FairCache: Introducing Fairness to ICN Caching Technical Report

Liang Wang[1]   Gareth Tyson[2]   Jussi Kangasharju[3]   Jon Crowcroft[1]

[1]University of Cambridge, UK   [2]Queen Mary University London, UK   [3]University of Helsinki, Finland

Abstract—**Information-centric networking (ICN) is a popular research topic. At its heart is the concept of in-network caching. Various algorithms have been proposed for optimising ICN caching, many of which rely on collaborative principles, *i.e.* multiple caches interacting to decide what to store. Past work has assumed altruistic nodes that will sacrifice their own performance for the global optimum. We argue that this assumption is flawed. We address this problem by modelling the in-network caching problem as a Nash bargaining game. We develop optimal and heuristic caching solutions that consider both performance and *fairness*. We argue that only algorithms that are fair to all parties will encourage engagement and cooperation. Through extensive simulations, we show our heuristic solution, FairCache, ensures that all collaborative caches achieve performance gains without undermining the performance of others.**

Index Terms—**Information-centric networking, network protocols, resources allocation, game theory, algorithm design, optimisation.**

I. INTRODUCTION

Information-Centric Networking (ICN) [1], [2] has been proposed to exploit the observation that much of today's Internet traffic is content distribution. ICN replaces the existing location-based Internet model with a content request/response model. One feature this enables is the capacity to cache content within the network. Whereas initial ICN caching approaches used traditional algorithms (*e.g.* Least Recently Used), there has been a number of novel proposals that attempt to specifically target ICN environments. These algorithms exploit things like inter-AS cooperation, request prediction and a priori topology maps to optimise performance [3]–[5].

A key outcome of this work has been the observation that collaborative caching usually outperforms locally optimised algorithms [5]–[8]. This is primarily caused by the nature of ubiquitous ICN caching, where nearby caches will often wastefully store the same objects [7]. To address this, a simple collaborative algorithm might involve two nodes strategically caching distinct objects [4], [9], [10]. Cache collaboration is therefore likely to play a role in any future ICN deployments [11]. In tandem, we are witnessing a fragmentation of cache ownership in the Internet, with large operators deploying separate infrastructures. Some of these providers exclusively host their own content (*e.g.* Google, Netflix), whilst other aggregate content from multiple sources (*e.g.* Akamai, ChinaCache). This adds an extra layer of complexity as its means that, even on an intra-domain level, we may begin to see multiple competing stakeholders operating caches within a single network. This will likely be accelerated by the growth of network function virtualisation, which will allow anybody to "spin up" caches within a network (we already see the availability of virtual cache services, *e.g.* Fortinet Virtual Cache). Hence, ensuring the cache collaboration can also work in this setting will become increasingly important.

A more extreme example of this fragmentation is within the expanding number of wireless community mesh networks, *e.g.* Guifi [12]; these are deployed by groups of individuals who each contribute wireless routers (*e.g.* mounted on their property). In a community network, *every* router/cache is operated by a separate individual. Hence, we predict that future ICNs will use caches that are provisioned not just by network operators, but also various distinct stakeholders at strategic locations. These observations, however, have the potential to undermine the key tenets of caching in ICNs: What if caches operated by separate entities pursue policies that do not include collaboration, the storage of competitor's content or the serving of specific users? This is currently the situation online today, and it is unlikely to change with the advent of ICN. Despite this, most ICN collaborative algorithm assumes altruistic nodes that simply strive to reach a global optimum [4], [5], [7], [8], [10], [13]–[15]. Whereas this is acceptabe in scenarios where a single organisation operates all caches, it ceases to be suitable in inter-domain scenarios or cases where caches in a single network are operated by multiple (third party) providers.

The reasons why a non-collaborative policy may be implemented are diverse. However, in this paper, we explore the topic from a utilitarian perspective. Intuitively, caches would wish to engage in a collaborative algorithm if they attain greater utility than if they were not to engage. This observation mandates some concept of *fairness*, where benefits are spread fairly across caches, and individuals are not expected to sacrifice personal performance by collaborating. Imagine, for instance, the above community network example; an individual who sees his/her own performance consistently detrimented by collaboration would (rationally) cease. We therefore argue that collaboration should be based on fairness, which may or may not reduce global performance. While a global optimum sounds attractive, we argue it is more important, from a practical perspective, that every node is better off by collaborating together than working alone. In this paper, we design a collaborative caching algorithm that embraces both high performance and fairness. Our focus is not to build a protocol that forces nodes to collaborate, or provides protection against malicious behaviour but, rather, to design underlying algorithms that can fairly share cache space across trusted collaborators. We first formulate the fair in-network

caching problem as a Nash bargaining game (§III) before describing optimal algorithms for allocating objects to caches (§IV). We then propose a heuristic collaborative caching algorithm (§V) with fairness at its core: FairCache. Through extensive simulations, we show that FairCache achieves in excess of 90% accuracy compared to the optimal solution, at a fraction of the overheads (§VII). Importantly, we show that, when using FairCache, *all* nodes improve their performance via cooperation. It can be deployed across small subsets of collaborating caches or, alternatively, globally without change to design. We conclude by extracting key lessons learnt (§X).

II. MOTIVATION AND SYSTEM MODEL

To underpin our design, we begin with a motivational example before outlining our system model. For convenience, Table I contains the notations used throughout the paper.

A. Motivational Example

We use the simple toy caching system described in Figure 1 as a motivating example. Imagine two routers with a cache capacity of one object. They each serve a nearby set of users and, consequently, it is desirable that they collaborate to decide which objects should be cached (*e.g.* to avoid storing the same object). To decide which object to store, the caches locally inspect the request rates they receive, depicted in Figure 1 (as a Demand Matrix). Intuitively, each cache would wish to selfishly optimise some concept of individual "utility". For simplicity, we measure their utility as the number of cache hits they get. We also allow nodes to redirect requests to the other cache; if a hit is attained there, a utility of 0.5 is given to the node performing the redirect (factored down due to the extra delay, overhead etc.). We consider three caching strategies:

Case 1: Greedy Strategy, where each cache locally and selfishly optimises its performance. As our comparison baseline, Greedy strategy is a perfect LFU that keeps track of all the objects. Cache 1 chooses to hold A since it is the most popular content of demand 90, which leads to $U_1 = 90 + \frac{31}{2} = 105.5$. Similarly, Cache 2 caches B which leads to $U_2 = 83 + \frac{5}{2} = 85.5$. Therefore, we have the aggregated utility $U_{Total} = U_1 + U_2 = 105.5 + 85.5 = 191$.

Case 2: Global Strategy, where each cache tries to maximise the aggregated utility U_{Total} of the whole system. By caching C and D on Cache 1 and 2 respectively, U_{Total} reaches its theoretical maximum, namely $U_{Total} = 126 + 85 = 211$. However, if we examine the individual performance and compare them to the Greedy Strategy, we notice that the increase in utility for Cache 1 results in a utility decrease for Cache 2.

Case 3: Fair Strategy, where caches attempt to collaborate fairly, in a way that does not reduce utility for any party. Cache 1 stores E and Cache 2 stores F. Although this does not achieve the global optimum (*i.e.* the aggregated utility U_{Total} drops from 211 to 208.5), it ensures that *both* caches improve their respective performance whilst also improving upon the local Greedy Strategy. This solution is Pareto efficient, and ensures both parties are incentivised.

The above reveals a stark mismatch. Attaining a global optimum often disadvantages some parties [16]. Thus, nodes that
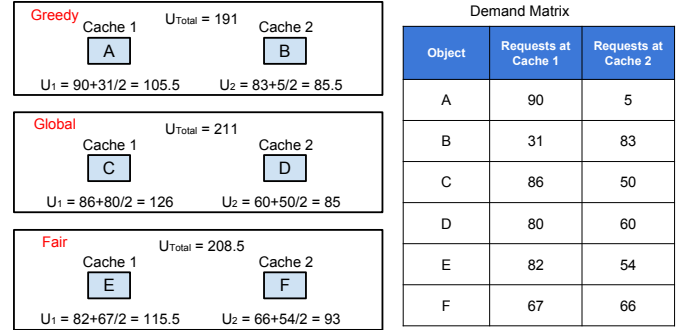

Fig. 1: A mini caching system with two caches and six objects. Three strategies (Greedy, Global and Fair) are presented.

Notation	Meaning
$G = (V, E)$	a graph G of node set V and edge set E
O	content set, o_k represents the k^{th} item
O'	a reduced content set of O by removing unpopular items
$\mathbf{s}$	s_k is the size of the k^{th} content item o_k
$\mathbf{w}$	demand matrix, $w_{i,k}$ is the demand of o_k on v_i
$\mathbf{U}$	utility vector, U_i is the utility of node v_i
$\mathbf{u^0}$	initial disagreement vector, u_i^0 is the disagreement value of v_i
$\mathbf{x}$	$x_{i,k}$ represents the decision whether v_i caches o_k locally
$\mathbf{y}$	$y_{i,j,k}$ represents the decision whether v_i retrieves o_k from v_j
$\mathbf{\Omega}$	solution space of all caching games, $\mathbf{\Omega^e}$ is the Pareto frontier
$\mathbf{\Psi}$	solution space of all fair caching games
$\mathbf{\lambda}$	dual variable associated with the constraint (7): $y_{i,j,k} \leq x_{j,k}$
$\mathcal{L}(\cdot)$	Lagrangian associated with the objective function (4)
$d(\cdot)$	Lagrangian dual function of the objective function (4)
$h(\cdot)$	subgradient of the dual function $d(\cdot)$
$\mathbf{\Phi}$	overall communication complexity at system level
N_i	neighbourhood of v_i, average size is denoted as $|\overline{N}|$
N_i^+	set of nodes having v_i in their neighbourhoods
r	average search radius, r_i uniquely defines N_i of v_i
n_1, n_2	average size of one-hop and two-hop neighbourhood
t	current iteration index while running subgradient descent
ξ	ξ_t is the step size of a subgradient algorithm at iteration t
θ', θ	current utility improvement and the stopping threshold

TABLE I: Table of main notation used in the paper

are unfairly exploited by other caches' redirects (at the cost of their own performance) are unlikely to continue collaboration: Caching should balance the need for high performance against the need for fair usage across caches.

