The Internet relies on cooperative endpoints to react to signals from the network that congestion is occurring. In particular, TCP interprets packet loss as a signal of congestion. However there are many new non-cooperative protocols in use which attempt to exploit the network aggressively and do not reduce their demands when the network signals congestion. We propose the aggregate control of ‘‘fluxes” defined by policies at individual routers. Each router can then calculate an optimal allocation of bandwidth to each flux contending for a given output link. We propose a combined hill climbing and convex programming method for this optimization, which we call HCCP. HCCP is designed to punish greedy fluxes rather than just regulating them: such fluxes may find their bandwidth allocation reduced to zero if they are sufficiently aggressive. Our results show that HCCP is effective at regulating a wide range of rather generally characterized transport protocols. We explore the use of both throughput maximization and proportionally fair allocation and recommend the latter because the former often leads to the situation where one or more fluxes receive zero bandwidth.