Brief Announcement: The Overlay Network Content Distribution Problem

Chip Killian, Michael Vrable, Alex C. Snoeren, Amin Vahdat, and Joseph Pasquale
University of California, San Diego
{ckillian,mvrable,snoeren,vahdat,pasquale}@cs.ucsd.edu

Many overlay multicast protocols have been designed and deployed across the Internet to support content distribution. To our knowledge, however, none have provided a rigorous analysis of the problem or the effectiveness of their proposed solutions. We define the *Overlay Network Content Distribution* (OCD) problem to allow such analyses.

Categories and Subject Descriptors: C.2.1 Network Architecture and Design: Packet-switching networks C.2.2 Network Protocols: Routing protocols
General Terms: Algorithms, Design, Performance
Keywords: overlay networks, content routing

Problem Description. We assume, without loss of generality, that all content (e.g. files) is subdivided as a set of unit-sized *tokens*. Tokens start out at one or more nodes (senders), and the goal is to transfer them to another set of nodes (receivers). For simplicity we consider topologies that consist of end hosts all capable of storing and duplicating tokens. We also assume that link capacities are constant and that latency does not change with load.

The input is a simple, weighted directed graph $G = (V, E)$, a set of tokens T, and two functions $h : V \to 2^T$ and $w : V \to 2^T$. Here, 2^T denotes the power set of T. Let $c : E \to \mathbb{N}$ denote the weight (capacity) function of each arc. The h (have) function denotes the initial set of tokens that each vertex possesses. The w (want) function indicates the tokens each vertex wants.

We define a move as an assignment of a token to an arc, and a *timestep* to be a set of simultaneous moves. A *distribution schedule* proceeds as a sequence of timesteps. At each timestep, the tokens assigned to an arc are limited by capacity $c(u, v)$, and the set of tokens available from the incident vertex at the start of the timestep.

More formally, let $t \in \mathbb{N}$ denote the length of a schedule (number of timesteps). The schedule is defined by a collection of functions $s_i : E \to 2^T$, where $0 \leq i < t$; the function s_i gives the set of tokens sent across arc (u, v) during timestep i. A schedule requires a set of functions $p_i : V \to 2^T$ for $0 \leq i \leq t$ that specify which tokens a vertex possesses at each timestep, and is subject to the following restrictions:

$$p_i(v) = \bigcup_{u \in V : (u,v) \in E} p_{i-1}(v) \cup s_{i-1}(u, v)$$

Additionally, $p_0(v) = h(v)$ (initial), $|s_i(u, v)| \leq c(u, v)$ (capacity), and $s_i(u, v) \subseteq p_i(u)$ (possession). A schedule is said to be *satisfying* if $w(v) \subseteq p_t(v)$ for all vertices $v \in V$. There are two obvious characteristics of satisfying schedules that are of interest: how many timesteps they contain, and their total bandwidth consumption.

File distribution times. Define the *Fast Overlay Content Distribution* (FOCD) problem as determining a satisfying distribution schedule of minimum length, τ. The problem is satisfiable if there exists some τ for which this is true. The *Decisional FOCD* (DFOCD) problem takes as input a FOCD problem and τ^*, and determines whether the OCD problem is satisfiable in τ^* steps. Here, τ is the maximum completion time over all content retrievals. This metric is often referred to as *makespan* in the scheduling literature. We show that DFOCD and FOCD are in NPC by reduction from Dominating Set. However, we note that since this reduction reduces Dominating Set problems to instances of FOCD problems with tokens from distinct distributions (different set of sources and/or receivers), it does not prove the single content/file distribution problem is in NPC.

Bandwidth constraints. A variation on FOCD, *Efficient Overlay Content Distribution* (EOCD), is defined as determining a satisfying distribution schedule with minimum bandwidth ($\sum_{\substack{0 \leq i < t \\ (u,v) \in E}} |s_i(u, v)|$). We show that EOCD is in NPC by reduction from the generalized Steiner tree problem. We have developed a formulation of EOCD as a time-indexed integer program; full details can be found in the TR.

On-line approximation using local knowledge. To make OCD and variants more realistic, we consider limitations stemming from lack of global knowledge for making decisions and for coordinating actions. *Local-knowledge Overlay Content Distribution* (LOCD) is defined as follows: Let $k_i(v)$ denote the *knowledge* of vertex v at the start of timestep i. We require that $k_0(v)$ be computed by a deterministic function of the neighbors of vertex v, the capacity of each edge incident upon v, $h(v)$, and $w(v)$. At timestep i, for any edge (v_1, v_2), $s_i(v_1, v_2)$ must be a function only of $k_i(v)$. $k_{i+1}(v)$ must then be computable by a deterministic function that takes as input only previous knowledge ($k_i(v)$) and knowledge acquired from neighbors ($k_i(u)$). The local-decision problem suggests these questions: Will all tokens eventually be adequately distributed? Is there a relative bound on the time taken? Are there (time) bounds independent of the algorithm chosen? We show that there exists no c-competitive local on-line algorithm for FOCD for any fixed constant c.

A full version of this paper is available as University of California, San Diego technical report CS2005-0824 from `http://www.cs.ucsd.edu/~snoeren`.

PODC'05, July 17–20, 2005, Las Vegas, Nevada, USA.
ACM 1-59593-994-2/05/0007.