B. System Model

We model the network as a graph, $G = (V, E)$, where V is the set of nodes and E is the set of edges. V could consist of all caches in a network or, alternatively, a subset of collaborating partners. These could be owned by one or more separate organisations. We follow an NDN [1] model, whereby hosts generate requests that get deterministically routed to sources that reply with content objects. Each node in the network, $v_i \in V$, is equipped with cache of size C_i. We denote O as the global set of content objects. For each $o_k \in O$, we associate two parameters: s_k, which is the object size and $w_{i,k}$, which is its aggregated demand (requests per second) observed from all the clients connected to v_i. We focus on a subcategory of caching: collaborative algorithms. Because of resource constraints, we assume that nodes are limited in the number of nearby nodes they can cooperate with. We

use r_i to represent v_i's search radius measured in hops. r_i defines a neighbourhood for each v_i, which we denote as $N_i = \{v_j | l_{i,j}^* \leq r_i, \forall v_j \in V, v_i \neq v_j\}$, where $l_{i,j}^*$ measures the length of the shortest path between v_i and v_j.

A collaborative caching algorithm can be decomposed into "caching decisions" and "retrieval decisions". These two parts solve "what to cache" and "where to fetch". The latter is necessary to allow nodes to redirect requests to other caches (opposed to forwarding it along the default route to the original source). This means that caches that do not locally store an object retain the flexibility to exploit objects stored elsewhere (*i.e.* collaboration). To model such a caching strategy, we use two vector decision variables: $\mathbf{x}$ and $\mathbf{y}$. $x_{i,k} \in \{0, 1\}$ denotes whether v_i caches o_k, and $y_{i,j,k} \in \{0, 1\}, \forall i \neq j$ denotes whether v_i retrieves the object o_k from v_j. Formally, we say:

Definition 1. *A caching strategy for a network G is a tuple of functions $(\mathbf{x}, \mathbf{y})$ where $\mathbf{x} : V \times O \rightarrow \{0, 1\}$ and $\mathbf{y} : V \times V \times O \rightarrow \{0, 1\}$. The family of all such tuples is denoted as Ψ, which represents the whole space of all caching strategies.*

Definition 2. *A caching strategy for a node v_i is defined as $(\mathbf{x_i}, \mathbf{y_i})$, where $\mathbf{x_i} : \{v_i\} \times O \rightarrow \{0, 1\}$ and $\mathbf{y_i} : \{v_i\} \times V \times O \rightarrow \{0, 1\}$ are the partial functions of $\mathbf{x}$ and $\mathbf{y}$ with domains restricted to $\{v_i\} \times O$ and $\{v_i\} \times V \times O$ respectively.*

Note that $\times$ above represents the Cartesian product. We strive for a caching strategy that is (*i*) Pareto efficient; (*ii*) has well-defined fairness achieved amongst the nodes; and (*iii*) attains high performance. From a utilitarian perspective, this combination of attributes will lead to stable cooperation.

III. FUNDAMENTALS OF BARGAINING GAMES

A bargaining game is a model for analysing how parties collaborate to obtain certain utility values. We model collaborative caching as a bargaining game, in which we aim to achieve both high performance (utility) and fairness. Ideally, a solution is considered fair if it satisfies certain axioms [17], [18]: (*i*) Pareto optimality; (*ii*) Scale invariance; (*iii*) Symmetry; (*iv*) Independence of the irrelevant alternatives; (*v*) Monotonicity. Nash proved that there is one unique solution which satisfies axiom (*i*)-(*iv*), termed the Nash Bargaining Solution (NBS) [17]. The NBS can be extended to multiple players. On the other hand, the Kalai-Smorodinsky Solution (KS) [18] satisfies axiom (*i*)-(*iii*)and (*v*). These two solutions lead to two fairness metrics. Compared to NBS, KS often does not have a closed-form expression. Hence, we focus on NBS.

A. Game Definitions

In game theory, each node attempts to optimise its personal "utility". In caching, utility for a node, U_i, can be measured by the delay to respond to a client. Each cache aims to serve its clients with the lowest possible delay. Consequently, serving a request from the local cache produces the greatest utility, but redirecting a request to another nearby neighbour also increases it (rather than forwarding to the original source). As such, a selfish cache strives to maximise its utility through a combination of local caching and redirects to nearby neighbours.[1] Of course, if utility can be maximised solely through local caching then a node will cease to collaborate. NBS is an axiomatic solution for solving the following problem:

$$\max \prod_{v_i \in V} (U_i - u_i^0) \qquad (1)$$

Eq. (1) is called the Nash product. As mentioned, U_i is node i's utility. u_i^0 is the initial disagreement value of i. The disagreement value is defined as the worst utility payoff a node would accept for collaboration. In practice, a node sets its disagreement value to the maximum value achieved by optimising locally as a standalone cache, *e.g.* using Least Recent Used. In the following, we give the formal definition of our in-network caching game and its solution.

Definition 3. *An in-network caching game is a tuple $(\mathbf{\Omega}, \mathbf{u^0})$, where $\mathbf{\Omega} \subset \mathbb{R}^{|V|}$ contains all the utility values obtainable via collaboration, and $\mathbf{u^0} \subset \mathbb{R}^{|V|}$ contains all the disagreement values leading to a negotiation breakdown.*

Let $\mathbf{\Omega^e} \subset \mathbf{\Omega}$ be the Pareto frontier of set $\mathbf{\Omega}$, *i.e.* the potential Pareto efficient solutions. We assume that $\mathbf{\Omega^e}$ is also a concave function with a closed compact convex domain. A game is considered fair iff its outcome is fair. Therefore, we have:

Definition 4. *A fair caching game is a game $(\mathbf{\Omega}, \mathbf{u^0})$ with a Nash bargaining solution, i.e. a function $f : \mathbf{\Omega^e} \rightarrow \mathbf{\Psi}$ such that $f(\mathbf{\Omega}, \mathbf{u^0}) = (\mathbf{x}, \mathbf{y})$ uniquely maximises $\prod_{v_i \in V}(U_i - u_i^0)$.*

By taking the logarithm of the objective function (1), we have $\ln(\max \prod_{v_i \in V}(U_i - u_i^0)) = \max \ln(\prod_{v_i \in V}(U_i - u_i^0))$. By taking the negation, NBS can be obtained equivalently by:

$$\max \sum_{v_i \in V} \ln(U_i - u_i^0) \implies \min \sum_{v_i \in V} -\ln(U_i - u_i^0). \quad (2)$$

B. Fairness Definitions

We argue that collaboration should follow the intuitive concept of fairness, such that all caches receive fair utility improvements through collaboration. This is critical to ensure that node owners do not feel exploited and do not disengage from the collaboration. Being Pareto efficient, alone, does not achieve this. To attain fairness, it is necessary to formalise the concept. Three well-defined fairness metrics are often referred to in the literature [19]–[21], *i.e. Egalitarian (EF)*, *Max-min (MF)* and *Proportional (PF)* fairness. *EF* pursues an equal amount of improvement on every node, which usually creates Pareto inefficiency (and is thus seldom used in practice). Both *MF* and *PF* have axiomatic foundations and are widely used, *e.g.* in traffic engineering. *MF* is a generalisation of KS, while *PF* is a generalisation of NBS. Thus, we only focus on *PF*:

Definition 5. *Proportional Fairness (PF): A caching strategy $(\mathbf{x}^*, \mathbf{y}^*)$ is PF iff $\forall (\mathbf{x}, \mathbf{y}) \neq (\mathbf{x}^*, \mathbf{y}^*) \Rightarrow \sum_{v_i \in V} \frac{U_i - U_i^*}{U_i^* - u_i^0} < 0$.*

A cache allocation is considered *PF* if the re-allocation of any object would decrease the proportional utility gain (from

---

[1]Here, we assume that each individual node selfishly optimises. However, our model can also support collective self interest amongst multiple nodes, *e.g.* if several caches are owned by a single organisation.

collaboration) of a node by *less* than the respective aggregated increase for others. For example, imagine an object is re-allocated from Cache 1 to Cache 2. It would not be fair if this re-allocation reduces Cache 1's utility by 20%, so that Cache 2 could increase its utility by just 3%. However, it would be considered fair if Cache 2 could increase its utility by 60%. Importantly, to be considered *PF* (and to incentivise uptake), it is necessary for *all* caches to improve their performance over local optimisation (*e.g.* Least Recently Used). Otherwise collaboration would immediately breakdown in favour of local algorithms. If this were to occur, each node would simply select its own preferred algorithm (*e.g.* LRU). However, in our caching games, *PF* is guaranteed by NBS (the proof is trivial and available in [22]). It is also trivial to show whenever a Pareto efficient solution achieves EF, it also achieves MF, *i.e.* a fair solution achieves all three fairness metrics in NBS given it is EF. We later show that being collaborative brings benefits to almost all nodes, indicating that the number of (rational) nodes who would revert to a local algorithm are very limited.

