MEMORANDUM

06 August 2000

TO: COPA Commission
FROM: Lawrence Lessig
RE: Proposed legislation to zone minors from material deemed harmful to minors

---

As you have requested, I have summarized my views about the trade-offs among various proposals for zoning minors from material deemed harmful to minors in cyberspace. I have drawn this analysis from my article with Paul Resnick, Zoning Internet Speech, 98 Michigan Law Review 395(1999). Any analysis of the constitutional issues raised by these proposals can be found in that article. My aim in this memorandum is simply to outline the alternatives, and the trade-offs among them.

As I said in my testimony, in my view your objective should be to identify techniques to enable parents to protect children, consistent with protecting the values of free speech. In my view, however, free speech is threatened both by bad law, and by bad code. My aim has been to identify a response that minimizes the effect of bad code. I offer Proposal (4) as an example.

INTRODUCTION

To zone minors from material considered "harmful to minors," a system must know the (1) age of the recipient and (2) the content of material the recipient wants to view. If the recipient is a minor, and the content is harmful to minors, then the system should block access; if the recipient is not a minor, or the content is not harmful to minors, then the system should not block access.

To facilitate such zoning, proposals to date have been of two general sorts. First, there have been legislative proposals to require that adults carry identification when they desire to get access to material that is harmful to minors.[1] (I will refer to proposals of

---

[1] The first federal proposal required identification whenever the adult sought access to "indecent" material, but the constitutional standard has only ever justified conditioning access based on whether material is "harmful to minors."

this sort as Proposal (1).) Second, there have been nonlegislative proposals to facilitate the rating and filtering of content on the Internet, thereby enabling parents to block access by their children to material that is harmful to minors. ("Proposal (2)").

Proposals of the first sort have not been successful in federal courts. The burden on adults to carry age-identification is significant; the burden on sites to verify the identification presented is also high. These two burdens have been considered too great in light of less burdensome alternatives. Every federal court to review these statutes has concluded they are unconstitutional.

Proposals of the second sort have also been met with great skepticism, though this skepticism is of more recent origin. Technologies for rating and filtering content on the Internet are inherently flawed. They universally reach beyond the narrow category of harmful to minors material. They therefore facilitate a far greater blocking of access to material than the government's legitimate interests reach. And while this blocking is done by individuals, and not governments, the effect of these proposals on access to controversial speech, even by minors, should be relevant in evaluating the merits of these proposals.

The solutions, in my view, are either proposals that (3) facilitate a less burdensome kind of identification, or proposals that (4) induce a less extensive form of rating and filtering. Proposals of type (3) depend upon systems that certify that the user is a minor, not that the user is an adult. And proposals of type (4) identify simply whether content is harmful to minors, and not anything more.

In the analysis that follows, I first describe proposals (3) and (4). Within each description, I identify the strengths and weaknesses of each proposal. I then describe how each proposal is complicated if the "harmful to minors" standard is different within different geographic communities.

PROPOSAL (3): IDENTIFYING MINORS

Imagine a browser that gave users the option to establish a "profile" that governed the preferences of the browser for that

user.[2] That profile would be protected by a password, so that when the user "logged onto" the browser, he or she would have to supply a password. Once the identity of the user is verified, the browser would then select the bookmarks, and user preferences desired.

Imagine further that in setting up the user profile, there was an option to designate that the user was a minor. If that option were selected, then the browser would not permit the transmission of personal data to a web site.[3] It would also, if requested, certify to a web site that the user was a minor.

Finally, imagine that a law required web sites serving material deemed "harmful to minors" first verify whether the user was a minor by "querying" the user's browser about whether the user was a minor or not. That query would simply be a request to the browser that it transmit whether the profile of the user was marked as a minor; the browser would answer in the affirmative if it was so marked. If the client answered affirmatively, then this law would forbid the server from serving that material to the minor. If the client did not answer affirmatively, then the server would be free to serve the material without legal liability.

This configuration of technological capacity and legal responsibility would facilitate, to some degree, the zoning of minors from material deemed harmful to minors. Browsers are essentially free. The modifications required to facilitate the identification of minors would be trivial. And the software to enable servers to query and block sites based on that code would be relatively easy to implement as well.

