Revisiting Colored Networks and Privacy Preserving Censorship

Yvo Desmedt
BT Chair of Information Security
University College London
UK

Yongge Wang
University of North Carolina
Charlotte
USA

Mike Burmester
Florida State University
USA

Friday 1 September, 2006

UCL

Yvo Desmedt was also partially supported by EPSRC EP/C538285/1.
He is also a courtesy professor at Florida State University (USA).

This paper was inspired by someone asking how to apply earlier research for preventing censorship.

OVERVIEW

Censorship in Western Societies

Traditional networking model

Color adversary structure

Computational complexity

Secure censoring

Conclusions and open problems

1. Censorship in Western Societies

We all heard about censorship in many non-Western societies.

Censorship in the West is not so uncommon.

Australia: The Australian Communications Minister Helen Coonan has suggested to censor an internet & TV program “Big Brother”. This made news in, e.g.

Canada:

http://www.cbc.ca/story/arts/national/2006/07/05/big-brother.html

UK: news.bbc.co.uk/2/hi/entertainment/5151248.stm

Note: Accordingly to www.censorwatch.co.uk/cw0606.htm the following books are censored in Australia:

- Defence of the Muslim Lands
- Join the Caravan

Belgium: the Information Minister Peter Vanvelthoven is looking into:

censoring websites with illegal content or with illegal services

(translated from the official Belgian memorandum at http://presscenter.org/archive/20060623/64e4b6b15afc76fdf9f1db8c3 or at least to:

inform customers that they entered a black listed site

Critics remember that before 1966 it was hard in small Belgian villages to buy books that were on the Vatican “Index Librorum Prohibitorum” blacklist.

France: Hitler’s “Mein Kampf” is censored in France and some other countries (e.g. Germany).

USA: the Rolling Stones performance during the 2006 superbowl on 5 February 2006 was partially censored.

In other countries monitoring measures are introduced. For example, in the UK the government has the right to know:

> who you phoned, who phoned you, your mobile phone location, email addresses contacted and websites visited.

Texts describing in details the construction of atomic bombs, or other classified information, are also censored.

Whether censorship is a benefit to mankind or not, is a non-scientific topic, and therefore not the focus of the presentation. In this talk we discuss methods that can be used to censor networks.

2. TRADITIONAL NETWORKING MODEL

The classical results:

If an adversary can destroy t nodes, then $t+1$-vertex disjoint paths are needed and sufficient to communicate from node A to node B.

If any two non-destroyed nodes want to communicate, it is necessary and sufficient that the graph must be $t+1$ connected.

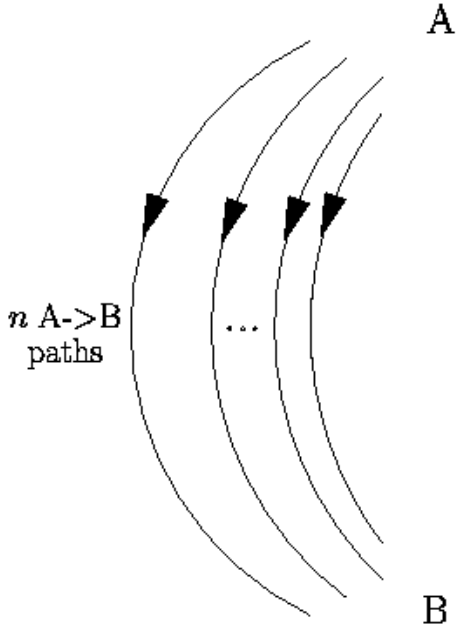

A polynomial time algorithm exists to find:

- the connectivity of the graph
- a separator, i.e. for any sender A and receiver B, one can find a subset of nodes such that A and B are disconnected.

In our context this means that anybody who knows the network can easily find the separator.

Disadvantage: as easy for a limited adversary to perform a denial of service as for the authorities to censor the internet!

Goal: possible for authorities to censor the internet, but hard for cyber terrorist (or hacker) to disrupt.

3. COLOR ADVERSARY STRUCTURE

In the traditional model the adversary can control at most t nodes. Accordingly to this model:

attacking $t+1$ machines running the same platform is hard, but attacking t machines running different platforms is easy!