IV. Solving a Fair In-Network Caching Game

We next devise both centralised and decentralised optimal solutions for achieving fair caching. We later use these to evaluate our heuristic solution, *FairCache*. We avoid presenting all the standard mathematical details but rather focus on the key mechanisms (remaining are available in [22]).

A. Defining a Utility Function

In this paper, we assume that a cache's utility is generated from serving its users' demand with low delay. For edge nodes, this demand comes directly from clients, whereas for core nodes this comes from their downstream customers. In either case, utility could be improved by a router using its local cache, or by redirecting a request to a nearby collaborative cache. Both improve delay compared to following the deterministic route to the origin. More precisely, v_i's utility is defined as:

$$U_i = \sum_{o_k \in O} s_k w_{i,k} x_{i,k} + \sum_{o_k \in O} \sum_{v_j \in N_i} \frac{s_k w_{i,k}}{l_{i,j}^* + 1} y_{i,j,k} \quad (3)$$

Both terms show that the utility is a non-decreasing function of demand and content size. The second term shows that the utility of retrieving remote content decreases as the distance increases. It indicates that a node prefers fetching from the closest source to reduce latency and traffic footprint. Although this affine utility function is used throughout the paper, any other metric (*e.g.* bandwidth) or affine function can be used to model the utility without change to our model.

B. Centralised Solution

We begin by outlining the optimal solution, which can be computed centrally (*e.g.* on a controller). Without loss of generality, we assume unit object size $s_k = 1$,[2] also let

[2]In practice, object size could either be varied per-object or, alternatively, objects can be separated into smaller fixed size units

$l_{i,j} \triangleq l_{i,j}^* + 1$ for simplicity of expression. Plugging in Eq. (3) and Eq. (2), we define the optimisation problem as:

$$\min \sum_{v_i \in V} -\ln(\sum_{o_k \in O} w_{i,k} x_{i,k} + \sum_{o_k \in O} \sum_{v_j \in N_i} \frac{w_{i,k}}{l_{i,j}} y_{i,j,k} - u_i^0).$$
$$(4)$$

Subject to

$$\sum_{o_k \in O} x_{i,k} \leq C_i, \quad \forall v_i \in V \quad (5)$$

$$\sum_{v_j \in N_i} y_{i,j,k} \leq 1, \quad \forall v_i \in V, \forall o_k \in O \quad (6)$$

$$y_{i,j,k} \leq x_{j,k}, \quad \forall v_i, v_j \in V, o_k \in O \quad (7)$$

$$x_{i,k} \in \{0,1\}, \quad \forall v_i \in V, o_k \in O \quad (8)$$

$$y_{i,j,k} \in \{0,1\}, \quad \forall v_i, v_j \in V, o_k \in O \quad (9)$$

Constraint (5) means the content stored at a node cannot exceed its cache capacity. Constraint (6) simplifies the data scheduling and avoids requesting redundant content by constraining a node to retrieve a maximum of one complete object in a cache period. Constraint (7) says v_i can retrieve o_k from v_j only if v_j has cached it; it also says v_i cannot get more than v_j can offer. Constraints (8) and (9) impose the domain of decision variables.

The above optimisation problem is a typical *Integer Programming* program which is NP-Complete. By applying *Linear Programming relaxation*, we relax constraints (8) and (9) by letting $x_{i,k} \in [0,1]$ and $y_{i,j,k} \in [0,1]$. We later round up/down $x_{i,k}$ and $y_{i,j,k}$ to construct caching strategies. Such relaxation renders a suboptimal solution hence is considered as the lower bound of the actual performance. Regarding Eq. (4), since all the affine functions are log-concave, their composite with logarithmic functions preserves concavity. Thus, the problem (4) becomes a convex optimisation problem defined over a set of compact and convex constraints as Lemma 1 shows, which leads to the unique Pareto efficient solution, which is trivially followed by the existence of the equilibrium in NBS by definition [17].

Lemma 1. *The problem (4) is a convex optimisation problem.*

The centralised solution can be derived by applying standard convex optimisation techniques (see [22]). The solver needs the demand matrix of each cache, cache size, content object set, whole network topology *etc.* as inputs. The whole equation system has $3|O| \cdot |V|^2 + 2|O| \cdot |V| + |V|$ variables and the same number of equations.

Theorem 1. *In a fair collaborative game, for the optimal caching strategy $(\mathbf{x_i^*}, \mathbf{y_i^*})$ of node v_i, there exist non-negative vectors $\boldsymbol{\alpha} \succeq 0$, $\boldsymbol{\beta} \succeq 0$, $\boldsymbol{\gamma} \succeq 0$, $\boldsymbol{\delta} \succeq 0$ and $\boldsymbol{\lambda} \succeq 0$, such that*

$$x_{i,k}^* = \frac{1}{\alpha_i + \gamma_{i,k} - \sum_{v_j \in N_i^+} \lambda_{j,i,k}} - \frac{\tau_{i,k}}{w_{i,k}} \quad (10)$$

$$y_{i,j,k}^* = \frac{1}{\lambda_{i,j,k} + \beta_{i,k} - \delta_{i,k}} - \frac{l_{i,j} \tau'_{i,k}}{w_{i,j}} \quad (11)$$

where $\tau_{i,k} = U_i - u_i^0 - w_{i,k} x_{i,k}$ and $\tau'_{i,k} = U_i - u_i^0 - \frac{w_{i,k}}{l_{i,j}} y_{i,j,k}$.

The centralised solution can be easily derived by applying standard convex optimisation techniques, as its closed form

expression shown in Theorem 1. The proof and actual equation system can be found in appendix. In practice, we can choose any modern LP solver to calculate the solution, but the computation of solving such a large system is non-trivial.

C. Distributed Solution

The optimal centralised solution has obvious drawbacks in its actual use: (*i*) it suffers from high computation complexity; (*ii*) it creates a single point of failure; and (*iii*) it is not adaptive under network dynamics. Hence, we next translate it into a distributed solution using decomposition techniques.

To solve an equation system, each node can be viewed as a subsystem. If they simply optimise locally, all the calculations in each subsystem are independent from those in others. However, due to collaboration, there are variables and constraints, which are referred to as *complicating variables and constraints* [23]. These make calculations interdependent and couple a subsystem with others. In problem (4), the only complicating constraint is (7).

To decompose Eq. (4), we apply Lagrangian dual relaxation. Lagrangian dual relaxation provides a non-trivial lower-bound of a primal. The difference between the dual and the primal is called the *duality gap*, which can be zero if certain conditions are met as we show below. The Lagrangian $\mathcal{L}(\cdot) : \mathbb{R}^{2|O||V|^2} \to \mathbb{R}$ associated with objective (4) is defined as follows:

$$
\mathcal{L}(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda}) \tag{12}
$$
$$
= \sum_{v_i \in V} [-\ln(U_i - u_i^0) + \sum_{v_j \in N_i} \sum_{o_k \in O} \lambda_{i,j,k}(y_{i,j,k} - x_{j,k})].
$$

$\boldsymbol{\lambda} \succeq 0$ is the dual variable associated with constraint (7) of objective function (4). Then the Lagrangian dual function $d(\cdot) : \mathbb{R}^{|O||V|^2} \to \mathbb{R}$ is as follows:

$$
d(\boldsymbol{\lambda}) = \inf_{\mathbf{x} \in X, \mathbf{y} \in Y} \mathcal{L}(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda}). \tag{13}
$$

Given $\boldsymbol{\lambda}$, let $\mathbf{x}^*$ and $\mathbf{y}^*$ be the unique minimizers for the Lagrangian (12) over all $\mathbf{x}$ and $\mathbf{y}$. Then the dual function (13) can be rewritten as $d(\boldsymbol{\lambda}) = \mathcal{L}(\mathbf{x}^*, \mathbf{y}^*, \boldsymbol{\lambda})$. By maximising the dual function, we can reduce the duality gap. The Lagrangian dual problem of the primal (4) is defined as:

$$
\max_{\boldsymbol{\lambda} \in \mathbb{R}^{|O||V|^2}} d(\boldsymbol{\lambda}) = \mathcal{L}(\mathbf{x}^*, \mathbf{y}^*, \boldsymbol{\lambda}). \tag{14}
$$

The constraints for the dual are the same as those of the primal except constraint (7) which has already been included in the dual objective function (14). Because (4) is convex and all the constraints (5)(6)(8) and (9) are affine, Slater's condition holds given a solution exists, and the duality gap is zero. Thus, when the dual (14) reaches its maximum, the primal also reaches its minimum. The optimal solution for primal problem (4) can be derived from the optimal solution for dual problem (14).

