

How Much Bandwidth Can Attack Bots Commandeer?

Michael B. Greenwald*, Sanjeev Khanna[†], and Santosh S. Venkatesh[†]

*Bell Laboratories, Murray Hill, NJ 07974

Email: greenwald@research.bell-labs.com

[†]University of Pennsylvania, Philadelphia, PA 19104

Email: sanjeev@cis.upenn.edu, venkatesh@ee.upenn.edu

Abstract—In a shared channel model for internet links, bandwidth is shared by principled users who abide by communal principles for sharing and using bandwidth and unprincipled scofflaws who seek to commandeer as much of the bandwidth as possible to effect disruptions such as spam and DoS attacks. Attacks are magnified by the spread of bots that surreptitiously take over the functioning of legitimate users. In such settings the natural filtering by router policies at ingress nodes and the rate of growth of link capacities towards the backbone play key roles in determining what fraction of the bandwidth is eventually commandeered. These considerations are presented in detail for a tree topology with users scattered at the leaves and with varying link capacity assignments and idealised router policies.

I. FROM THE DOLEV-YAO MODEL TO THE SHARED CHANNEL MODEL AND BEYOND

Threats to the integrity and confidentiality of communications have traditionally been studied in the now classical Dolev-Yao framework [1] in which an adversary is assumed to be able to exercise complete control over the communication network. The key result of these studies is that cryptographic guarantees can provide a level of immunity to these twin threats even under this draconian model of adversarial control of the communication infrastructure. With the rapid proliferation of the Internet, however, a new security threat has emerged: denial of service (DoS) attacks seek not to eavesdrop on or spoof communications but are content merely with disrupting interactions by making it impossible to communicate to select users. The classical cryptographic approaches to providing integrity and confidentiality do not provide much guidance on how to deal with this new threat.

It is clear that the Dolev-Yao model cedes too much power to an adversary in this setting. Certainly, if an adversary has control of the network—and can consequently select and drop packets at will—he is already assured a DoS capability. The *shared channel model* proposed in recent work of the authors [2], [3] is an effort to relax some of the Dolev-Yao provisions while still permitting adversaries substantial, but bounded, resources. The model assumes that a legitimate user and an attacker share a packet communication channel

to a receiver. While the details of the model are not important for our purposes here, two key features are relevant: a (single) adversary is assumed to have large, but bounded, resources in terms of computation and access to channel bandwidth; and each legitimate user is assumed to have access to a certain minimum bandwidth even in the worst case of an all out attack by the adversary. It seems reasonable to presuppose these conditions in a network model for the investigation of DoS attacks and defences as the violation of either of these will result in a *prima facie* DoS capability for the adversary. Within the framework of the shared channel model, the results of the studies [2], [3] were encouraging in so far as it was shown that in some settings DoS attacks could be thwarted at low cost.

Even a cursory consideration of any real network suggests, however, that the shared channel model, for all its appeal, is somewhat simplistic in its blithe allocation of bandwidth. Nor does it take into account a possible plethora of inimical sources, zombies, and bots who take over unwary legitimate users. In such settings the natural filtering by router policies at ingress nodes and the rate of growth of link capacities towards the backbone play key roles in determining what fraction of the bandwidth is eventually commandeered. These considerations are presented in detail here for a tree topology with users scattered at the leaves and with varying link capacity assignments and idealised router policies. We present a framework, in a somewhat sanitised and idealised setting, in which one can analyse how the density of bots affects the fraction of the available bandwidth that adversarial traffic can commandeer. For the tree structure considered here, as we shall see, exact results and prescriptions can be obtained with a modicum of effort. Our longer term goal is to understand how to apply our results to more general settings.

II. A REGULAR TREE TOPOLOGY

Networks are built with arbitrarily complex topologies. However, we may want to look at the network from the point of view of a single server, considering only the subgraph consisting of paths from the clients to the server. In the Internet, at any instant in time, the shape of

such a subgraph is generally a tree. (In the common case, packets are routed based on their destination address. Thus, two packets from distinct sources that “meet” anywhere in the network, will share the remainder of their paths to the server. The shape of the subgraph is then a tree, with the server at the root.)

Accordingly, consider a tree topology consisting of an edge-weighted, regular d -ary tree of depth n . The d^k vertices at depth k in the tree form the vertices at generation k , with the root of the tree identified with generation 0 and the terminating d^n vertices forming the leaves of the tree at generation n . We identify the internal vertices of the tree with *routers*. The sources (or *users*) are all arrayed in the leaves of the tree with information flowing upwards from the leaves towards the root. Thus, each internal vertex of the tree has d incoming edges, one from each of its d children, and one outgoing edge to its parent. An outgoing edge is also attached to the root of the tree which may be considered to be the gateway to, say, a backbone or a sink.

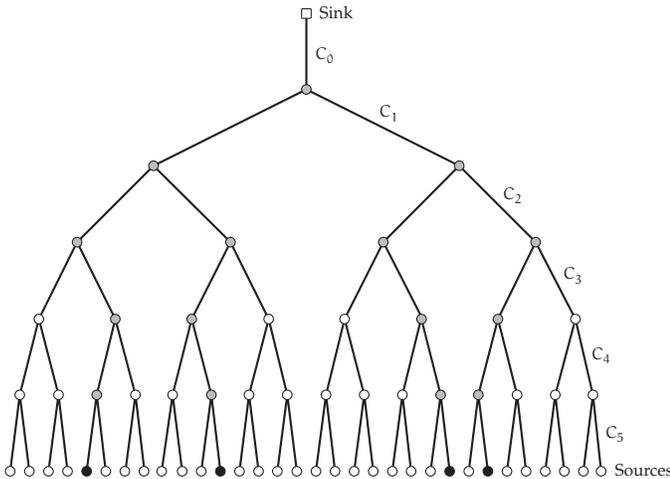


Fig. 1. The dark vertices among the leaves denote compromised users (bots). The lightly shaded vertices show the routers that are contaminated.

We associate a positive edge weight C_k with the outgoing edge of each vertex in generation k . These weights represent *link capacities*. (We may impose the natural majorisation $C_0 \geq C_1 \geq \dots \geq C_{n-1} \geq C_n$ with routers possessing larger and larger capacities the closer they get to the backbone, though this kind of monotone relation between link capacities is not essential for our results.)