This model is clearly not realistic. A weakness of one router/computer can easily be exploited on another one if it runs the same platform. Indeed, using viruses and worms one can replicate an attack!

Burmester-Desmedt (2004) proposed the t-color adversary structure. Vertices are given colors. t colors can be corrupted. It

allows to model routers that run the same platform, i.e. have the same weakness, to be assigned the same color.

Color adversary structure is interesting to understand counter-intuitive arguments: i.e.: color separable is not linked to vertex disjoint paths.

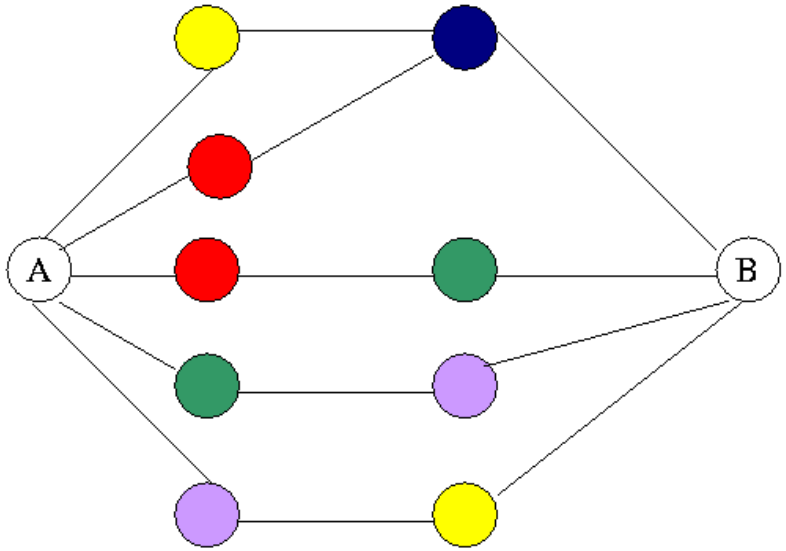

Definition 1. Let $G(V, E)$ be a directed graph, A, B be nodes in $G(V, E)$, and $\mathcal{Z}_{C,t}$ be a t-color adversary structure on $V \setminus \{A, B\}$, where C is the set of colors.

- A, B are called $\mathcal{Z}_{C,t}$-separable in G, if there is a set Z of nodes of at most t different colors such that all paths from A to B go through at least one node in Z. We say that Z separates A and B.

- A, B are called $(\mathcal{Z}_{C,t} + 1)$-*connected* if they are not $\mathcal{Z}_{C,t}$-separable in G.

4. COMPUTATIONAL COMPLEXITY

Deciding whether a vertex colored graph with C the set of colors, is $\mathcal{Z}_{C,t}+1$-connected is co-**NP**-complete.

Proof We demonstrate the complementary problem is **NP**-complete. The reduction is from the Vertex Cover problem.

INSTANCE: A graph $G=(V,E)$ and a positive integer $k \leq |V|$.

QUESTION: Is there a vertex cover of size k or less for G, that is, a subset $V' \subseteq V$ such that $|V'| \leq k$ and, for each edge $(u,v) \in E$, at least one of u and v belongs to V'?

We now construct a network between A and B. Assume $E=\{e_1, e_2, \ldots, e_{m-1}, e_m\}$. Let us define:

- $E_1 = \{e_{(1,1)}, e_{(2,1)}, \ldots, e_{(m-1,1)}, e_{(m,1)}\}$ a set of nodes, and
- similarly E_2
- a bijection f_1 from E to E_1 such that $f_1(e_i) = e_{(i,1)}$, and
- similarly f_2 maps e_i into $e_{(i,2)}$.

We now construct the following new graph G_c:

