historical references; this forms the basis of the hyperlinked news organization. We see this in nascent form in the blogging world today.

In this way, aggregation and sorting are woven into this vision.

Bush's vision was that of the lone scientist sitting at his memex accessing information. Phase Three is as much about social networking as it is about information retrieval.

Users will be able to associate content from a wide variety of sources, including filtered access to other users, through a common interface. Online communities will be available to discuss any “article” at any time of the night or day. A user lands on the article and discussions about that article are continuous. News reports become a focal point for social networking. Again, we see this bubbling-up in web logs today.

In this vision, two very important boundaries which all of us have grown up with, know and understand change. The first is authority. In a newspaper or in a television news broadcast, the journalistic experience is one-way and bounded. In this new experience, sources range and are scored.

A professionally crafted restaurant review will –in my opinion – bring the same kind of authority to the user in the digital world as it does today. But this review will be augmented by a range of opinions from others; ranging from other professionals to bloggers to audience. We can already glimpse these scoring mechanisms – and this is just one example – in our rate and review area in our Movies section. As better reputation and recommendation management systems emerge, readers will “empower” new kinds of creators who may emerge from outside of traditional institutions.

New interfaces will develop that can support a multiplicity of “views” into a particular geography or issue. Today, production technologies and bandwidth limitations constrain our “view” to a linear tightly edited broadcast format. Even CPSAN, with several fully dedicated channels, can only provide us with one view at a time.

But imagine taking a world like World of Warcraft – designed for massive numbers of videogame players – and apply it to the real world, where the players are reporting from all corners of the planet. This is a vibrant, interactive real-time view of the world. Perhaps we are glimpsing an early version of this in Second Life.

Users in this context can zoom into the ongoing storyline taking place in dozens or even hundreds of locations. In this context, there is not a simply John Burns reporter in Baghdad. There is a kind of ongoing John Burns channel that brings with it a continuous record. Many people make the mistake of thinking that a more distributed news capability obviates the need for professional journalists.

On the contrary, the judgment, synthesis and analysis journalists do become even more valuable as the information flow and

complexity increase. The role of the Fourth Estate remains at the center of our Democracy. It must, or God help us.

However, the range of this journalistic flow extends in several directions at once, portending new aggregation schemes. Not only do we have better access to global resources, but we also have an almost infinite array of niche subjects that can be mined. The limitations of the news hole are blown apart.

In this context, the audience becomes far more integral to the brand experience. We call our readers the “knowledge audience” because our research shows very clearly that the psychographic attribute that ties them together is an emphasis on education, travel a love for culture and lifelong learning.

Newspapers have always created a kind of sociology around them, mostly community or geographically defined. Most newspapers

have a Letters column, for example. But in this new vision, the social extensions of the journalism become deeply embedded into the product itself.

Reputation systems are central to this idea, because while deception is possible even in the most controlled context (think about Jayson Blair at The Times), the casual online encounter is rife with potential fraud and spam. Bad actors abound. That's why highly ranked sellers on eBay can charge more than sellers without similar rank.

People often ask me why we register our users --- identity and reputation go hand-in-hand. Answers can come from anywhere in the network, but users need a way to verify the source. Ultimately, journalism is about trust. In this new product vision, the aggregator's brand must remain trusted, but the reputation systems that surround and augment it must also take into account these enhanced social effects.

Ultimately this new form of journalistic expression extends in several directions at once:

- The means of production come into existence to empower a new creative class to tell stories that fluently combine all media types. We can this in early form with our Cory Lidle example.

- This storytelling is done in a non-linear, hyperlinked format, as databases are woven seamlessly into the fabric of the story, unencumbered by news hole or broadcast time frame. We see this in nascent form on our topic pages.
- These stories are continuous and are aggregated into a kind of universe where users can join “views” from around the world. These views are an audio visual window into place and topic, with multiple sub-views branching off the main view. The sub-views can involve different places, analyses and discussions. This is happening in its own way at Wikipedia.

- The role of the aggregator is to build and maintain this news universe. To paraphrase photographer Henri Cartier-Bresson, the heart will always beat faster in some places than in others. The aggregator must always be responsible for this editorial judgment. But now he can draw from so much more.
- Finally, the boundaries of the medium extend into social networks supported by online reputation systems. The character of a publication is respected as much for its audience interactivity as it is for its content.

I started by describing four phases --- and by saying that the fourth phase, the one in which devices, operating systems and interfaces come together to create massive institutional dislocation --- is five years or so away

Perhaps the iPod and iTunes together form a kind of pre-historical view of this change, as these new interfaces and business models wreak havoc on the traditional music industry. But I'll leave phase IV for another time. After all, this is 100 year arc and we have plenty of time left. If you invite me back ten years from now, it will be in full flower.

Dick Brass, a retired Microsoft Vice President and friend, once postulated that the last newspaper edition of The New York Times would be published in 2016. I think he's missed the point. I'm not sure that our new, digital journalism will ever really replace the extraordinary package that is the printed newspaper. No, it does something different, and quite unique.

The more important strategic point is that as participants in this rapidly changing landscape we imagine new possibilities and prepare to recognize the sparks that lead to big fires and fuel them when they ignite. In this way, we will be far better prepared for what should be the most exciting and liberating era yet to come for journalism.

Thank you very much for listening.