Each legitimate (or *uncompromised*) user offers a (mean) load of β (per unit time) to its parent vertex in generation $n - 1$. We may suppose, without loss of generality, that $\beta \leq d^{-(n-k)} C_k$ for each $0 \leq k \leq n$. (Else legitimate users will suffer packet drops even under clean operating conditions and we may as well set

$$\beta \leftarrow \min\{d^{-(n-k)} C_k, 0 \leq k \leq n\}.$$

which will reduce the problem to the stated case.)

For some $0 \leq r \leq n$, suppose that a fraction $d^{-(n-r)}$ of the legitimate users are subverted with *attack bots* taking over their functioning. We suppose that each bot (or compromised user) offers a load $\alpha := C_n$ to its parent vertex limited only by the outgoing link capacity at the leaves.

We will suppose that the d^r bots are distributed uniformly across the leaves, one per each group of d^{n-r} leaves. More precisely, suppose that each vertex at generation r has precisely one bot (i.e., compromised user) among its descendants. This will correspond to a worst-case scenario where the bot distribution is maximally effective in disrupting traffic.

Say that a vertex is *contaminated* if it has a compromised descendant, and *uncontaminated* otherwise. Then all vertices in generations $0 \leq k \leq r$ are contaminated.

We assume a *fair queueing protocol* where each router at generation k guarantees a minimum bandwidth of up to C_k/d to each of its d children. Subject to the availability of bandwidth under this fair queueing model, each router at generation k makes a best-effort attempt at delivering any offered load up to its link capacity C_k . In this model, after all children whose offered load is $\leq C_k/d$ are accommodated, any excess bandwidth is allocated uniformly across all children who have loads in excess of C_k/d . If the offered load on any incoming edge is in excess of the total allocated bandwidth then the router drops packets on the offending link until the offered load matches the allocated bandwidth. We make the simplifying assumption that packet level granularity is sufficiently fine for the flow on each edge to be thought of as a fluid. In a bow to convention we will continue to refer to the components of a flow as “packets” though with the understanding that the flow is arbitrarily divisible and, at need, any given fraction of a flow can be dropped to match the bandwidth allocated to that flow.

Our variation on the fair queueing theme is idealised and deviates somewhat from protocols used in practice. Ideal “fair queueing” mechanisms allocate capacity equally among all of the *end-to-end* flows. In practice, protocols such as Stochastic Fair Queuing (SFQ) are deployed, where each link is divided up into a fixed number of bins, and capacity is allocated evenly between the bins. In contrast, our variant of fair queueing, allocates outgoing capacity fairly between incoming *links* regardless of the number of individual flows on each link. Our riff on the general theme reflects the behaviour of stochastic fair queueing under overload, where there are more flows than bins at each router, and has the great advantage of simplicity. At the cost of some added complexity, we can apply a similar analysis to the cases of SFQ or Fair Queueing.

Random drop models are both natural and easy to implement at the routers and we consider two variants. To eschew unnecessary complexity at this stage we

begin with two stylized drop models that ignore random effects in trimming flows.

- *Agnostic drop model.* The router makes no judgement about the presence or absence of hostile elements in an incoming flow that exceeds allocated bandwidth and merely drops packets uniformly across the flow until the load offered by the link matches the allocated bandwidth. The proportionate mix of users in the trimmed flow is kept in the same proportion as in the original incoming flow.
- *Pernicious drop model.* Packets in the contaminated flow that originate with legitimate users are dropped first (uniformly across all legitimate packets in the flow under consideration) with bot generated packets dropped only after all legitimate packets are exhausted. In the portion (if any) of the trimmed flow that is occupied by legitimate users, the users appear in the same proportion vis à vis each other as they did in the incoming flow. Attacker friendly drops of this nature may be thought of as an attempt to model situations where an attacker exploits knowledge of the router protocol by timing, replay, or other attacks.

The fraction of the flow of any user that is dropped in the above proportionate drop models may be identified with the *expected* fraction of the flow that is dropped in the corresponding random drop models. As we shall see, for reasonable capacity link assignments, the results may be interpreted in terms of expected loads for random drop protocols.

On notation: It will be convenient to introduce the notations $x \wedge y := \min\{x, y\}$ and $x \vee y := \max\{x, y\}$ for minimum and maximum, respectively. As is usual, we write

$$x_+ := \begin{cases} x & \text{if } x \geq 0, \\ 0 & \text{if } x < 0, \end{cases}$$

to denote the positive part of x .

III. EDGE FLOWS

Ignoring random effects for the nonce, suppose that each legitimate user offers a load exactly equal to β .

Write φ_k for the load carried by the outgoing edge of any vertex in generation k . We call φ_k the *flow emanating from* a vertex in generation k . We write $\varphi_k = \varphi_k^c$ if the vertex is contaminated, and $\varphi_k = \varphi_k^u$ if it is uncontaminated.

Vertices in generations $r \leq k \leq n$ may be either uncontaminated or contaminated. We consider these in turn.

The flow emanating from uncontaminated vertices satisfies the recurrence $\varphi_k^u = d\varphi_{k+1}^u$ (recall that $\beta d^{n-k} \leq C_k$ for each k and that the bandwidth allocation model supports fair queueing). Together with the boundary condition $\varphi_n^u := \beta$, this leads to the solution

$$\varphi_k^u = \beta d^{n-k} \quad (r+1 \leq k \leq n). \quad (1)$$

(Bear in mind that all vertices in generations $k \leq r$ are contaminated.)