After decomposition, each node v_i now only needs to optimise its utility locally for a given $\boldsymbol{\lambda}$ by calculating:

$$
\min \mathcal{L}_i(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda})
$$
$$
= -\ln(U_i - u_i^0) + \sum_{v_j \in N_i} \sum_{o_k \in O} \lambda_{i,j,k}(y_{i,j,k} - x_{j,k}).
$$

We use the standard projected subgradient method [23] to derive the algorithm. Let $h(\boldsymbol{\lambda})$ and $\partial d(\boldsymbol{\lambda})$ denote the subgradient and subdifferential of dual function $d(\cdot)$ at point $\boldsymbol{\lambda}$ respectively. Then for every $h_{i,j,k} \in h(\boldsymbol{\lambda})$ we have:

$$
h_{i,j,k} = y_{i,j,k}^* - x_{j,k}^* \implies h(\boldsymbol{\lambda}) \in \partial d(\boldsymbol{\lambda}).
$$

Gradient $\mathbf{h} \triangleq h(\boldsymbol{\lambda})$ points to the direction where $d(\cdot)$ increases fastest. In each iteration, node v_i solves the local subsystem (15) to update the dual variable $\boldsymbol{\lambda}$. t represents the t^{th} iteration. ξ_t is the step-size in the t^{th} iteration which can be determined by several standard methods [23]. The projected subgradient method projects $\boldsymbol{\lambda}$ on its constraint (*i.e.* $\boldsymbol{\lambda} \succeq 0$) in each iteration, and we use $(\cdot)_+$ as a shorthand for the Euclidean projection. Eventually $\boldsymbol{\lambda}^{(t)} \to \boldsymbol{\lambda}^*$ when $t \to \infty$. The primal solution can be constructed from the optimum $\boldsymbol{\lambda}^*$. Combining the above, we refer to Eq. (15) as the *distributed optimal algorithm*:

$$
\begin{cases}
\mathbf{x}_i^{(t)}, \mathbf{y}_i^{(t)} = \arg\min_{\mathbf{x}, \mathbf{y}} \mathcal{L}_i(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda}^{(t)}) \\
\mathbf{h}^{(t)} = -(\mathbf{x}_i^{(t)} - \mathbf{y}_i^{(t)}) \\
\boldsymbol{\lambda}^{(t+1)} = (\boldsymbol{\lambda}^{(t)} + \xi_t \sum_{v_j \in N_i \cup \{v_i\}} \mathbf{h}_j)_+
\end{cases} \tag{15}
$$

Theorem 2. *Optimal algorithm converges to its optimum as the sequence $\{\boldsymbol{\lambda}^{(1)}, \boldsymbol{\lambda}^{(2)} ... \boldsymbol{\lambda}^{(t)}\}$ converges, if a diminishing step size is used such that $\lim_{t \to \infty} \xi_t = 0$ and $\sum_{t=1}^{\infty} \xi_t = \infty$.*

The above theorem guarantees convergence [22]. $\lambda_{i,j,k}$ can be viewed as the *"shadow price"* of transferring o_k from v_j to v_i, which is a "cost" for v_i but an "income" for v_j. Given an ideal network, the distributed optimal algorithm will converge faster than the centralised one due to its parallel computations. It is worth emphasising that although the optimal algorithm above distributes the calculations over nodes, the overall computations are not reduced. At the same time, the communication cost increases due to exchanging "shadow price" information. However, the overall communication complexity remains the same as that of the centralised solution as we will later show in Section VI.

V. FAIRCACHE: A LOW-COMPLEXITY HEURISTIC DESIGN

The distributed algorithm, although optimal, comes with high overheads. To mitigate this, we propose FairCache, a heuristic algorithm which does not require global knowledge regarding the content and network topology. We emphasise that FairCache is a decentralised algorithm for fairly sharing cache capacity across multiple *trusted* stakeholders. It is not intended to be a protocol, by which malicious behaviour (*e.g.* falsifying content demand) can be prevented. Hence, we assume trusted parties who faithfully execute the algorithm, much like is assumed within existing Internet routing schemes.

A. Overview of Heuristics

To understand the rationality behind our heuristic, we first give a verbal explanation on the mechanisms of the optimal algorithm expressed in Eq. (15). Recall $\boldsymbol{\lambda}$ represents the shadow price of transmitting an object between two nodes. Each node hence maintains a list of prices for any given object from any given node. In each iteration of the optimisation,

a node tries to minimise its total cost using $\boldsymbol{\lambda}^{(t)}$. During the optimisation, the node adjusts its local caching strategy (via $\mathbf{x_i}$ and $\mathbf{y_i}$) and price list on other nodes (via $\boldsymbol{\lambda}$ and $\mathbf{h}$). Namely, a node may decide to cache an object if it brings significant improvement, or stop retrieving an object from another node due to high cost. Meanwhile the node adjusts the price on how to charge its neighbours by offering help. Then the node collects the price adjustments from *all others in the network* and updates its own list. The procedure continues until the performance converges based on certain well-defined criteria (as described next). Future updates are periodically shared to address changes in content popularity. As, generally, popularity changes are relatively slow to occur (hours, rather than minutes), this does not create considerable overheads. The mechanisms above indicate that we can approximate the optimal algorithm in the following ways:

(*i*) **Cut out unpopular content**: This approximation takes advantage of the highly skewed content popularity distribution. It is well-known that the popularity distribution has a long and heavy tail and most content fall into the tail. Removing the tail can significantly reduce the size of the exchanged messages. Meanwhile, the results will not be significantly influenced because of their marginal contribution to the overall utility (whilst also reducing signalling overheads dramatically). Thus, requests for unpopular content will be forwarded towards the origin (as with vanilla NDN [1]).

(*ii*) **Cut out distant nodes**: This approximation takes advantage of topological locality. Since the utility of retrieving distant content is a decreasing function of the hop count between two nodes, the value quickly diminishes as path length increases. It is more likely to find the requested content in nearby nodes due to content spatial locality [24], [25]; removing remote nodes should not have significant impact on the result.

(*iii*) **Reverse direction**: This approximation takes advantage of the behaviours of gradient methods. In the optimal algorithm, the neighbourhood (r) gradually shrinks from the network diameter to its optimum (as a result of minimising the cost function). However, most elements in $\mathbf{y_i}$ are already set to zero by the gradient method in the beginning phase of the optimisation. Exchanging messages between nodes that are not going to collaborate is a waste of resources. By growing the neighbourhood set outwardly, instead of shrinking it, we can avoid unnecessary message exchange.

B. FairCache Algorithm

We embed the above heuristics in our algorithm, FairCache, presented in Algorithm 1. It takes several inputs. $\mathbf{w}$ is the local demand matrix. r is for tracking the current number of hops that defines a node's neighbourhood radius. θ' is used for recording the utility improvement by increasing the radius from r to $r+1$, while θ is the threshold below which FairCache should stop growing the neighbourhood size. $\boldsymbol{\lambda}$ is the list for tracking the shadow prices; this needs to be exchanged amongst nodes (via price adjustment $\mathbf{h}$) in a neighbourhood. Algorithm 2 shows how heuristic (*i*) is implemented to derive a reduced content set.

Algorithm 1 Fair in-network caching (FairCache) on v_i

1: **Input:**
2:     Demand matrix $\mathbf{w}$
3:     Dual variables $\boldsymbol{\lambda}$
4:     Search radius $r = 0$
5:     Improvement threshold θ, θ' ($\theta = 10^{-2}; \theta < \theta'$)
6: **Output:**
7:     Caching decision $\mathbf{x}_i$
8:     Collaboration decision $\mathbf{y}_i$
9: **while** $\theta' \geq \theta$ **and** $r <$ network diameter **do**
10:     $r = r + 1; t = 0;$
11:     **while** $t < t_{stop}$ **do**
12:         $\mathbf{x}_i, \mathbf{y}_i = \arg\min_{\mathbf{x},\mathbf{y}} \mathcal{L}_i(\mathbf{x}, \mathbf{y}, \boldsymbol{\lambda})$
13:         $\mathbf{h} = \mathbf{y}_i - \mathbf{x}_i$; trim $\mathbf{h}$ for $\forall v_j \in N_i^+$
14:         $\mathbf{h} = \mathbf{h} + \sum_{\forall v_j \in N_i} \mathbf{h}_j$
15:         $\boldsymbol{\lambda} = (\boldsymbol{\lambda} + \xi\mathbf{h})_+$
16:         $t = t + 1$
17:     **end while**
18:     Update θ' with current improvement
19: **end while**

Algorithm 2 Construct reduced content set O_i' on v_i

1: **Input:**
2:     Demand matrix $\mathbf{w}$
3:     Complete content set O
4: **Output:**
5:     Reduced content set O_i'
6: Sort O in decreasing order based on $\mathbf{w}$;
7: set $O_i' = \emptyset$;
8: **for each** $o \in O$ **do**
9:     **if** size (O_i') + size (o) $\leq v_i$'s cache capacity
10:     **then** add o to O_i'
11:     **else** break
12: **end for each**

Algorithm 3 On sending the price update $\mathbf{h}_i$ on v_i

1: **Input:**
2:     Neighbourhood N_i^+
3:     Caching strategy $\mathbf{x}_i, \mathbf{y}_i$
4: set $\mathbf{h}_i = \mathbf{y}_i - \mathbf{x}_i$;
5: **for each** $v_j \in N_i^+$ **do**
6:     construct $\mathbf{h}_i'$ from $\mathbf{h}_i$ based on $(O_i' \cup O_j')$;
7:     send $\mathbf{h}_i'$ to v_j;
8: **end for each**

Algorithm 4 On receiving all the $\mathbf{h}_j'$ on v_i

1: **Input:**
2:     Neighbourhood N_i
3:     Caching strategy $\mathbf{x}_i, \mathbf{y}_i$
4: set $\mathbf{h}_i = \mathbf{y}_i - \mathbf{x}_i$;
5: **for each** $v_j \in N_i$ **do**
6:     receive $\mathbf{h}_j'$ from v_j;
7:     set $\mathbf{h}_i = \mathbf{h}_i + \mathbf{h}_j'$;
8: **end for each**

To apply the approximations, for node v_i, instead of making a complete price list, $\boldsymbol{\lambda}$, containing all the content and nodes in the network, node v_i makes a partial $\boldsymbol{\lambda}$ which only includes: (*i*) the most popular content that can be fit into its local cache (*i.e.* heuristic (*i*)); and (*ii*) the nodes in the neighbourhood defined by r (*i.e.* heuristic (*ii*)). It is possible that v_i observes other content in the $\mathbf{h}_j$, collected from neighbours while r grows (*i.e.* heuristic (*iii*)), then v_i dynamically adds those content into its own $\boldsymbol{\lambda}$. v_i can also remove items from $\boldsymbol{\lambda}$ if they are too expensive to retrieve. After local optimisation in each iteration, the price adjustment $\mathbf{h}$ will be trimmed before exchange by removing information that is not included in $\boldsymbol{\lambda}$; and removing the unchanged items, *i.e.* the zero values. Essentially, v_i only exchanges the trimmed $\mathbf{h}$ within its neighbourhood and $\boldsymbol{\lambda}$ only contains the aggregated popular content in the neighbourhood. Algorithm 3 and 4 detail the logic in lines 13–14 in Algorithm 1 whenever sending and receiving price updates in each iteration. Obviously, these approximations render incomplete information (due to removing unpopular content and distant nodes). To handle the missing $\lambda_{i,j,k}$ in the local optimisation, we let missing $\lambda_{i,j,k} = \infty$ ($i \neq j$), which indicates that the optimisation algorithm should neither exchange unpopular content nor exchange content with distant nodes.