Nonetheless, Proposal (3) would impose burdens on Internet speech. In the balance of this section, I describe these burdens. I then describe the legislation that would be needed to move the net

---

[2] While I have abstracted this description from the particulars of any specific existing technology, it is clear that there are many existing technologies that come close to the description I offer here. The Netscape browser permits different user profiles. The Mac OS 9 permits profiles specified at the operating system level. There is no reason these technologies could not be made more generally available.

[3] This is a complicated objective. Certainly it would be easy to ensure the browser itself does not send any of the personal data stored in its preference files. But it would be harder to interpret a web page to determine whether an email address or other personal information was being requested.

in the direction of this configuration. That legislation is what I will describe as Proposal (3).

The Burdens

The burdens of this configuration are two: first, the burden on any site to determine whether its content was "harmful to minors." Second, the risk of misuse of the identifying information that the user of a particular browser is a minor.

The burden of rating material "harmful to minors"

The first burden is no greater than exists under real space laws that restrict access to material harmful to minors, except to the extent geography becomes relevant. (I will discuss this qualification below). Sites offering material that is harmful to minors today must take steps in many states to identify that material, and keep it from children.

Nor is the burden any greater than exists under Proposals of type (1). They too require the site to determine whether it must block access based on age; that determination requires the same sort of judgment required by Proposal (3).

Moreover, relative to a world dominated by systems following Proposal (2), the effective burden of Proposal (3) on sites may be less. The risk with Proposal (2) is that third party ratings may mistakenly block sites. At least the owner of the site has control over whether the blocking occurs in a world with Proposal (3).

Nonetheless, except for the possible benefit of more accurate rating, forcing sites to identify whether their content is "harmful to minors" is a burden relevant to considering the constitutionality, and advisability, of such a proposal. [4]

The risk of misuse of the "minor" certificate

The more significant criticism of Proposal (3), however, is the risk that a signal that a user is a minor would increase the risk that

[4] Note that the burden of requiring labeling is not quite as significant as it is in real space. To an ordinary user viewing the site without a "kids-enabled" browser, the label would be invisible. The only people who know how the site is labeled are those that have enabled discrimination based on the label.

minors will suffer from illegal behavior.[5] Depending upon how the signal was constructed, it could be a simple matter for someone seeking children on the Internet to induce the client to identify that the user was a child. That information could then be used to facilitate abuse.

This risk could be minimized. For example, browsers could be coded to reveal the age of a user only to servers that have been certified to request that information. This would cut down on the improper querying of age information. Second, because it would be easier for law enforcement to identify users who are improperly querying the age identifier, Proposal (3) might well facilitate a better system for tracking down those who would abuse children.[6]

Nonetheless, this risk is a reason to be skeptical of Proposal (3), and to prefer another that might achieve the same benefits without this particular risk. This, in my view, is just what Proposal (4) would do.

The Necessary Legislation

The legislation necessary to realize the configuration I have described is relatively simple.[7] In my view, it would require two

---

[5] Some have argued that Proposal (3) is no different from Proposal (1), since in both cases age must be certified, and the costs of certifying would be the same under both proposals. This is a mistake. Under Proposal (1), age must be certified by some third party, because holding an adult ID gives users access to information to which they otherwise might be blocked from gaining access. There is an incentive, therefore, to lie in securing an adult ID. But a minor-ID would not create any incentive to lie. Indeed, there would be no reason not to allow people to lie about whether they were a minor. Anyone who would want to assure that they were not exposed to material deemed harmful to minors could simply so indicate. Since there is no reason to be certain that a person is truthfully indicating, there would be no need for a third party certification.

[6] Law enforcement, for example, could flood the net with clients pretending to be children, so increasing the odds that an offender would be identified that it would make the net a very dangerous place for child sex-offenders.

[7] All of the legislation that I will describe is civil regulation. In my view, there should not be, and possibly cannot be, criminal regulation in this context. It would be sufficient to impose civil fines on sites that violate the rules proposed here. At least Congress should begin with that assumption, and increase the penalties only upon a showing that sites are not generally complying.

parts. First, it would direct a regulatory agency (which I will assume is the FCC) to specify, in consultation with Internet standards bodies, (a) a minimal protocol to query a client about whether the user was a minor, and (b) a standard for answering such a query. Second, it would direct any server with a substantial custom coming from the United States to implement the protocol for querying and blocking based on age if that site is serving material that is "harmful to minors."