Now each contaminated vertex in generation k (with $r \leq k \leq n$) has precisely one contaminated child. Each such vertex will hence accommodate the entire offered load $\varphi_{k+1}^c + (d-1)\varphi_{k+1}^u$ if it is less than the link capacity C_k ; else it will accommodate the entire offered load from each of its uncontaminated children (as $\varphi_{k+1}^u = d^{n-k-1}\beta \leq C_k/d$) and limit the bandwidth allocated to the contaminated child to $C_k - (d-1)\varphi_{k+1}^u = C_k - (d-1)d^{n-k-1}\beta$ so that the total traffic carried by the outgoing link is at link capacity C_k . Accordingly, the flow emanating from a contaminated vertex in generation k satisfies the recurrence

$$\varphi_k^c = C_k \wedge \{\varphi_{k+1}^c + (d-1)\varphi_{k+1}^u\} \quad (r \leq k \leq n).$$

By induction, we quickly obtain

$$\begin{aligned} \varphi_k^c &= C_k \wedge \{C_{k+1} + (d-1)d^{n-k-1}\beta\} \\ &\quad \wedge \{C_{k+2} + (d-1)(d+1)d^{n-k-2}\beta\} \\ &\quad \wedge \{C_{k+3} + (d-1)(d^2 + d + 1)d^{n-k-3}\beta\} \wedge \dots \end{aligned}$$

which leads to the general solution

$$\begin{aligned} \varphi_k^c &= \bigwedge_{j=k}^n \{C_j + (d^{j-k} - 1)d^{n-j}\beta\} \\ &= \bigwedge_{j=k}^n \{C_j + (d^{n-k} - d^{n-j})\beta\} \quad (2) \end{aligned}$$

for $r \leq k \leq n$.

To finish up, consider vertices in generations $0 \leq k \leq r-1$. As all such vertices are contaminated, we obtain the recurrence

$$\varphi_k = \varphi_k^c = C_k \wedge (d\varphi_{k+1}^c)$$

and induction again quickly yields

$$\begin{aligned} \varphi_k^c &= \bigwedge_{j=k}^{r-1} (d^{j-k}C_j) \wedge (d^{r-k}\varphi_r^c) = \left(\bigwedge_{j=k}^{r-1} d^{j-k}C_j \right) \\ &\quad \wedge \left(\bigwedge_{j=r}^n \{d^{r-k}C_j + (d^{n-k} - d^{n-j+r-k})\beta\} \right). \quad (3) \end{aligned}$$

for $0 \leq k \leq r-1$. The *convention* $\bigwedge_{\emptyset} := \infty$ allows us to unify our results and write

$$\begin{aligned} \varphi_k^c &= \left(\bigwedge_{j=k}^{r-1} d^{j-k}C_j \right) \wedge \left(\bigwedge_{j=k \vee r}^n \{d^{(r-k)+}C_j \right. \\ &\quad \left. + (d^{n-k} - d^{n-j+(r-k)+})\beta\} \right) \quad (0 \leq k \leq n). \end{aligned}$$

It follows *a fortiori* that the flow emanating from the root satisfies

$$\varphi_0 = \varphi_0^c = \left(\bigwedge_{j=0}^{r-1} d^j C_j \right) \wedge \left(\bigwedge_{j=r}^n \{d^r C_j + (d^n - d^{n-j+r})\beta\} \right).$$

IV. ATTACK RATIOS AND COOPTED BANDWIDTH

Let α_k denote the bot generated portion of the flow emanating from a contaminated vertex at generation k . We will refer to bot generated traffic as *attack packets* and in keeping with this vivid terminology we may identify α_k as the *attack flow* emanating from a contaminated vertex in generation k . The *attack ratio* $\rho_k := \alpha_k / \varphi_k^c$ then connotes the fraction of the total flow emanating from the vertex that is bot driven.

AGNOSTIC DROP MODEL: Begin with the recurrence base

$$\alpha_n = \alpha = C_n.$$

Now consider the flow emanating from any contaminated vertex in generation k for $r \leq k \leq n-1$. Recall that each such vertex has exactly one incoming contaminated flow with the remaining $d-1$ incoming flows uncontaminated. If the total flow entering the vertex, $\varphi_{k+1}^c + (d-1)\varphi_{k+1}^u$, is no larger than the link capacity C_k then the vertex passes on the accumulated incoming flow onto its outgoing link so that $\varphi_k^c = \varphi_{k+1}^c + (d-1)\varphi_{k+1}^u$ and $\alpha_k = \alpha_{k+1}$. If, on the other hand, the net incoming flow $\varphi_{k+1}^c + (d-1)\varphi_{k+1}^u$ exceeds link capacity C_k then, as fair queueing mandates that the entire flow from uncontaminated links is passed on untrammelled, the bandwidth allocated to the incoming contaminated link is $C_k - (d-1)\varphi_{k+1}^u$ and a random selection of packets is dropped from the offending incoming link until its offered load matches the allocated bandwidth. As the fraction of attack packets in the contaminated incoming flow is $\rho_{k+1} = \alpha_{k+1} / \varphi_{k+1}^c$, under a model of uniform agnostic drops it follows that this is also the fraction of the allocated bandwidth that is usurped by attack packets and, consequently, $\alpha_k = (C_k - (d-1)\varphi_{k+1}^u) \rho_{k+1}$. Putting the two cases together we obtain the recurrence

$$\begin{aligned} \alpha_k &= \alpha_{k+1} \wedge (C_k - (d-1)\varphi_{k+1}^u) \rho_{k+1} \\ &= \{1 \wedge (C_k - (d-1)\varphi_{k+1}^u) / \varphi_{k+1}^c\} \alpha_{k+1} \end{aligned}$$

valid for all $r \leq k \leq n-1$. An easy induction now leads to the general solution

$$\alpha_k = C_n \prod_{j=k}^{n-1} \{1 \wedge (C_j - (d-1)\varphi_{j+1}^u) / \varphi_{j+1}^c\} \quad (r \leq k \leq n), \quad (4)$$

the usual convention $\prod_{\emptyset} := 1$ allowing us to extend the result to the base case $k = n$ as well.

The product on the right suitably normalises by a quantity no larger than one. For instance, we may verify

$$\begin{aligned} \alpha_{n-1} &= C_n \{1 \wedge (C_{n-1} - (d-1)\varphi_n^u) / \varphi_n^c\} \\ &= C_n \wedge \{C_n (C_{n-1} - (d-1)\beta) / C_n\} \\ &= C_n \wedge \{C_{n-1} - (d-1)\beta\} \end{aligned}$$

where the right-hand side is in accordance with what we may write down by direct observation.