Looking more closely, FairCache consists of two loops. The outer loop (lines 9–19) increases the search radius r by one hop in each iteration. The outer loop stops when the current improvement, θ', drops below the threshold θ (*i.e.* $\theta' < \theta$) due to enlarging the neighbourhood. The inner loop (lines 11–17) finishes the local optimisation —- this is the calculation in Eq. (15) for the given neighbourhood defined by the current radius r. The communication overhead come from the operation in line 14 which collects the price adjustments $\mathbf{h}_j$ from the neighbourhood N_i. Line 15 adjusts the local shadow price list and updates the $\boldsymbol{\lambda}$ by removing or adding items.

VI. FairCache Complexity Analysis

A. Overview

The distributed solution requires exchange of control messages between nodes, whilst the centralised solution requires the distribution of the derived solution to *all* nodes (once it has been centrally computed). The optimal solution comprises of $|O|$ matrices of size $|V|^2$. The aggregation of each row i (in all matrices) represents one specific optimal strategy for the corresponding node v_i in the system. To transmit these rows to the nodes, the aggregated communication complexity is $\Theta(|O| \cdot |V|^2)$. This is the same for the distributed solution, as this performs the same computation (shared across multiple nodes). The rest of this section will delineate the complexity of FairCache. Table II presents an overview of the computational, communication, and space complexity of all three solutions.

B. FairCache Heuristic Complexity

FairCache's computation, communication, and space complexities all correlate with the size of the content set and the neighbourhood set. Herein we evaluate the complexity savings of FairCache's three heuristics.

Heuristic (*i*) is used to reduce the content set size ($|O'|$ from $|O|$). The level of reduction depends on both the popularity distribution of content and the percentage of requests we want to capture in the request streams. It is well documented that content popularity follows a Zipf-like distribution [26], [27]. If we assume a Zipf distribution with $\alpha = 1.0$, for a content set of 10^6 equal sized objects, we are able to cover 72.8% of the requests by caching only 2% of $|O|$. In other words, $|O'| = 0.02 \cdot |O|$ indicates a 98% reduction. Even if we only cache 1000 objects (*i.e.* 0.1% of $|O|$), we can still capture 52% of requests, whilst leading to a significant reduction in $|O'|$, *i.e.* up to 99.9%.

Heuristic (*ii*) is used to reduce the size of collaboration neighbourhood, which is $|V|$ for the distributed solution. By investigating the FairCache algorithm, we can see that the communication complexity is due to exchanging $\mathbf{h}$ in order to update the local shadow price $\boldsymbol{\lambda}$, *cf.* the last equation in Eq.(15) and its corresponding line 14 in Algorithm 1. The overhead therefore consist of two parts. The first part is induced by replying the queries on $\mathbf{h}$ from the nodes having v_i in their neighbourhood, namely N_i^+. The second part is induced by collecting $\mathbf{h}$ from the nodes in v_i's own neighbourhood, namely N_i. Given that the communication complexity is measured by the number of exchanged messages, the complexity ϕ_i of node v_i can be calculated as

$$\phi_i = c \cdot |O'| \cdot (|N_i^+| + |N_i|) \tag{16}$$

Scalar c in Eq. (16) represents a constant factor for communication complexity, and can be understood as message size or other protocol-dependent factors. Now we show how to calculate the overall complexity Φ at system level from ϕ_i in the following.

Given any $v_j \in N_i$, a neighbourhood relation can be written as a tuple (v_i, v_j). Calculating $\sum_{v_i \in V} |N_i|$ is equivalent to counting how many tuples there are in the whole system. Obviously, $v_i \in N_j \iff v_j \in N_i^+, \forall v_i, v_j \in V$, in other words, as long as there is a tuple (v_i, v_j) for N_i, there must be a tuple (v_j, v_i) for some N_j^+, and vice versa. Hence we can show $\sum_{v_i \in V} |N_i| = \sum_{v_i \in V} |N_i^+|$. Because one tuple (v_i, v_j) represents one message exchange from v_i to v_j, by using double counting technique, it is obvious that each message will be counted twice in the calculation if we try to aggregate all ϕ_i: *i.e.* (v_i, v_j) is in both v_i's N_i and v_j's N_j^+. Therefore, overall complexity Φ can be calculated as half of the aggregated ϕ_i of all v_i in V as below.

$$\Phi = \frac{1}{2} \sum_{v_i \in V} \phi_i = \frac{c}{2} \cdot |O'| \cdot (\sum_{v_i \in V} |N_i^+| + \sum_{v_i \in V} |N_i|) \tag{17}$$

$$= c \cdot |O'| \cdot \sum_{v_i \in V} |N_i| \tag{18}$$

The intuitive explanation of eliminating N_i^+ in calculating Φ is that we only need to count the aggregated messages sent out from a node (*i.e.* the size of a node's own neighbourhood) since for each message there is always a correspondence in the network to receive it. Clearly, if we denote the average neighbourhood size as $|\overline{N}|$, then the system level communi-

cation complexity of FairCache is $\Phi = \Theta(|O'| \cdot |V| \cdot |\overline{N}|)$ as below.

$$\Phi = c \cdot |O'| \cdot \sum_{v_i \in V} |N_i| = c \cdot |O'| \cdot \sum_{v_i \in V} |\overline{N}| \quad (19)$$

$$= \Theta(|O'| \cdot |V| \cdot |\overline{N}|) \quad (20)$$

Similarly in the distributed solution Eq.(15), since each node has the whole network of size $|V|$ as its neighbourhood and uses the complete content set O in the optimisation. Then its complexity can be calculated as $\Phi = \Theta(|O| \cdot |V|^2)$ by plugging the term $|\overline{N}| = |V|$ and $|O'| = |O|$ into Eq.(20). As we can see, the derived complexity is exactly the same as that of the centralised solution.

Heuristic (*iii*) is used to minimise the size of the neighbourhood by increasing, rather than decreasing, the neighbourhood size. As we can see from the analysis above, the complexity of FairCache is proportional to its average neighbourhood size $|\overline{N}|$. Reducing its size is therefore beneficial. The authors in [24] have proposed a neighbourhood model to calculate $|\overline{N}|$ in its closed form on any general network topologies. With its most general form, the overall complexity of the network of average r-hop neighbourhood can be calculated as

$$\Phi = |O| \cdot |V| \cdot n_1 \cdot \left(1 + \left[\frac{n_2}{n_1}\right] + \left[\frac{n_2}{n_1}\right]^2 + \cdots + \left[\frac{n_2}{n_1}\right]^{r-1}\right)$$

$$= \Theta\left(|O| \cdot |V| \cdot \left[\frac{n_2}{n_1}\right]^r\right) \qquad \forall n_1, n_2 \in \mathbb{N}, n_2 > n_1$$

where n_1 and n_2 denote the average number of one-hop and two-hop neighbourhoods, and r represents the search radius. If we assume $n_2 > n_1$ (which holds for all scale-free networks), the complexity of FairCache will grow exponentially if the search radius r increases.[3] Hence, by ensuring a small r, Heuristic (*iii*) dramatically improves scalability. Importantly, it has been shown that r is very small in most network optimisation problems [24] (confirmed in Section VII).

Besides the memory for storing the actual content objects, extra space is needed to store the input parameters of the optimisation problem (*i.e.* **w** and λ). The space complexity specifically refers to such extra space. Obviously, the space complexity is also decided by the size of content set and the size of neighbourhood. Hence FairCache on each node needs $\Theta(|O'||\overline{N}|)$ memory space to keep track of the information needed for optimisation. If we assume the average object size is 800 KB (much smaller than the average YouTube video size: 8 MB [26]) and each node has 8 GB memory installed. Each node needs to maintain roughly 10^4 objects. If we allocate 64-bit space to store the information for each object, given a neighbourhood of 10 nodes, even in the worst case wherein all these content sets are disjoint, each node only needs approximately 800 KB extra memory. However, spatial locality is very common [24] in content networks, which means there is a significant overlap regarding the stored content among nearby routers. If we assume that only 10% of objects are overlapping, the 800 KB can be further reduced to 651 KB (by summing up a geometric sequence to calculate). In

[3]Note that most natural graphs like Internet, ISP networks, and social networks are all scale-free [28]

Complexity comparison of three different solutions

Type of complexity	Centralised	Distributed	FairCache
Computation	$\Theta(|O||V|^2)$	$\Theta(|O||V|^2)$	$\Theta(|O'||V||\overline{N}|)$
Communication	$\Theta(|O||V|^2)$	$\Theta(|O||V|^2)$	$\Theta(|O'||V||\overline{N}|)$
Memory space	$\Theta(|O||V|^2)$	$\Theta(|O||V|^2)$	$\Theta(|O'||V||\overline{N}|)$
Individual (worst)	$\Theta(|O||V|^2)$	$\Theta(|O||V|)$	$\Theta(|O'||\overline{N}|)$

TABLE II: The overall and (worst) individual complexity of the three solutions. The communication complexity of the centralised solution does not occur during the computation but comes afterwards when distributing the solution to all nodes. The distributed solution does not reduce any complexity of the centralised solution but, instead, shares it across multiple nodes. FairCache significantly reduces all three complexities by applying its heuristics.

the end, the extra memory overhead (*i.e.* the ratio between the memory space for parameters and memory space for objects) is less than 0.01%.