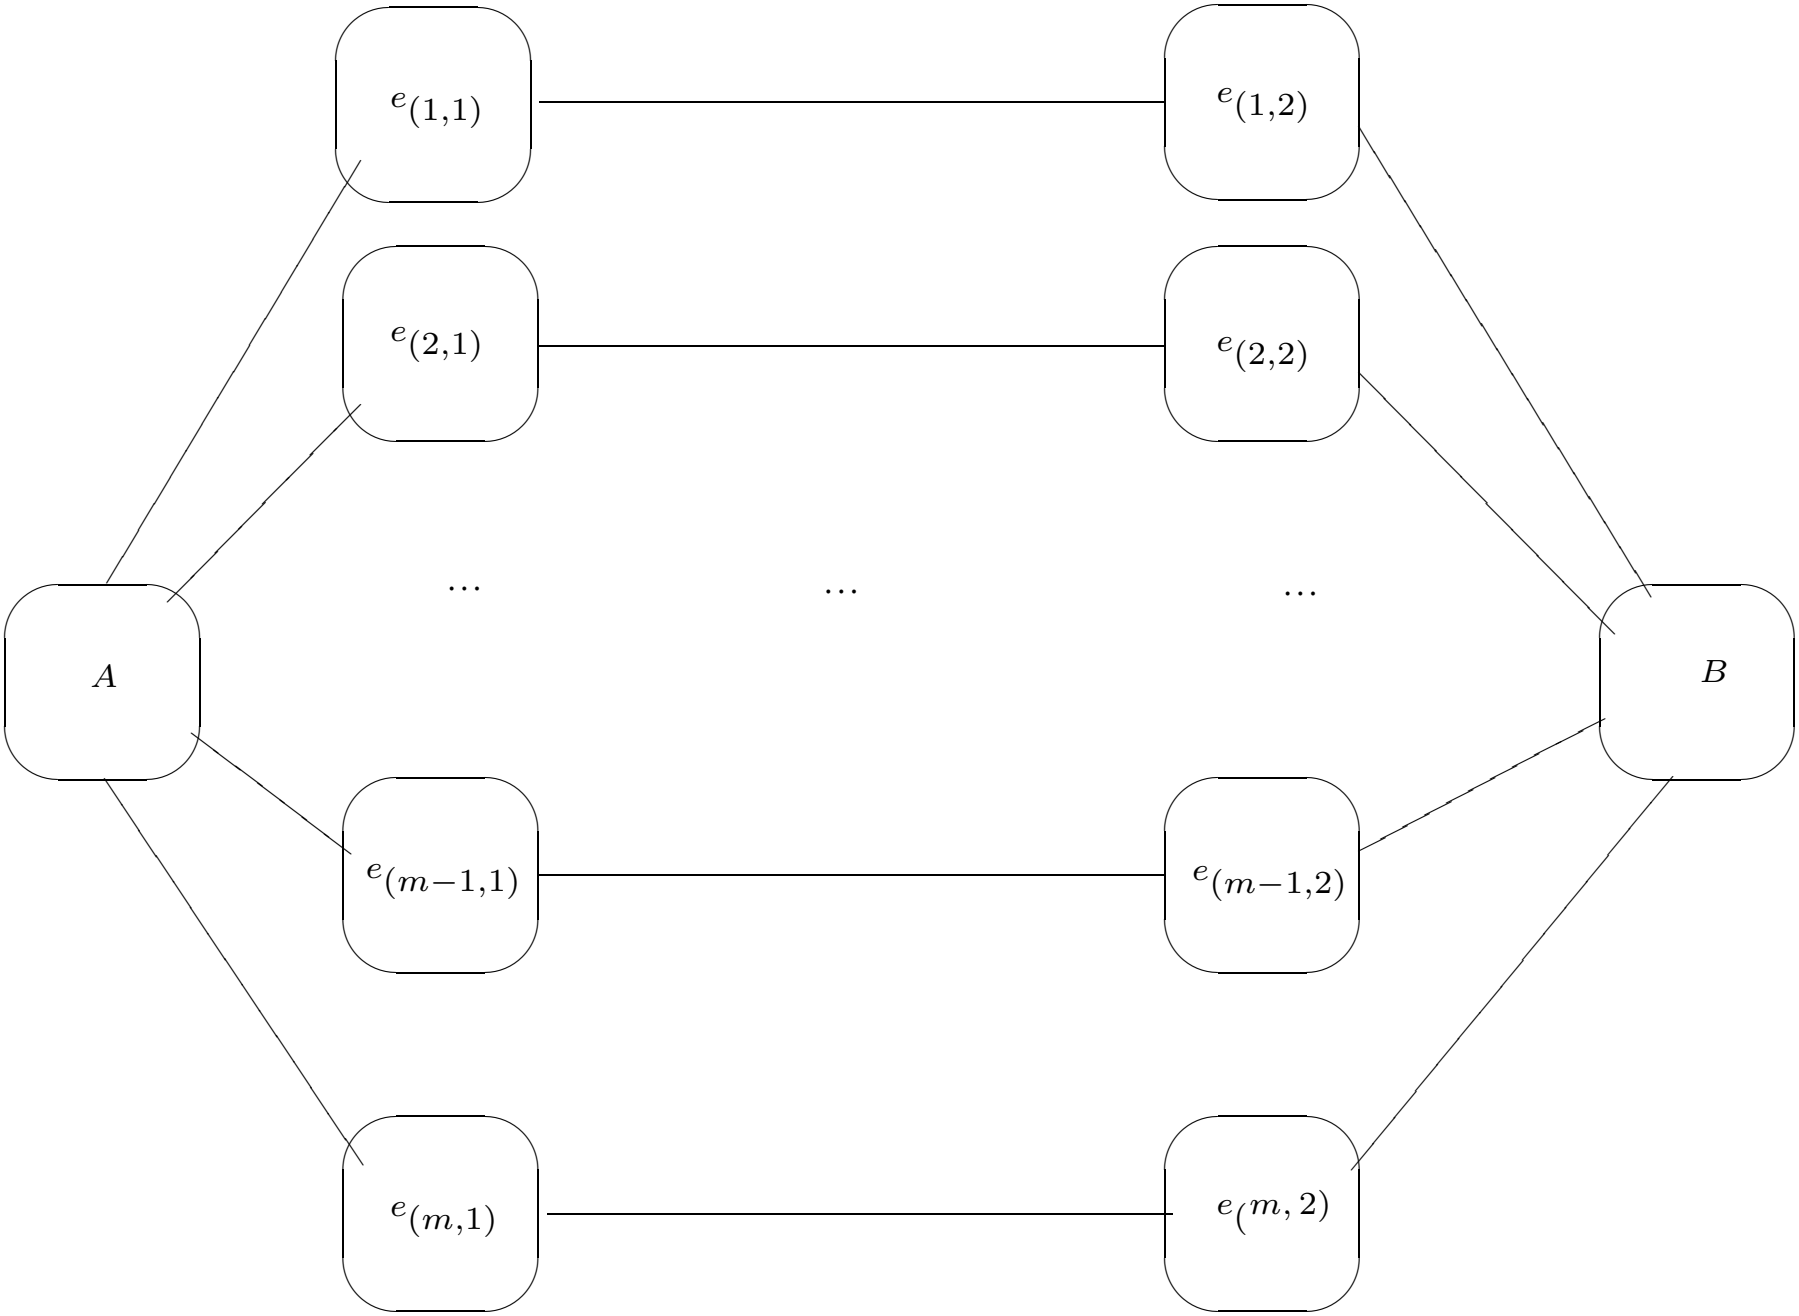
$e_{(1,1)}$
$e_{(1,2)}$
$e_{(2,1)}$
$e_{(2,2)}$
...
...
...
A
B
$e_{(m-1,1)}$
$e_{(m-1,2)}$
$e_{(m,1)}$
$e_{(}m, 2)$