To complete the recursive specification, each vertex in generation k for $0 \leq k \leq r-1$ sees d contaminated incoming edges, each carrying traffic $\varphi_{k+1}^c = \alpha_{k+1} + (\varphi_{k+1}^c - \alpha_{k+1})$ of which α_{k+1} is the portion coopted by the attack flow. If $d\varphi_{k+1}^c < C_k$ then the entire offered load $d\varphi_{k+1}^c$ is accepted and passed on to the outgoing link whence $\alpha_k = d\alpha_{k+1}$. If $d\varphi_{k+1}^c > C_k$, on the other hand, the incoming traffic is pruned back with traffic on each incoming link allocated bandwidth C_k/d whence $\alpha_k = d(\rho_{k+1}C_k/d) = \rho_{k+1}C_k$. It follows that

$$\alpha_k = (d\alpha_{k+1}) \wedge (\rho_{k+1}C_k) = \{d \wedge C_k / \varphi_{k+1}^c\} \alpha_{k+1}$$

which leads to the inductive solution

$$\alpha_k = \alpha_r \prod_{j=k}^{r-1} \{d \wedge C_j / \varphi_{j+1}^c\} \quad (0 \leq k \leq r-1). \quad (5)$$

Pooling results, we obtain the desired expression

$$\begin{aligned} \alpha_k &= C_n \left(\prod_{j=k}^{r-1} \{d \wedge C_j / \varphi_{j+1}^c\} \right) \\ &\quad \times \left(\prod_{j=k \vee r}^{n-1} \{1 \wedge (C_j - (d-1)\varphi_{j+1}^u) / \varphi_{j+1}^c\} \right), \end{aligned}$$

with the convention on products over empty sets allowing us to extend the result to the cases $r \leq k \leq n$ as well. In particular, the attack flow at the root is given by

$$\begin{aligned} \alpha_0 &= C_n \left(\prod_{j=0}^{r-1} \{d \wedge C_j / \varphi_{j+1}^c\} \right) \\ &\quad \times \left(\prod_{j=r}^{n-1} \{1 \wedge (C_j - (d-1)\varphi_{j+1}^u) / \varphi_{j+1}^c\} \right). \end{aligned}$$

PERNICIOUS DROP MODEL: The recurrence base is unaffected with

$$\alpha_n = \alpha = C_n$$

as before. For a vertex in generation k with $r \leq k \leq n-1$, if the net incoming flow $\varphi_{k+1}^c + (d-1)\varphi_{k+1}^u$ is no larger than the outgoing link capacity C_k then the entire incoming flow is passed on to the outgoing link, $\varphi_k^c = \varphi_{k+1}^c$ and, *a fortiori*, $\alpha_k = \alpha_{k+1}$. If, on the other hand, the net incoming flow exceeds link capacity then, under the aegis of fair queueing, all uncontaminated incoming flows are allocated bandwidth sufficient to handle their flows with the bandwidth allocated to the contaminated incoming flow restricted perforce to $C_k - (d-1)\varphi_{k+1}^u$. Under the pernicious drop model, the bot generated portion of the contaminated incoming flow gets first access to this bandwidth (up to the allocated maximum) whence $\alpha_k = \alpha_{k+1} \wedge \{C_k - (d-1)\varphi_{k+1}^u\}$. Pooling cases, we have

$$\alpha_k = \alpha_{k+1} \wedge \{C_k - (d-1)\varphi_{k+1}^u\} \quad (r \leq k \leq n-1).$$

An easy induction mops up and we obtain

$$\alpha_k = \bigwedge_{j=k}^{n-1} \{C_j - (d-1)\varphi_{j+1}^u\} \wedge C_n \quad (r \leq k \leq n), \quad (6)$$

with the convention on the minimum over an empty set allowing us to extend the identity to the base case $k = n$.

Arguing as before, for $0 \leq k \leq r-1$, all incoming flows to a vertex in generation k are contaminated. If $d\varphi_{k+1}^c \leq C_k$ then the entire incoming flow can be accommodated: $\varphi_k^c = d\varphi_{k+1}^c$ and $\alpha_k = d\alpha_{k+1}$. If $d\varphi_{k+1}^c > C_k$ then each of the incoming flows is allocated bandwidth C_k/d . As the attack segments of these flows grab the allocated bandwidth first up to the allocated maximum, $\alpha_k = d(\alpha_{k+1} \wedge C_k/d) = (d\alpha_{k+1}) \wedge C_k$. It follows that

$$\alpha_k = (d\alpha_{k+1}) \wedge C_k = \bigwedge_{j=k}^{r-1} (d^{j-k}C_j) \wedge (d^{r-k}\alpha_r) \quad (7)$$

for $0 \leq k \leq r-1$.

Unifying the two cases, the attack flow emanating from a contaminated vertex in generation k is given by

$$\alpha_k = \bigwedge_{j=k}^{r-1} (d^{j-k}C_j) \wedge \bigwedge_{j=k \vee r}^{n-1} d^{(r-k)_+} \{C_j - (d-1)\varphi_{j+1}^u\} \wedge \{d^{(r-k)_+}C_n\}$$

for $0 \leq k \leq n$. In particular, the attack flow emanating from the root is given by

$$\alpha_0 = \bigwedge_{j=0}^{r-1} (d^jC_j) \wedge \bigwedge_{j=r}^{n-1} d^r \{C_j - (d-1)\varphi_{j+1}^u\} \wedge \{d^rC_n\}.$$

It is instructive to compare these results for the pernicious drop model with the corresponding expressions obtained for the agnostic drop model.

V. BANDWIDTH UTILISATION

Of the traffic entering the sink, the total bandwidth utilised by legitimate users is $\varphi_0^c - \alpha_0$. It is not difficult now to obtain a finer resolution of bandwidth utilisation across different categories of legitimate users.

Introduce the natural metric on the vertices of the tree by setting the the distance between any two vertices to be the hop count to their closest common ancestor. For each $1 \leq t \leq n-r$, let U_t be the equivalence class of users at distance t from the nearest bot. It is easy to see that a given bot has $d-1$ users at distance 1 from it, d^2-d users at distance 2 from it and, in general, d^t-d^{t-1} users at distance t . By symmetry, it follows perforce that $\text{card } U_t = d^t(d^t-d^{t-1})$.