C. Complexity Summary

Table II summarises our previous analysis of the complexity of the three different solutions. We wish to highlight the following key points:

- The computation and space complexity of the centralised and distributed solutions are identical. It is always $\Phi = \Theta(|O| \cdot |V|^2)$. However, because the distributed solution is able to share the complexity across each node in the network, in practice, it leads to a more scalable design and avoids the single point of failure.
- FairCache is able to reduce all complexities by limiting the sizes of the content set and its collaborative neighbourhood. The complexity of FairCache is only a small fraction of the distributed solution, more precisely $\frac{|O'|}{|O|} \cdot \frac{|\overline{N}|}{|V|}$. We have shown $|O| \gg |O'|$, also in reality $|V| \gg |\overline{N}|$ (as showed in [24]). We also confirm this experimentally in Section VII), hence the improvement is significant and improvements should grow as the network size increases.
- FairCache on each node needs $\Theta(|O'||\overline{N}|)$ memory space to keep track of the information needed for optimisation (*i.e.* **w** and λ). In practice this introduces only negligible overhead. As the average object size is bigger, the overhead becomes even smaller.

VII. FAIRCACHE EVALUATION

A. Methodology

To evaluate FairCache, we perform extensive simulations using the publicly available LiteLab platform [29]. We use several topologies. First, we use real topologies collected by the Rocketfuel project [28]; namely, two ISP router-level topologies: Sprint (604 nodes, 2,279 edges) and AT&T (631 nodes, 2,078 edges). This embodies the usecase where an individual network has allowed the deployment of caches in their network by multiple parties (*e.g.* via network function virtualisation). Second, we use traces from Guifi [12]. Guifi is the largest open wireless community mesh network in the

world. It allows any user to purchase equipment and become part of the network. We use its core network topology in the Catalunya region (735 nodes, 1,059 edges). Third, we use the CAIDA ITDK trace [30] which takes a snapshots of the Internet AS topology every 24 hours. We use the 7/01/2017 trace (25,107 nodes, 49,458 edges). This captures the situation where each AS operates its own caches, but chooses to collaborate with neighbouring ASes. Fourth, to allow us to vary key graph parameters, we also generate synthetic networks based on two models: the Barabási-Albert (BA) model and the Erdős-Rényi (ER) model. Four parameter sets are used for these synthetic networks: $\{BA_1 : m = 2\}$, $\{BA_2 : m = 4\}$, $\{ER_1 : p = 1.1 \cdot \log(n)/n\}$ and $\{ER_2 : p = 1.5 \cdot \log(n)/n\}$. If there are multiple components in a synthetic network, we only use the largest one. $\log(n)/n$ in the setting guarantees there is only one single giant component in the network with high probability [31]. We have also analysed many recent AS-level topologies using the CAIDA ITDK traces [30]. Since most AS networks have an average degree of 4, we use BA_2 configuration to reflect this fact. Therefore, we note that our synthetic networks represent both a per-router and per-AS topologies. For each topology we attach a single client to each edge router (*i.e.* with degree of 1). For example, this results in 161 clients in Sprint; 207 clients in AT&T; and 200 in Guifi. We then randomly select between 10 and 20 distinct routers to attach a source to. Each router is then allocated a given cache capacity, which we vary; the default is 4 GB, which is $< 0.1\%$ of the corpus. We select a low value to be representative of feasible cache capacity in a wide area network with a large corpus.

Using the above topologies, clients generate requests at each simulation tick, which are then routed through the network to either a content source or an intermediate cache based on the strategy employed. We base our content set on the Youtube trace from [26]. This contains 1,687,506 objects (average size is 8.0 MB and aggregated size is 12.87 TB). We use the view count information to fit a *Zipf* (α) distribution $(\alpha = 0.9537)$ to model the overall content popularity. To explore the impact of different request patterns, we also perform sensitivity analysis on α. Throughout this section, we use our distributed optimal algorithm (*i.e.* Eq. (15)) as an optimal benchmark to compare FairCache against. Each result is averaged over 50 runs; errorbars are not plotted if they are sufficiently small ($< 5\%$).

B. Scalability

We start by exploring scalability, measured by FairCache's convergence rate, *i.e.* how many rounds of message exchange it takes the algorithm to bootstrap. This happens once at initiation: future dynamics are addressed using periodic low cost updates that are algorithmically trivial. Figure 2a compares the convergence rates of the optimal (the upper figure) and FairCache (the lower figure) on the various topologies. The optimal needs significantly longer time to bootstrap than FairCache. Unsurprisingly, larger cache sizes also lead to a longer convergence time, as more state must be exchanged. To investigate how network size impacts the convergence rate, we use synthetic topologies with 4 GB caches. The lines in

Figure 2b are clearly divided into two groups: the upper one is the optimal (with hollow markers) and the lower one is Fair-Cache (with filled markers). The convergence rate degrades as the network size grows. Importantly, though, the increase in convergence time only grows sub-linearly, stabilising at networks of size 1k; we experimented with topologies of up to 9k nodes to find consistent results.

We also measure FairCache's scalability by its traffic overheads. Clearly, it is undesirable to generate large amounts of control messages to bootstrap. In this experiment, we measure the aggregated size of control messages for both the distributed optimal and FairCache as C_O and C_F respectively. We then calculate the traffic reduction as $\frac{C_O - C_F}{C_O}$. Figure 2c presents a box plot of the results for both BA_1 (upper boxes) and ER_1 (lower boxes) topologies. It shows that FairCache is able to achieve over 80% traffic reductions, even on small networks of 100 nodes. As the network size increases, the benefit of using FairCache becomes more obvious. In a network of 900 nodes, FairCache attains 95% reductions. This equates to significant traffic volumes; in one iteration, a 500 node network with 10^3 objects can save 887 MB of control traffic via FairCache (leaving only 66.8 MB). With FairCache, on average, each cache only introduces 136 KB traffic overhead in an iteration. We can therefore combine the above message overhead and convergence measurements to calculate the convergence time. If we configure the rate of control messages to 100 KB/s, FairCache takes 11 minutes to bootstrap. This is just 4.6% of the time taken by the distributed optimal algorithm. Given a saturated 54 Mbit link, the FairCache control messages would therefore consume just 1.4% of bandwidth. Importantly, this is only a bootstrap process; changes in request patterns are addressed with low cost updates within each node's neighbourhood. Even in highly dynamic situations where demands are volatile, we can let nodes exchange demand information in the background while running FairCache. Once the new demand matrix is constructed, FairCache simply re-calculates the new solution. Recall our setting in Section VI where each node needs to maintain 651 KB demand information. To guarantee FairCache continuously runs, a total of 651×10 KB needs to be exchanged among 10 nodes within 11 minutes, these updates constitute under 10 KB/s. These values can be configured to reflect the operating environment (we anticipate that in many cases operators would transfer state at much higher rates). Overall, we believe these overheads are more than acceptable for the overall performance gains (§VII-E)

C. Accuracy

FairCache significantly reduces the convergence time and messaging overhead of fairly allocating caching responsibilities. These improvements potentially come at the cost of accuracy (*i.e.* lower utility than optimal solution). We next inspect the accuracy sacrifice required to obtain these improvements.

To measure the accuracy of FairCache, we compare it against the optimal algorithm using multiple topologies of different sizes. We first run the optimal algorithm and measure the utility U_i for every node i. Similarly, we run FairCache and measure the utility U_i'. We then calculate the accuracy of

(a) Convergence on ISP networks.   (b) Convergence on synthetic networks.   (c) Reduction of traffic overhead.   (d) The accuracy of FairCache.
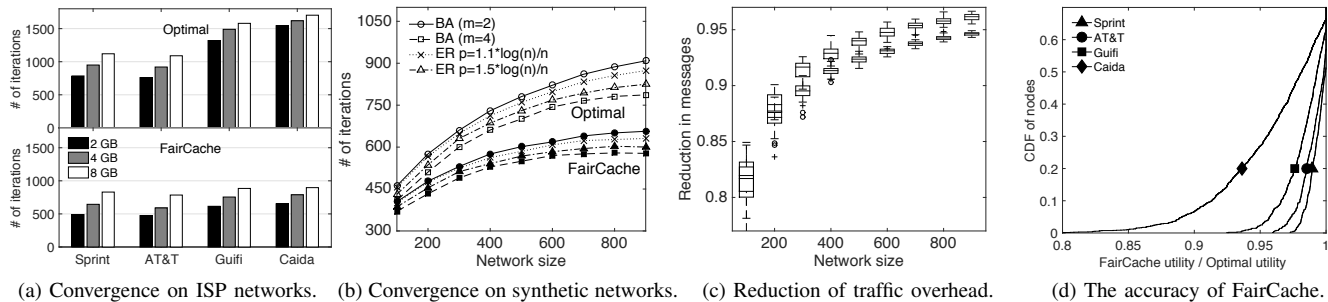
Fig. 2: Compared to the optimal algorithm, FairCache is more scalable on both real and synthetic networks. FairCache has a faster convergence rate and generates less traffic overhead than the optimal. Meanwhile FairCache achieves high accuracy.