In my view, no legislation would be required to induce compliance on the client side. If there were a simple protocol to query and block based on age, and if sites were required to implement this protocol, then software providers would have a significant incentive to develop tools to implement this protocol and enable parental choice. The legislation, in other words, would create a market that software providers would have an adequate incentive to serve. There would therefore be no need to regulate either the makers of browsers, or the suppliers of operating systems for computers. That part of Proposal (3) would, in a sense, take care of itself.

PROPOSAL (4): THE HARMFUL TO MINORS LABEL

Proposal (4) differs from Proposal (3) in one small, but significant, way. Under both Proposal (3) and (4), sites carrying material harmful to minors would have to rate that material. But while under Proposal (3), the site would block access if the client indicated the user was a minor, under Proposal (4), it is the client that blocks access if the site signals that it is serving material harmful to minors. The critical difference then is that the client does not reveal that the user is a minor; therefore the risks of that revelation are avoided.

This proposal imagines the following configuration:

First, that there was a simple protocol for sites to signal that they were carrying material deemed harmful to minors.

Second, that web browsers were configured as described above, to facilitate different password protected user profiles, as well as the ability to mark that the user of a particular profile was a minor.

Third, that when a client browser using a profile that indicates the user is a minor comes across a site that signals that it is carrying material harmful to minors, the browser blocks access to that site.

With this configuration of technology, parents who wanted to protect their kids from access to material harmful to minors could do so by using a browser so configured – assuming, of course, that suppliers of material harmful to minors displayed a common label indicating as much. Proposal (4) would induce that display, by mandating that servers with material harmful to minors indicate that fact by adopting a common, or specified, label.

In the balance of this section, I consider the benefits and costs of this proposal.

Burdens

The burdens of this configuration of technology and legal requirements are, in my view, the least among the four proposals. Like proposals (1) and (3), this proposal would require sites to label their content. But again, as with Proposal (3), this self-labeling would reduce the risk of mislabeling by third parties. Thus while this requirement would no doubt be a burden on sites carrying material deemed harmful to minors, it would not be a burden that was disproportionate to other proposals, or to the burden on providing such content in real space.

This proposal too would require modification of browser code to enable minor-marked profiles and the blocking of sites that identify themselves as carrying material harmful to minors. But again, both changes in code would be trivial. And if sites generally complied with a requirement to label harmful to minor material, then the market would create a significant incentive for suppliers of browsers or operating systems to facilitate such blocking. Thus legislation effecting this requirement would create a market for software authors to develop child protective software.

Legislation Required to Effect Proposal (4)

The legislation required to bring Proposal (4) into effect is simpler than the legislation necessary to bring into effect Proposal (3). The legislation would direct both an agency and web sites. But the task of both would be simpler under Proposal (4) than under Proposal (3).

Direction to the FCC

Under Proposal (4), an agency would, in consultation with Internet standards bodies, determine a label that a web site could transmit when initiating contact with a client to signal that con-

tent on a particular page was harmful to minors. This protocol could in principle be a simple label, <htm>, </htm>. But how best to implement this would be a judgment initially made by Internet standards bodies.

Direction to web sites

Web sites that carried material harmful to minors would then be required to signal that fact upon connection with a client. The web site would not be required to implement any logic for dealing with the client (as in Proposal (3)). Like a label that indicated that food contained sugar, thereby enabling a diabetic to properly respond, this label would simply signal to a user the fact that the site has judged the material on that page to be harmful to minors. And again, as this label would be buried in the code of a web page, the user would not realize a site is so labeled unless his or her browser was enabled for minor-rated browsing.

Results

If web sites complied with this requirement, then a significant market would develop to take advantage of this additional information being provided by servers. Suppliers of browsers or operating systems would market updates to their technologies so that parents would be able to take advantage of this information. Schools as well could use this information to restrict access on the Internet for computers within their control. No regulation of browser or operating system manufacturers would therefore be required. As with Proposal (3), the market, in a sense, would solve this part of the proposal itself.