For each t and k , let $\beta_{t,k}$ denote the flow originating from a user in class U_t that is admitted into the outgoing link of its ancestor in generation k . Now fix any $1 \leq t \leq n-r$ and consider any user in the class U_t . As the nearest contaminated ancestor is t hops away, fair

queueing dictates that the entire user flow is preserved through the lowest $t+1$ generations whence

$$\beta = \beta_{t,n} = \beta_{t,n-1} = \dots = \beta_{t,n-t}.$$

Suppose $k \geq r$. Then the subtree with root at any contaminated vertex in generation k has d^{n-k} leaves consisting of one bot, $d-1$ users in the class U_1 , d^2-d users in the class U_2 , ..., and $d^{n-k}-d^{n-k-1}$ users in the class U_{n-k} . It follows that the flow emanating from the vertex may be decomposed into

$$\begin{aligned} \varphi_k^c &= \alpha_k + \sum_{t=1}^{n-k} (d^t - d^{t-1})\beta_{t,k} \\ &= \alpha_k + (d-1) \sum_{t=1}^{n-k-1} d^{t-1}\beta_{t,k} + (d-1)d^{n-k-1}\beta \end{aligned}$$

as $\beta_{n-k,k} = \beta$.

Now, fix t and consider any given user in class U_t . For $r \leq k \leq n-t-1$, the k th generation ancestor of the tagged user sees $d-1$ uncontaminated incoming flows in addition to a single contaminated flow containing the tagged user's packets. The bandwidth allocated by the vertex to the contaminated flow is hence $\varphi_k^c - (d-1)\varphi_{k+1}^u$ of which α_k is coopted by the attack flow. The residual bandwidth $\varphi_k^c - \alpha_k - (d-1)d^{n-k-1}\beta$ is spread proportionately among the commingled flows of the legitimate users in the contaminated incoming flow

$$\varphi_{k+1}^c = \alpha_{k+1} + \sum_{s=1}^{n-k-1} (d^s - d^{s-1})\beta_{s,k+1}.$$

Accordingly,

$$\begin{aligned} \beta_{t,k} &= \frac{\beta_{t,k+1}}{\sum_{s=1}^{n-k-1} (d^s - d^{s-1})\beta_{s,k+1}} \\ &\quad \times \{\varphi_k^c - \alpha_k - (d-1)d^{n-k-1}\beta\} \\ &= \frac{\beta_{t,k+1}}{\varphi_{k+1}^c - \alpha_{k+1}} \{\varphi_k^c - \alpha_k - (d-1)d^{n-k-1}\beta\}, \end{aligned}$$

so that a ready induction yields

$$\beta_{t,k} = \beta \prod_{j=k}^{n-t-1} \frac{\varphi_j^c - \alpha_j - (d-1)d^{n-j-1}\beta}{\varphi_{j+1}^c - \alpha_{j+1}} \quad (r \leq k \leq n), \quad (8)$$

the usual convention on empty products allowing us to extend the range of validity of the above equation to $n-t \leq k \leq n$, as well.

To finish up, consider any vertex in generation k with $0 \leq k \leq r-1$. Such a vertex is necessarily contaminated and the subtree rooted at the vertex has d^{n-k} leaves consisting of d^{r-k} bots, $d^{r-k}(d-1)$ users in class U_1 , $d^{r-k}(d^2-d)$ users in class U_2 , ..., $d^{r-k}(d^{n-r}-d^{n-r-1})$ users in class U_{n-r} . Accordingly,

$$\varphi_k^c = \alpha_k + d^{r-k} \sum_{t=1}^{n-r} (d^t - d^{t-1})\beta_{t,k}.$$

As all d incoming flows to the vertex are contaminated, the vertex allocates bandwidth φ_k^c/d to each incoming flow of which a portion α_k/d is coopted by the attack segment of that flow. Thus, the legitimate users commingled in each contaminated flow have available a bandwidth $(\varphi_k^c - \alpha_k)/d$. Arguing as above, any class \mathcal{U}_t descendant of the tagged vertex hence receives a proportionate bandwidth

$$\beta_{t,k} = \frac{\beta_{t,k+1}}{\varphi_{k+1}^c - \alpha_{k+1}} \left\{ \frac{\varphi_k^c - \alpha_k}{d} \right\},$$

and, for $0 \leq k \leq r-1$, another easy induction shows

$$\beta_{t,k} = \frac{\beta_{t,r}}{d^{r-k}} \prod_{j=k}^{r-1} \frac{\varphi_j^c - \alpha_j}{\varphi_{j+1}^c - \alpha_{j+1}} = \frac{\beta_{t,r}}{d^{r-k}} \left(\frac{\varphi_k^c - \alpha_k}{\varphi_r^c - \alpha_r} \right) \quad (9)$$

as the product telescopes. Pooling our results we obtain

$$\begin{aligned} \beta_{t,k} &= \frac{\beta}{d^{r-k}} \left(\prod_{j=k}^{r-1} \frac{\varphi_j^c - \alpha_j}{\varphi_{j+1}^c - \alpha_{j+1}} \right) \\ &\quad \times \left(\prod_{j=k \vee r}^{n-t-1} \frac{\varphi_j^c - \alpha_j - (d-1)d^{n-j-1}\beta}{\varphi_{j+1}^c - \alpha_{j+1}} \right) \\ &= \frac{\beta}{d^{r-k}} \left(\frac{\varphi_{k \wedge r}^c - \alpha_{k \wedge r}}{\varphi_r^c - \alpha_r} \right) \\ &\quad \times \prod_{j=k \vee r}^{n-t-1} \frac{\varphi_j^c - \alpha_j - (d-1)d^{n-j-1}\beta}{\varphi_{j+1}^c - \alpha_{j+1}}, \end{aligned}$$

the usual conventions allowing us to extend the identity to all values of k . In particular, the bandwidth used at the root by a user in class \mathcal{U}_t is given by

$$\beta_{t,0} = \frac{\beta}{d^r} \left(\frac{\varphi_0^c - \alpha_0}{\varphi_r^c - \alpha_r} \right) \prod_{j=r}^{n-t-1} \frac{\varphi_j^c - \alpha_j - (d-1)d^{n-j-1}\beta}{\varphi_{j+1}^c - \alpha_{j+1}}$$

for $1 \leq t \leq n-r$. The specific form of these expressions depends on whether the agnostic or the pernicious drop protocol is in force.

VI. APPLICATIONS

The nature and impact of the DoS threat can vary substantially depending on the specification of the link capacities. Some examples in natural settings may serve to illustrate the extremes of behaviour.

EXAMPLE 1 (SOCIALISED CAPACITY) Suppose all links have the same capacity which we may assume, without loss of generality, to be unit: $C_0 = C_1 = \dots = C_n = 1$. For definiteness, consider a binary tree with $d = 2$. To ensure that no legitimate packets are lost under normal, average operating conditions, we require $\beta 2^n < C_0 = 1$. Accordingly, write $\beta 2^n = 1 - \epsilon$ where the positive ϵ denotes the capacity held in reserve at the root.