FairCache as its ratio to the optimal for every node, *i.e.* $\frac{U'_i}{U_i}$. Figure 2d plots the per node CDF of this ratio. We can see that FairCache achieves very high accuracy. For large networks like Guifi, all the nodes achieve an accuracy of over 92%. For medium size networks like Sprint, all the nodes have at least 97% accuracy and about 50% of the nodes reach 100% accuracy. Besides Figure 2d, we also measure the *aggregated* accuracy using $\frac{\sum_i U'_i}{\sum_i U_i}$. We find it is always above 95% for medium-sized networks, whilst it decreases slightly for larger networks (*i.e.* 3% drop from Sprint to Guifi). Even for the very large CAIDA topology (over 25k nodes), the average accuracy is still $\approx 95\%$ with fewer than 6% nodes that have an accuracy drop of $10\% \sim 20\%$. These findings are consistent across the other topologies. To validate that these benefits continue to be enjoyed by large topologies, we also repeated the experiments presented in Figure 2b on topologies ranging from 1k–9k nodes. Again, we find high levels of accuracy, stabilising at 95%.

The results confirm the rationale behind the heuristics used by FairCache. The approximation introduces almost negligible degradation in the accuracy. The main reason is that the highly skewed content popularity means that the bulk of caching decisions are limited to the most popular objects. This means that FairCache can attain high accuracy without requiring to share information about all objects (unlike the optimal). Further, by localising interactions to neighbouring nodes, FairCache can scale-up easily, without be overly affected by increasing network sizes.

D. Price of Fairness

FairCache aims to realise fair collaboration amongst nodes, which could cause a degradation in aggregated global utility. We use the *Price of Fairness* (PoF) to measure the loss in utility. The PoF is calculated as the ratio between the aggregated utility loss of all nodes using FairCache and the global optimal that does not consider fairness [16]. A higher PoF value indicates a larger utility sacrifice.

Figure 3a and 3b plot the PoF results of using both real and synthetic networks with three cache sizes. Both figures convey the same information, which is that the PoF increases as network size increases. We experiment with both realistic and synthetic networks of up to 9k nodes (Table III summarises the

Scalability test on PoF and accuracy using real ISP topologies.

ASN	#nodes	#edges	Accuracy	PoF (2GB)	PoF (4GB)
#1755	295	544	97.3%	16.5%	11.3%
#3356	1620	6743	96.2%	20.8%	13.7%
#1221	2669	3181	95.4%	22.6%	14.9%
#2914	4670	7618	94.8%	23.9%	15.4%
#1239	7337	9924	94.7%	24.3%	15.6%
#7018	9430	11682	94.8%	24.1%	15.2%
Caida	25107	49458	94.6%	24.4%	15.4%

TABLE III: Key metrics for increasing network sizes (295 - 25k nodes). Networks are real ISP topologies taken from Rocketfuel and CAIDA AS-level trace.

results), to find that the PoF stabilises after reaching a size of $\approx$3k nodes with a maximum PoF of 24% (and a maximum of 20% in the real topologies, *e.g.* Guifi). This is not negligible, but is likely not significant enough to dissuade caches that are interested in fairness from using FairCache. Interestingly, our results also show that increasing the cache size is an effective way to ameliorate the loss in efficiency. In Figure 3a, a 4 GB cache significantly improves the PoF. Using a 2 GB cache, the PoF increases by 11% when the network size increases from 100 to 900, whereas the PoF only increases by 3% if a 4 GB cache is used. Figure 3b shows similar properties with, for example, a 57% improvement in PoF when increasing the cache size from 2 GB cache to 4 GB in Guifi.

Overall, we believe that an average PoF of $< 8\%$ is a cost worth paying for those concerned by a need for fairness. Moreover, FairCache exhibits good scalability regarding both accuracy and PoF, as the results in Table III show. The accuracy only slightly degrades (1.5%) from 1000 nodes to 9000 nodes, and is always above 94.7%. The accuracy stabilises after 2000 nodes. Similarly, PoF stabilises after reaching around 3000 nodes. For the 2GB cache configuration, PoF is capped by 24.3%; 4GB by 15.6%; 8GB by 11.4% (not included in Table III due to space limit). The scalability of FairCache can be explained as follows: only small neighbourhoods play an important role in deciding a node's overall performance. As the size of the neighbourhood is relatively unaffected by the network size, this property always exists.

E. Caching Performance

The previous section has shown that utility is reduced by considering fairness. Next, we explore performance from the

(a) PoF on synthetic networks.   (b) PoF on ISP networks.   (c) Byte hit rate, 4GB cache.   (d) Footprint reduction, 4GB cache.
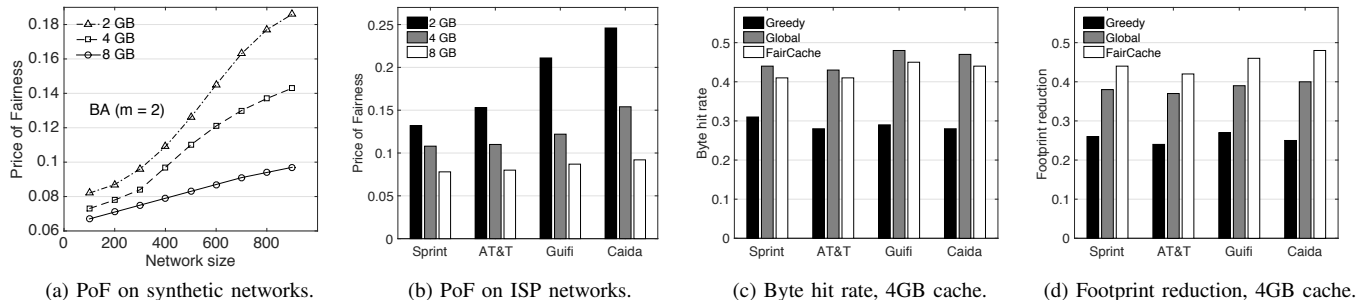
Fig. 3: FairCache achieves fairness by trading off some efficiency. However, a large cache size can effectively reduce PoF. In reality, FairCache is able to achieve very similar performance as Global, and is superior to Greedy in all cases.

perspective of traditional metrics: byte hit rate and footprint reduction. Hit rate is a conventional metric to measure saving on inter-domain traffic, whilst footprint reduction is the reduction on the product of traffic volume and distance.

We compare FairCache against two other strategies: (*i*) *Greedy*, which computes the local optimal for each cache without collaboration; and (*ii*) *Global*, which maximises the aggregated utility. Figure 3c and 3d plot the results on the real networks using 4 GB caches. Naturally, Figure 3c shows that Global achieves the best hit rates due the fact that it optimises the overall network. That said, FairCache only performs slightly worse, with a 5%–10% performance degradation. Compared to Greedy, FairCache is consistently superior with at least a 28% improvement. This shows that, regardless of fairness, FairCache can offer significant performance improvements over local algorithms (note that Greedy is the theoretical upper bound of algorithms such as Least Recently Used). When inspecting the traffic footprint reduction, performance is even higher. FairCache is superior in all networks. Although the reasons are intuitive for Greedy, which sees nodes locally optimising, it is more surprising in Global. The reason is that FairCache only requests from nearby caches (limited by r). In contrast, Global uses *any* node in the network. This increases hit rates, but results in more traffic.

To have a closer look how utility is spread across caches, we select the AT&T network and study the utility distribution in the network (*i.e.* how are traffic savings distributed across caches). Figure 4a plots the CDF of normalised utility values across each node (normalised by the top value per simulation). By comparing Greedy and FairCache, we see that *every* node is better off through collaboration using FairCache (note this is also the case across all other topologies and cache sizes). On the other hand, the Global strategy intersects with both Greedy and FairCache, *i.e.* some caches in Global get lower utility than Greedy. The area between the lines indicates the percentage of caches that are worse off due to global optimisation. The Global strategy leads to 13% of nodes getting worse off compared to Greedy, and 20% compared to FairCache. With Global, these nodes should rationally cease to cooperate. Again, regarding the aggregated utility, Global is only about 5% better than FairCache. Moreover, the CDF curve of Global is more stretched than that of FairCache, which indicates there

are much larger variations in nodes' utilities when using the Global strategy, *i.e.* benefits are not evenly distributed.
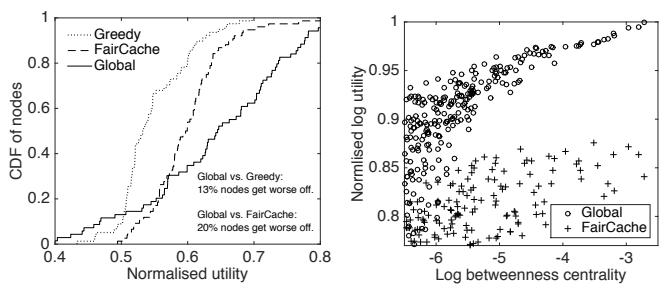
(a) Cumulative distribution of utilities.   (b) Betw. centrality vs. utility.

Fig. 4: Comparison of strategies on AT&T, 4 GB cache size.

Figure 4b shows the log-log plot of nodes' normalised utility as a function of betweenness centrality [7], [14]. Nodes with a high betweenness are core routers, whilst those with low betweenness are usually found at the edge. Interestingly, when nodes use the Global strategy, a node's utility strongly correlates with its position in the network: core nodes gain the highest utility. This is because the Global optimal tends to place all the popular (*i.e.* high value) content at the core to reduce duplicates — a theoretically attractive, but practically infeasible approach. In contrast, FairCache significantly weakens this correlation. This is beneficial as it means that utility is also increased at the edge caches. As well as improving fairness, it also reduces load in the backbone and provides consumers with lower delay access to object. This also contributes to FairCache's high traffic reductions, as hits are pushed closer to clients.