The Proposal Compared

Proposal (4) is preferable to, in my view, each of the other three proposals, and to doing nothing at all. In the balance of this section, I sketch reasons why.

Advantages over Proposal (1)

Like Proposal (1), Proposal (4) depends upon a form of identification — that the user is a minor. But unlike Proposal (1), there is no need under Proposal (4) for users to secure costly third party identification. Nor, for the reasons I described above, is there any need for web sites to engage in costly verification of the identification. The assertion made under Proposal (4) (that the user is a mi-

nor) is not a claim that anyone has a reason falsely to assert, or if they do, no one has a reason to correct that falsity. Proposal (4) is better than (1), then, in that it reduces the cost of identification.

Advantages over Proposal (2)

Like Proposal (2), Proposal (4) makes the choice to block content an individual's. No site is required, under this proposal, to block content on its own. But unlike Proposal (2), Proposal (4) would not necessarily lead to labels or filters beyond the narrow class that the government has a legitimate interest in regulating. Individuals may still desire a more comprehensive set of tools for restricting access to Internet content. But the absence of an effective minimum would not artificially increase the demand for more extensive measures.

Advantages over Proposal (3)

Like Proposal (1) and (3), Proposal (4) depends upon a form of identification. Like Proposal (3), it depends upon a form of identifying that the user is a minor. But unlike Proposal (3), that information is not made available to others on the network. The fact that a user is a minor affects just what his or her browser does; it does not signal that fact to other sites. Thus the proposal would not create the risk of abuse for children using the net, though it would, if properly implemented, increase the protection for children.

Advantages over doing nothing

Thus, in my view Proposal (4) trumps each of the three other proposals for zoning minors from material harmful to minors on the Internet. So too does it, in my view, trump the proposal of doing nothing. The consequence of doing nothing is to increase the demand for products based on Proposal (2). As organizations such as the ACLU, and Peacefire, have made abundantly clear, these technologies have imposed a significant cost on free speech on the Internet. The demand for such products would be limited, in my view, if a viable and less restrictive alternative were available. That provides an affirmative reason to prefer regulation over doing nothing.

THE COMPLICATION OF GEOGRAPHY

The one complicating factor in the whole of this analysis is the effect of community standards upon any solution. In principle, it is possible that what is "harmful to minors" in one area of the country is not "harmful to minors" in another. This is possible, at least, though it is by no means necessary. Movies rated "R" or "X" are not rated differently depending upon the part of the country in which they are being played. It is not clear why Internet content would have to be any different.

This is an uncertain issue jurisprudentially, simply because the case that ratified the "harmful to minors" standard, *Ginsberg v. New York*, 390 U.S. 629 (1968), described such material as "obscenity for children." The case was decided, however, before the modern standard for determining obscenity was finally settled upon. Thus it is unclear to what extent the "harmful to minor" standard must be adjusted to different communities. If, as the Third Circuit recently indicated, it does, then this would increase the complexity for all four proposals.

Proposal (4) could incorporate a geographically based difference, though it would raise the costs of the proposal significantly. Rather than simply providing a harmful to minors label, the label would have to indicate harmful to minors in X, where X was a geographic location. That would then set a standard that the client would have to judge relative to. If the jurisdiction of the child were more conservative than site X, then the fact that something was harmful to minor in X would entail it was harmful to minors in the client's jurisdiction. The contrary, however, would not necessarily follow.

The Supreme Court has not finally resolved this question of geography. If they resolve the question in favor of community standards, then this may make *any* regulation too cumbersome. For the reasons I have offered in favor of some regulation over none, in my view, that would be unfortunate.

CONCLUSION

The aim of policy making in cyberspace must be to consider the interaction between law and technology, and to recommend regulation only for that part of a policy problem that will not take care of itself. My aim in this analysis has been to suggest the least invasive form of regulation that will avoid the apparent conse-

quence of no regulation – the spread of "censorware" technologies, or Proposal (2) technologies. Proposal (2) technologies are, in my view, as harmful to free speech values as bad law could be. My aim has been to identify good law that might avoid this bad code.