Edge flows: For uncontaminated edge flows we have from (1) that

$$\varphi_k^u = 2^{n-k}\beta = (1 - \epsilon)2^{-k} \quad (r+1 \leq k \leq n).$$

On the other hand, (2) and (3) show that all contaminated flows satisfy

$$\varphi_k^c = 1 \quad (0 \leq k \leq n).$$

Attack ratios under agnostic drops: From (4) we obtain

$$\begin{aligned} \alpha_k &= \prod_{j=k}^{n-1} \left(1 - \frac{\varphi_{j+1}^u}{\varphi_{j+1}^c} \right) = \prod_{j=k}^{n-1} (1 - 2^{n-j-1}\beta) \\ &= \prod_{i=0}^{n-k-1} (1 - 2^i\beta) \quad (r \leq k \leq n), \end{aligned}$$

while (5) shows that the attack flow emanating from each router remains unaltered in size for generations prior to r :

$$\alpha_k = \alpha_r \prod_{j=k}^{r-1} \{2 \wedge 1\} = \alpha_r \quad (0 \leq k \leq r-1).$$

It follows in particular that the attack flow emanating from the root satisfies

$$\begin{aligned} \alpha_0 &= \prod_{i=0}^{n-r-1} (1 - 2^i\beta) < \exp \left\{ -\beta \sum_{i=0}^{n-r-1} 2^i \right\} \\ &= \exp \{ -\beta(2^{n-r} - 1) \} = \exp \{ -(1 - \epsilon)2^{-r} + \beta \} \end{aligned}$$

the second step following by virtue of the elementary inequality $1 - x < e^{-x}$. The upper bound is in fact tight asymptotically and we may prove the following *proposition*: If $r = r_n$ increases unboundedly with n then $\alpha_0 \sim \exp \{ -(1 - \epsilon)2^{-r_n} \} = 1 - \mathcal{O}(2^{-r_n})$ as $n \rightarrow \infty$. Indeed, as $\log(1 - x) = -x + \mathcal{O}(x^2)$ as $x \rightarrow 0$, we have

$$\begin{aligned} \alpha_0 &= \exp \left\{ \sum_{i=0}^{n-r_n-1} \log(1 - 2^i\beta) \right\} \\ &= \exp \left\{ -\beta \sum_{i=0}^{n-r_n-1} 2^i + \mathcal{O} \left(\beta^2 \sum_{i=0}^{n-r_n-1} 2^{2i} \right) \right\} \\ &= \exp \{ -\beta(2^{n-r_n} - 1) + \mathcal{O}(\beta^2 2^{2(n-r_n)}) \} \\ &= \exp \{ -(1 - \epsilon)2^{-r_n} + \mathcal{O}(2^{-n} + 2^{-2r_n}) \} \end{aligned}$$

as $n \rightarrow \infty$, to complete the proof.

Thus, even very modest attacks with only $\log n$, or even $\log^* n$, cannily placed attackers, (corresponding to $r_n = \log \log n$ and $\log^* n$, respectively) will end up coopting all but a vanishing fraction $\mathcal{O}(2^{-r_n})$ of the bandwidth at the root even though the number of attackers is dwarfed by the exponential number of regular users.

Attack ratios under pernicious drops: Proceeding similarly, (6) yields

$$\begin{aligned} \alpha_k &= \bigwedge_{j=k}^{n-1} \{1 - 2^{n-j-1}\beta\} \wedge 1 \\ &= 1 - 2^{n-k-1}\beta = 1 - \frac{1}{2}(1 - \epsilon)2^{-k} \quad (r \leq k \leq n), \end{aligned}$$

and substituting α_r from this result into (7) gives

$$\begin{aligned}\alpha_k &= \bigwedge_{j=k}^{r-1} 2^{j-k} \wedge 2^{r-k} (1 - \frac{1}{2}(1 - \epsilon)2^{-r}) \\ &= 1 \wedge (2^{r-k} - \frac{1}{2}(1 - \epsilon)2^{-k}) \quad (0 \leq k \leq r-1).\end{aligned}$$

It follows in particular that

$$\alpha_0 = 1 \wedge (2^r - \frac{1}{2}(1 - \epsilon)) = \begin{cases} 1 - (1 - \epsilon)/2 & \text{if } r = 0, \\ 1 & \text{if } r \geq 1. \end{cases}$$

With pernicious drops, even a single attacker (corresponding to $r = 0$) will cause half the legitimate packets to be dropped, while two attackers (corresponding to $r = 1$), one placed on each half of the leaves of the tree, will coopt the entire bandwidth at the root.

It should be borne in mind, however, that each attack bot is allowed the capability of imposing an exponentially larger load than the typical user. This large initial attack capability of even a few bots is fostered by the substantial excess capacity in the downstream links. ■

The identical link capacities of the previous example result in all contaminated links operating at maximum capacity. For an illustration in the opposite direction, consider the following familiar setting.

EXAMPLE 2 (DIGITAL TELEPHONY TREES) If each router in the tree has sufficient capacity to simultaneously handle all incoming flows at maximum intensity we are led to a model where link capacities increase exponentially towards the sink. This is the model for digital telephony in the DSL hierarchy and the Bell Labs T1, T3, ... system. Accordingly, suppose

$$C_k \geq dC_{k+1} \quad (0 \leq k \leq n-1).$$

If the load offered by legitimate users is bounded we may suppose without loss of generality that the maximum offered load is unit with the mean load $\beta < 1$. In this setting we may normalise by setting $C_n := 1$, whence $C_k \geq d^{n-k}$ for each k .

