Covert channels and anonymity

Steven J. Murdoch

UNIVERSITY OF CAMBRIDGE
Computer Laboratory
Security Group

Covert channels

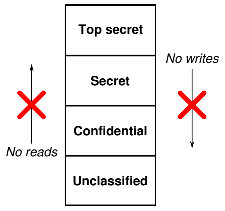

When military and intelligence organisations computerised their operation, they were concerned that confidential data might leak out. To reduce this risk, they required that computers enforce their existing document handling rules. In the simplest case, files are classified as either *top secret*, *secret*, *confidential* or *unclassified*. Users are assigned to levels and the operating system forbids *reading* information above a user's clearance. The system also prevents *writing* information from a high-level file to a lower-level one.

Normal storage and communication mechanisms, such as files and networks, can be made to respect these restrictions. But there are other ways to transmit information that the system designer might not have considered, called *covert channels*. For example, a high-level program could send data by changing its CPU usage while a low-level one observes CPU load to recover the data.

Attacking Tor

Covert channels can be used to learn along which path messages are sent through the anonymising network Tor. Here, the attacker operates a webserver (corrupt server) that the user (victim) is accessing. The webserver inserts a distinctive load pattern into the network, and a second server (corrupt Tor node), controlled by the attacker, measures the performance of each Tor relay in turn:

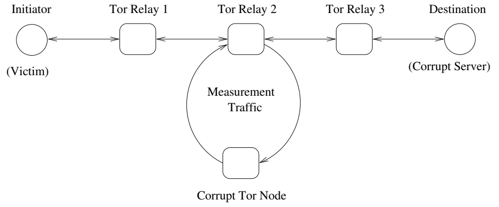

When the corrupt Tor node connects to one of the Tor relays carrying the victim's traffic, the attacker will see the pattern, and so learns that the path goes through the node being measured.

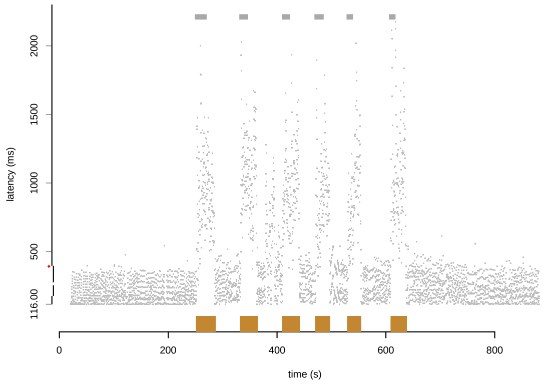

CPU temperature as a covert channel

In addition to anonymous web browsing, Tor allows users to run servers without giving away their identity. To find out who is operating a particular hidden server, an attacker could increase its CPU load by connecting through the Tor network, while simultaneously connecting to each hidden server directly and measuring load:

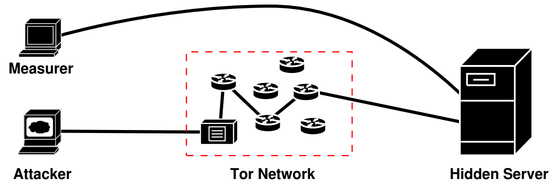

The frequency of clock crystals, as used in computers for measuring time, varies with temperature. As the load of a computer increases, the temperature will too, so by remotely measuring changes in clock frequency, the CPU load of a system can be estimated, even if it cannot be measured directly.

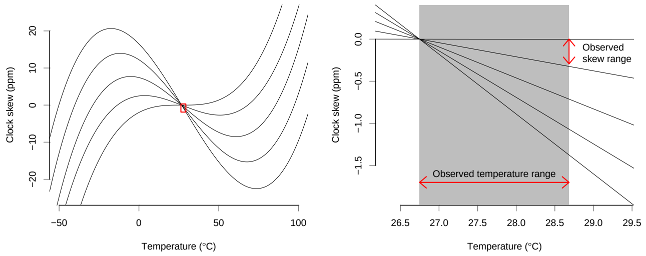

If the attacker can see a match between the load pattern induced and the clock frequency, he knows that the correct hidden server has been found. In the graph below, the grey bars show the load pattern, the brown circles show the temperature and the blue triangles the clock frequency. The load pattern can be seen to alter the clock frequency, so here the attacker has identified the server.

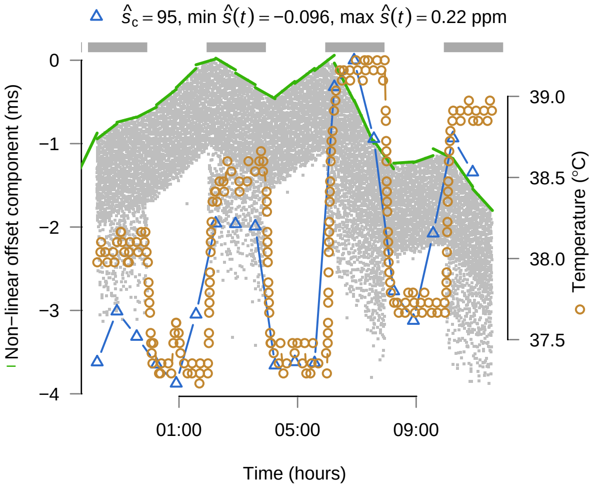

For more information see `http://www.cl.cam.ac.uk/users/sjm217/projects/anon/`

2006-07-06