

Backpressure Scheduling

Tassiulas & Ephremides ¹⁹⁹³

MaxWeight Algorithm: a policy to schedule links over time in such a way that any arrival rate vector in the stability region of the network

can be scheduled without causing any queue to become unstable.

Consider a given network, represented by a graph $G = (V, E)$ and a conflict graph $H = (V_H, E_H)$ where $V_H = E$, E_H is set of edges between "nodes" $E \in V_H$ s.t. the corresponding links have a conflict.

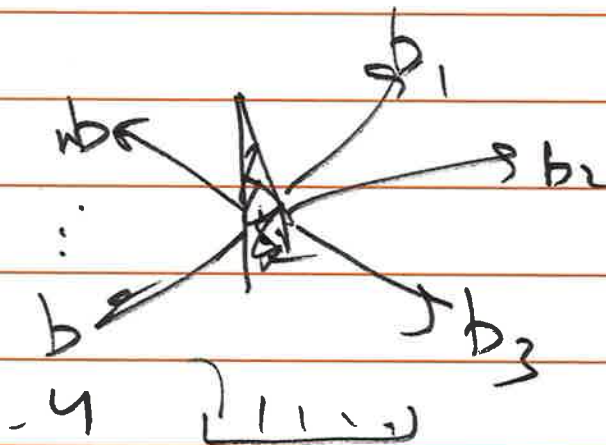
Scott Moeller (PhD student at USC) developed BCP - Backpressure Collection Protocol. A routing protocol for data collection in wireless sensor networks, that is derived from / inspired by backpressure scheduling (MaxWeight).

BCP algorithm/protocol:

each node in a distributed fashion, estimates ETX_{ij} & in every packet, includes in the header its own (Q_i) queue size & by logging neighbor's packets, figures out (Q_i)

There is only 1 queue at each node, because all nodes are sending to the same sink

A special case where MaxWeight is easy: single transmitter.



e.g. downlink channel

for this setting:
MaxWeight =

serve the ~~most~~ K -largest

throughput/users if
 K is the # of available channels.

if links are the same

Else serve

The K largest throughputs ~~at~~
with respect to $\underbrace{Q_i}_S \times \underbrace{R_i}_{\text{link rate}}$

There are a set of
(say unicast) flows in the network
(1 ... n) w/ corresponding
sources & destinations:
 $(s_1, t_1), (s_2, t_2) \dots (s_n, t_n)$

The arrival rate vector is a
set of rates for each flow:
 (r_1, \dots, r_n)

$r_i \leftarrow$ how many pkts/sec.

arrive (on average) at source
 s_i , to be sent to t_i .

For simplicity assume exactly
1 pkt can be sent on each
slot on any link.


A scheduling policy, at each time,
can pick a set of non-interfering links,

and more at most one pkt
on each of those links.

At each node, we allow there
to be n different queues:
 $Q_i^1, Q_i^2, \dots, Q_i^n$

The scheduling policy also gets
to decide, on link $i-j$

which flow's packet should be
served.




assume the conflict graph is a fully-connected /
complete graph, i.e. every link interferes w/
every other link.

At the sink for each flow,
 i.e. at t_i , the pressure
 is assumed to be always 0.
 i.e.: $Q_{t_i}^i = 0$

Recall all queues are unlimited
 in size (ideal assumption),
 however we wish to avoid
 queues growing without bound

(i.e. becoming unstable).



MaxWeight can guarantee stability
 $\forall r \in \Delta$. Sometimes called
 a "Throughput-Optimal" policy.

The MaxWeight Scheduling Policy:

consider independent sets S —
a set of edges that do not
conflict w/ each other, i.e. can be
co-scheduled.

pick the independent set S
at each time, which maximizes
the weight:

$$\max_S \sum_{e_{ij} \in S} w_{ij}^*$$

link
rate of the
edge $i-j$

where

$$w_{ij}^* = \max_f (Q_i^f - Q_j^f, 0) \cdot R_{ij}$$

(queue backpressure
or queue differential)

Repeat the above selection at
each time.

Properties of Max-Weight:

1. guarantees stability $\forall r \in \Delta$
2. works for stochastic arrivals
3. works for any size network, any conflict definition, any set of flows.
(not necessarily unicast-)

Requirements (Idealized):

- needed for analysis, not for implementation
- unlimited queue size
 - Global time synchronization
 - Calculating the maximum weight independent set is NP-hard
 - Centralized scheduling that needs global information about all queues at all times!

Each node i , computes for each ~~at~~ neighbor j , the weight-
link cost

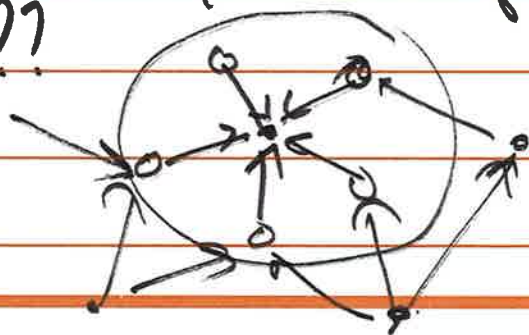
$$w_{ij} = Q_i - Q_j - V \cdot ETX_{ij}$$

queue backpressure \uparrow a tunable parameter

If $w_{ij}^* = \max_{j \in \text{neighbors}} w_{ij} > 0$

then node i sends its next pkt to j^* , i.e. the nbr j that maximizes w_{ij} .

Note that the above description makes no reference to scheduling multiple links w/o interference — why??




(where)

BCP is a layer 3 protocol
- frames just in which
neighbor each node should
forward a pkt to.

It lets L2 (MAC/link layer)
use CSMA to take care
of scheduling/avoiding interference.
(when to send)

[Now, do not need global
view or coordination or
time sync.]

calculation at each node is
trivial (max)



say $\epsilon_{T \times i j} = 1 \forall \text{ links}$
 $V = 2$

$$w_{ij} = \phi_i - \phi_j - 2 \cdot 1$$

$$\phi_i - \phi_j - 2$$

for sending a pkt on ij :

$$w_{ij} > 0 \quad \phi_i - \phi_j > 2$$


$$\phi_i > \phi_j + 2$$

FIFO queueing is the standard queue service scheme in networks

∴ FIFO queueing BCP has a high delay at low arrival rate!

Station:

FIFO → LIFO



completely solves the delay problem.

BUT ... a set of

pkts stay stuck forever.

The early control pkts that are used to learn the gradients will be stuck.

Can show:

stuck pkts at node i

$$= V \cdot \text{ETX}_{i \text{ to } \text{sink}}$$

↑
shortest path ETX cost.

Problem of Generalizing BCP:

for any to any (unicast) traffic, we would need the number of queues at each node = # of flows in the network

worst case:
N nodes, # flows = $O(N^2)$

$$w_{ij} = \phi_i - \phi_j - \sum_{f \in T_{ij}} x_{ij}^f$$

$$w_{ij}^f = \phi_i^f - \phi_j^f - \sum_{T \in T_{ij}} x_{ij}^f$$

$$w_{ij}^* = \max_{j, f} w_{ij}^f$$

doesn't

scale well
for large networks

Open problem! Can this be solved?

- Shortest ETX routing
- Anypath routing
- Cooperative Broadcast
- Backpressure scheduling.

next class: routing in mobile networks
transport layer for wireless networks
is TCP good enough?

final exam is on August 8