We now color the nodes in E_1 as following. Let $C = V$. Let $e_{(i,1)} \in E_1$. Let $(v_j, v_l) = f_1^{-1}(e_{(i,1)})$, where $j < l$. Color $e_{(i,1)}$ using color v_j. Similar for coloring the nodes in E_2, but we use v_l.

The graph G has a vertex cover of size k if and only if in G_c there are k colors which will disconnect A from B.

5. SECURE CENSORING

If the security model is an ordinary threshold one, then anybody knows who can/can not censor. If the color adversary structure is used, then the problem whether it is (at most) k-color connected, is **NP**-complete. So, the secret is a separator Z of at most k colors.

Why should this remain secret?

Advantage: it may be hard for the limited adversary to find the secret.

So, the question becomes:

> Can the designer prove in zero-knowledge the existence of a k-color separator?

Designing an efficient zero-knowledge proof seems rather trivial.

Here an idea:

Step 1 The prover permutes all the vertices, and permute all the colors and commits to these.

Step 2 The verifier asks a binary question.

Step 3 If the question is 0, then the prover opens all commitments, else he reveals a set V' that separates A and B in this isomomorphic graph.

Step 4 The verifier, in the first case, checks the commitment. In the else case, the verifier checks that the number of colors in V' is at most k and checks V' indeed separates.

Unfortunately, above protocol is not zero-knowledge. Indeed, it

leaks the size of V', which it should not. The knowledge of the size of V' may help the verifier to find the k colors. Moreover, it also leaks the multiplicity of each color, etc.

To solve this problem, we prove the following lemma.

Lemma 1. *Let $G_c = G_c(V, E, C, f)$ be a vertex-colored graph, where C is the set of colors and $f : V \to C$. Let $C' \subseteq C$ be such that $|C'| = k$ and $V' = \{v'_i : f(v'_i) \in C'\}$ separate A and B. Let k' be the maximum number of vertex disjoint paths in (V, E) ignoring the colors. Let P_1, P_2, . . . , $P_{k'}$ be these vertex disjoint paths. We then have that for each of these path P_i: $P_i \cap V' \neq \emptyset$. So, on each path P_i there exists a node of a color in C'.*

Proof: The proof follows trivially by contradiction. □

Zero-Knowledge interactive proof

Setting

Let $G = G(V, E, C, f)$ be a vertex-colored graph and $m = |C|$. For simplicity we assume $C = (1, 2, \ldots, m)$. Let C' and V' be as before.

Precomputation

First the verifier and the prover (separetely) compute:

- k', i.e. the maximum number of vertex disjoint paths ignoring colors.
- k' vertex disjoint paths P_1,P_1, P_2, …, $P_{k'}$.

This can be done in polynomial time. So both prover and verifier obtain the same k' vertex disjoint paths. Let l_i be the length of the path P_i minus one, and let us call the vertices, except A and B, on

this path $v_{(i,1)}$, $v_{(i,2)}$,. . . ,$v_{(i,l_i)}$.

Protocol

They repeat the following steps n times, where n is specified later. The randomness in each run is chosen independently.

Step 1 The prover chooses a permutation π of the colors, so $\pi \in_R sym(\{1, \ldots, m\})$. For each of the aforementioned paths P_i:

- the prover chooses a permutation $\rho_i \in_R sym(\{1, \ldots, l_i\})$, permutes the vertices (ignoring A and B) on the path P_i and sends the verifier a commitment for the permuted coloring of the permuted vertices, so formally, sends:

$$E_{(i,j)} = \text{commit}(\pi(f(v_{(i,\rho_i(j))})), r_{ij}) \text{ for } j = 1, \ldots, l_i,$$

where r_{ij} is chosen independently uniformly random, and

- for each $c_h \in C'$ ($h = 1, \ldots, k$) sends $E'_h = \mathrm{commit}(\pi(c_h), r'_h)$, where r'_h is chosen independently uniformly random.

Step 2 The verifier flips a coin q_1 and also chooses randomly a value $q_2 \in_R \{1, \ldots, k'\}$ and sends the prover the query (q_1, q_2).

Step 3 *If* $q_1 = 0$, *then* the prover reveals π, all ρ_i and opens all commitments of the type $E_{(i,j)}$ (Note the prover does not open E'_h.),

else the prover decommits one (of the) permuted colors of the vertex set: $P_{q_2} \cap V'$. This is done by opening:

- exactly one $E_{(q_2,j')}$, and
- exactly one E'_h

such that $f(v_{(q_2,\rho_{q_2}(j'))}) = c_h$. (Note π is not opened, and neither is ρ_{q_2}).

Step 4 *If* $q_1 = 0$, *then* the verifier verifies that π and all ρ_i are permutations and all the decommitted values,

else the verifier checks that the two opened commitments and checks that they correspond to the same color.

Theorem 1. *When n is chosen such that $((k'-1)/k')^n$ is negligible, the protocol is a computational zero-knowledge interactive proof system for the color separable problem assuming that the commitment function* commit *is secure.*

6. Conclusions and open problems

Open problem: how to efficiently generate hard instances with a trapdoor. This means:

> How to generate a colored graph such that the provider can demonstrate to the authorities the existence of a t-color separator, while at the same time it is hard for the limited adversary to find these t colors.

Conclusion: we demonstrated that it may be hard for a limited adversary to perform a denial of service, while the provider can demonstrate to the authorities that censorship is possible.