In this model, each router will accept its entire offered load even under worst case settings in normal operation. It follows that the edge flows are given by

$$\varphi_k^c = d^{(r-k)_+} + (d^{n-k} - d^{(r-k)_+}) \quad (0 \leq k \leq n),$$

as we may also verify directly from (2) and (3). In particular,

$$\varphi_0^c = d^r + (d^n - d^r)\beta.$$

As all incoming flows are accepted whole at every router the same situation obtains for both agnostic and pernicious drops as can be verified by running through the corresponding pairs of equations (4, 5) and (6, 7). The corresponding attack flows hence satisfy

$$\alpha_k = d^{(r-k)_+} \quad (0 \leq k \leq n),$$

and, in particular,

$$\alpha_0 = d^r.$$

The fraction of the bandwidth coopted by attack bots at the root is hence given by the attack ratio

$$\rho_0 = \frac{\alpha_0}{\varphi_0^c} = \frac{d^r}{\beta d^n + (1 - \beta)d^r}.$$

For any fixed $0 < \beta < 1$, if $r = r_n \leq \kappa n$ for any $\kappa < 1$ then $\rho_0 = \mathcal{O}(d^{-(1-\kappa)n})$ decreases exponentially. Indeed, even if $r = r_n = n - \omega_n$ where ω_n increases slowly but without bound, $\omega_n \rightarrow \infty$, $\omega_n = o(n)$, then $\rho_0 = \mathcal{O}(d^{-\omega_n})$ and the attack flow at the root is an asymptotically negligible fraction of the the total traffic with most of the legitimate packets making it through.

Thus, if link capacities increase exponentially towards the sink then DoS effects are minuscule unless a constant fraction of the users are compromised [i.e., $n - r_n = \mathcal{O}(1)$]. As is intuitive, the limited link capacities at the downstream vertices functions as a gatekeeping mechanism which limits the severity of the initial attack—the ratio of the bandwidth accessed by an attack bot to that of the mean bandwidth occupied by a legitimate user, $1 : \beta$, may be large but is bounded. ■

The link capacity assignments in the previous example are conservative in that under attack-free conditions there is sufficient capacity to simultaneously handle the worst-case load that can be offered simultaneously by all legitimate users. In a somewhat less generous assignment of capacities we may seek to exploit statistical information about offered loads so that only rare load excursions lead to packet drops with typical loads handled effortlessly.

EXAMPLE 3 (STATISTICAL OVERPROVISIONING) Let X be a positive, bounded random variable with mean β and variance σ^2 . We shall suppose, without loss of generality, that X is bounded above by 1 so that β and σ^2 are both positive and less than one. Suppose the loads offered by legitimate users are independent random variables all with the common probability law of X . As each router in generation k has d^{n-k} leaves as its descendants, the total load S_k proffered by these descendants is the sum of d^{n-k} independent copies of X . It follows that $E(S_k) = \beta d^{n-k}$ and Bernstein's inequality suggests that large excursions above this value are unlikely:

$$\begin{aligned}\mathbf{P}\{S_k > \beta d^{n-k} + A\sigma d^{(n-k)/2}\} \\ \leq \exp\{-\frac{1}{2}A^2 / (1 + \frac{1}{3}Ad^{-(n-k)/2})\}.\end{aligned}$$

The right-hand side can be made less than δ , say, by choice of the positive constant A . Accordingly, the choice of link capacities

$$C_k = \beta d^{n-k} + A\sigma d^{(n-k)/2} \quad (0 \leq k \leq n)$$

will ensure that the fraction of time that any given router sees a buffer overflow under normal, attack-free conditions may be made as small as desired.

Edge flows: With these link capacity assignments the terms under the minimum in (2) are decreasing whence

$$\varphi_k^c = C_n + (d^{n-k} - 1)\beta = \beta d^{n-k} + A\sigma \quad (r \leq k \leq n). \quad (10)$$

On the other hand, the terms under the first minimum on the right of (3) are increasing and accordingly

$$\begin{aligned} \varphi_k^c &= C_k \wedge (d^{r-k} \varphi_r^c) = \beta d^{n-k} + A\sigma d^{\min\{r-k, (n-k)/2\}} \\ &= \begin{cases} \beta d^{n-k} + A\sigma d^{r-k} & \text{if } (2r-n)_+ \leq k \leq r-1, \\ \beta d^{n-k} + A\sigma d^{(n-k)/2} & \text{if } 0 \leq k < (2r-n)_+. \end{cases} \end{aligned} \quad (11)$$

In particular, the edge flow emanating from the root satisfies

$$\varphi_0^c = \beta d^n + A\sigma d^{\min\{r, n/2\}}.$$

Attack ratios under agnostic drops: As all incoming flows to routers in generations $r \leq k \leq n$ are accepted (as is easy to verify from (4)), we have

$$\alpha_k = C_n = \beta + A\sigma \quad (r \leq k \leq n).$$

On the other hand, (10) and (11) show that

$$d\varphi_{j+1}^c \begin{cases} \leq C_j & \text{if } (2r-n)_+ \leq k \leq r-1, \\ > C_j & \text{if } 0 \leq k < (2r-n)_+, \end{cases}$$

and, accordingly, (5) shows that

$$\alpha_k = \alpha_r d^{r-k} \quad \text{if } (2r-n)_+ \leq k \leq r-1,$$

while

$$\alpha_k = \alpha_r d^{r-(2r-n)_+} \prod_{j=k}^{(2r-n)_+-1} \frac{C_j}{\varphi_{j+1}^c} \quad \text{if } 0 \leq k < (2r-n)_+.$$

It will be convenient to set

$$\theta_j := \frac{A\sigma d^{(n-j-3)/2} (\sqrt{d}-1)}{\beta d^{n-j-1} + A\sigma d^{(n-j-1)/2}}.$$

Then we may compact our expressions into the form

$$\alpha_k = (\beta + A\sigma) d^{r-k} \prod_{j=k}^{(2r-n)_+-1} (1 + \theta_j) \quad (0 \leq k \leq r-1), \quad (12)$$

with the usual convention on empty products taking care of the cases $k \geq (2r-n)_+$. Two cases are indicated.

Low intensity attacks. If $r \leq n/2$ then

$$\varphi_0^c = \beta d^n + A\sigma d^r, \quad \alpha_0 = (\beta + A\sigma) d^r,$$

and all flows are accepted without truncation at any generation as in the previous example, the difference only in that the entering flows at the leaves are restricted

to $C_n = \beta + A\sigma < 1$. It follows that the attack ratio at the root is given by

$$\rho_0 = \frac{\alpha_0}{\varphi_0^c} \sim \left(1 + \frac{A\sigma}{\beta}\right) d^{-(n-r)} \quad (n \rightarrow \infty).$$

There are modest gains over the previous case (in a multiplicative factor of $1 + A\sigma/\beta$ instead of $1/\beta$) if the mean and standard deviation of the legitimate flows are comparable and both much less than the load maximum 1, this kind of large excursion being rare.

High intensity attacks. Suppose now that $r = r_n$ increases with n so that $\liminf r_n/n > 1/2$ and $\limsup r_n/n \leq 1$. Then,

$$\varphi_0^c = \beta d^n + A\sigma d^{n/2} \sim \beta d^n.$$

Now for $0 \leq j \leq 2r-n-1$, we have

$$\theta_j \sim \frac{A\sigma}{\beta\sqrt{d}} (\sqrt{d}-1) d^{-(n-j)/2} \quad (0 \leq j \leq 2r-n-1).$$

Taking logarithms of the product on the right of (12), we hence obtain

$$\begin{aligned} \sum_{j=0}^{2r-n-1} \log(1 + \theta_j) &\sim \frac{A\sigma}{\beta\sqrt{d}} (\sqrt{d}-1) d^{-n/2} \sum_{j=0}^{2r-n-1} d^{j/2} \\ &= \frac{A\sigma}{\beta\sqrt{d}} (d^{-(n-r_n)} - d^{-n/2}) \sim \frac{A\sigma}{\beta\sqrt{d}} d^{-(n-r_n)}. \end{aligned}$$

It follows that

$$\alpha_0 = (\beta + A\sigma) d^{r_n} \exp\left(\frac{A\sigma}{\beta\sqrt{d}} d^{-(n-r_n)} (1 + o(1))\right)$$

as $n \rightarrow \infty$. The attack ratio at the root hence satisfies

$$\rho_0 \sim \left(1 + \frac{A\sigma}{\beta\sqrt{d}}\right) d^{-(n-r_n)} \exp\left(\frac{A\sigma}{\beta\sqrt{d}} d^{-(n-r_n)} (1 + o(1))\right).$$

Write $r_n = n - \omega_n$. If $\omega_n \rightarrow \infty$ and $\omega_n = o(n)$ then $\rho_0 \sim (1 + \frac{A\sigma}{\beta\sqrt{d}}) d^{-\omega_n}$ and again the attack flow at the root is asymptotically negligible. Even much larger attacks do not take over the entire flow. Thus, if $r_n = n - K$ for any fixed K then $\rho_0 \rightarrow (1 + \frac{A\sigma}{\beta\sqrt{d}}) d^{-K} \exp(A\sigma d^{-K}/\beta\sqrt{d})$ and at most a constant fraction of the flow is coopted. The gains are largest when the flows are highly concentrated around the mean so that $\sigma/\beta \ll 1$. This situation should be compared with that of the previous example.

Attack flows under pernicious drops: The sequence of values $C_j - (d-1)\varphi_{j+1}^u$ is decreasing as j increases from r to $n-1$ and in consequence (6) shows that

$$\alpha_k = C_n \quad (r \leq k \leq n).$$

Likewise, for any $0 \leq k \leq r-1$, the sequence of values $d^{j-k} C_j$ increases as j increases from k to $r-1$. Accordingly, (7) yields

$$\alpha_k = C_k \wedge (d^{r-k} C_n) \quad (0 \leq k \leq r-1).$$

In particular, the attack flow emanating from the root satisfies

$$\alpha_0 = \{\beta d^n + A\sigma d^{n/2}\} \wedge \{(\beta + A\sigma)d^r\}.$$

Low intensity attacks. If $r = r_n = n - \omega_n$ where ω_n satisfies $\omega_n \rightarrow \infty$ and $\omega_n = o(n)$ then $\alpha_0 = (\beta + A\sigma)d^{r_n}$ eventually and the attack ratio at the root is given by

$$\rho = \frac{(\beta + A\sigma)d^{r_n}}{\beta d^n + A\sigma d^{r_n}} \sim \left(1 + \frac{A\sigma}{\beta}\right) d^{-(n-r_n)}$$

which matches the agnostic drop result. Thus, under a low intensity attack, an asymptotically negligible fraction of the bandwidth is coopted by the attack even under a pernicious drop protocol.

High intensity attacks. Suppose now that $r_n = n - K$ for some fixed K . In this case, all but a fixed fraction d^{-K} of the leaves is compromised. Then, for all sufficiently large n ,

$$\alpha_0 \sim \{\beta \wedge (\beta + A\sigma)d^{-K}\} d^n,$$

so that the attack ratio at the root is given asymptotically by

$$\rho_0 \rightarrow 1 \wedge \left(1 + \frac{A\sigma}{\beta}\right) d^{-K} \quad (n \rightarrow \infty).$$

If the deviations from the mean are not very small, the entire bandwidth gets coopted. If σ is comparable to β then for all sufficiently large K an asymptotically fixed fraction of the bandwidth is coopted. ■

VII. IN FINE

Even in this stylised setting, a key conclusion that emerges from the analysis is that, as is intuitive, there may be distinct advantages in combating DoS attacks to allow of early gatekeeping mechanisms at downstream routers by limiting ingress capacities at the leaves. The analysis presented here is, of course, simplistic in that it ignored random fluctuations in user traffic. It may be anticipated, however, that there is concentration of the results around the expressions obtained by using the mean values for the loads—at least for “reasonable” assignments of link capacities—; we reserve these considerations for elsewhere. In another direction, more complex network models pose more challenges in analysis because of a multiplicity of routes that become available between sender and receiver and the presence of crossing traffic which has the nuisance value of creating fluctuations in demand in the sub-net under consideration. The analysis now depends intimately on the topology, the routing protocols in force, and the nature and composition of the incidental traffic, as well as on the deployment of bots. We leave these issues for a later date.

REFERENCES

- [1] D. Dolev and A. C. Yao, “On the security of public-key protocols,” *IEEE Trans. Inform. Theory*, vol. 29, pp. 198–208, 1983.
- [2] C. A. Gunter, S. Khanna, K. Tan, and S. S. Venkatesh, “DoS Protection for Reliably Authenticated Broadcast,” *11th Annual Network and Distributed System Security Symposium*, San Diego, California, February 2004.
- [3] M. Sherr, M. B. Greenwald, C. A. Gunter, S. Khanna, and S. S. Venkatesh, “Mitigating DoS Attack Through Selective Bin Verification,” *Proc. Workshop on Secure Network Protocols (NPsec)*, Boston, Massachusetts, November 2005.