F. Sensitivity Analysis of Spatial/Content Locality

FairCache's heuristics take advantage of highly skewed spatial and content popularity localities. A natural question is how these localities impact the algorithm. To explore this, we perform sensitivity analysis across these two parameters to measure the robustness of our heuristics. Here, we solely present the Guifi topology due to space constraints. The reason we select Guifi is that the dataset contains geographic coordinates of each node, allowing much more fine grained
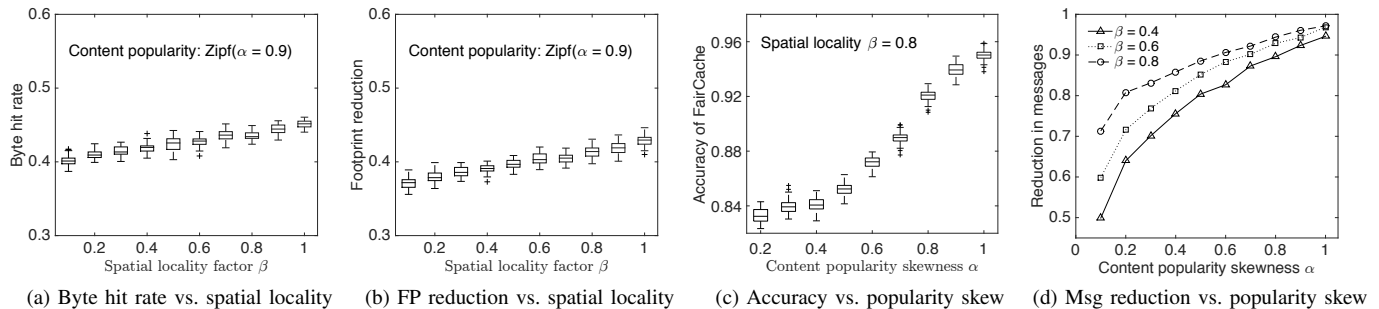
(a) Byte hit rate vs. spatial locality   (b) FP reduction vs. spatial locality   (c) Accuracy vs. popularity skew   (d) Msg reduction vs. popularity skew

Fig. 5: Experiments on the Guifi network, 4 GB cache size. We vary both content popularity skewness α and spatial locality factor β from 0.1 to 1. We observe a gradual and slow improvement in caching performance as spatial locality factor increases. Both spatial locality and content popularity skewness have significant impacts on the accuracy and the traffic reduction.

analysis of spatial locality. We have confirmed that the results are representative of the other topologies.

We use a Hawkes process-based algorithm [24], [25] to generate a user request trace. The algorithm is controlled by two parameters: a *content popularity skewness* α and a *spatial locality factor* β. α controls the overall content popularity which follows *Zipf(α)*. The spatial locality factor, $\beta = 0$, means the request pattern reduces to an *Independent Reference Model*; whilst $\beta = 1$ indicates very high spatial localisation (*i.e.* requests for an object often occur in the same locale).

First, we inspect their impact on the caching performance metrics. Figure 5 presents the results by varying both α and β in $(0, 1]$. From Figure 5a and 5b, we observe a shallow improvement on byte hit rate and footprint reduction as β increases. Specifically, they increase by only 6% and 8% respectively when increasing β from 0.1 to 1. This suggests that spatial locality is not a critical requirement for FairCache.

On the other hand, the popularity skew, α, has a more significant impact on the accuracy and message reduction of FairCache. Figure 5c shows that the average accuracy of FairCache improves from 85% to 97% by increasing α from 0.2 to 1. The speed of degradation of accuracy by decreasing α also slows down at certain point ($\alpha = 0.4$). The reason is because the general popularity distribution gets closer to a uniform distribution (due to a small α). Thus, items are randomly requested, which means that each object has a similar utility when being cached. Interestingly, this means the overall utility of a cache will not vary much, though the solution can be quite different from the optimal one.

Last, we inspect the messaging overhead of running Fair-Cache, presented as the reduction in comparison to the distributed optimal solution again. In Figure 5d, we see that both α and β have a notable impact. Higher α and β both result in lower overheads (*i.e.* higher reductions). The reason is that a smaller α value leads to a more uniform popularity distribution, which makes the demand matrices deviate more from each other, which further leads to larger exchanged messages for $\boldsymbol{\lambda}$ values. The smaller β values have almost the same effect on demand matrices as that of α. However, we also notice that β has more significant impacts when α is small.

G. Neighbourhood Size Distribution

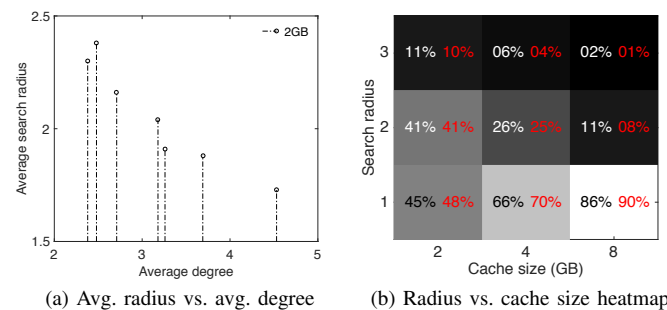(a) Avg. radius vs. avg. degree   (b) Radius vs. cache size heatmap

Fig. 6: The distribution of neighbourhood size (in terms of r) after FairCache converges. (a) uses seven ISP topologies; there is a relatively strong negative correlation between nodes' average degree and their average neighbourhood size. (b) uses Guifi topology with three cache configurations; the numbers in the grid show the percentage of the nodes and percentage of the control messages (in red).

As previously stated in § VI, maintaining some neighbourhoods is critical for ensuring scalability. This is because (i) it dictates the communication overhead; and (ii) it justifies the effectiveness of the heuristic (iii). In this section, we present our empirical study on the neighbourhood size. We launch a number of experiments using the default setup, whilst varying the topology (as this is what dictates the neighbourhood size). We measure the neighbourhood sizes after FairCache converges (represented by the search radius r).

Figure 6a plots the average search radius, r, as a function of the average degree of the topology. We use seven $r1$-level ISP topologies (*i.e.* ISP router-level topology with one-hop clients included), with a 2GB cache size. As the average degree increases, r decreases, thereby reducing network overhead. This is because network density increases, it becomes less necessary to create multi-hop neighbourhoods. Most important is the fact that even with a low degree, the search radius is very small on all topologies (around 2 hops).

Figure 6b presents results for experiments performed on the Guifi network. It shows a heatmap, which reveals the

percentage of nodes that have a certain radius (across three cache sizes). Each point in the grid shows the percentage of nodes in a simulation that have a certain search radius (lighter colour means more), as well as the percentage of control messages generated within each neighbourhood (shown in red). Figure 6b conveys two pieces of important information. First, most of the nodes end up with a neighbourhood of fewer than 3 hops. Even for the small 2GB cache sizes, 87% of the nodes have no more than 2-hop search radius. Second, using larger cache size makes the distribution even more skewed, which further leads to even smaller search radius. When we increase the cache size from 2GB to 8GB, the percentage of nodes with a 1-hop radius increases from 45% to 86%. With 8 GB cache, 97% of nodes' final neighbourhoods are no more than 2 hops. In addition, we also provide the distribution of control messages in each r-hop neighbourhood, plotted in red colour in the same grid. As we can see, the traffic distribution is even more condensed within a very small neighbourhood, which indicates most interactions are between directly connected neighbours. We have confirmed that these results are mirrored across all other topologies.

VIII. DISCUSSION AND LIMITATIONS

Deployment of FairCache raises a number of interesting questions. A key practical concern is the potential for parties to participate in FairCache in a malicious or non-cooperative manner. This is a possibility as nodes are expected to report shadow prices, which they could be manipulated. FairCache is not intended to force stakeholders to collaborate, or to protect against cheating; hence, we have assumed that all nodes adhere to the FairCache algorithm. However, if deployed, such complexities would need to be addressed. In current inter-domain network protocols this problem is handled using out-of-band trust establishment, alongside signing authorities to bind autonomous systems to trusted identities (*e.g.* RPKI [32]). Equally, we envisage FairCache could rely on similar principles, in which legally formed (potentially transitive) collaboration agreements are underpinned by public key cryptography.

There are also a number of alternative practical concerns. We implicitly assumed that the demand matrix is stationary, whereas in reality the demand will change over time. Recall that the space complexity (in Section VI) of storing an individual demand matrix in FairCache is $\Theta(|O'||\overline{N}|)$, hence the complexity of updating the whole matrix is bounded by $\Theta(|O'||\overline{N}|)$. Two other facts further help to reduce this overhead: (*i*) the spatial locality we have mentioned; and (*ii*) the various research [26] that has reported that such changes are gradual. Hence we can incrementally update the demand matrix to avoid unnecessary traffic. Furthermore, such occasional and incremental updates do not necessarily need to be synchronised especially when FairCache is deployed as an ever-running background process. In theory, the stale information may slow down the convergence in an iterative optimisation process. In practice, less-frequent and incremental updates can ameliorate such impacts [23]. Moreover, due to the nature of FairCache algorithm, the updates mostly affect the local neighbourhood and their cascading effects will drastically decrease out of the neighbourhood.

Although we have focussed on building an efficient algorithm, we also acknowledge that FairCache treats storage as the key bottleneck. There are also a number of other constraints that could be included [33]. For instance, hardware bottlenecks can render servers useless even whilst in possession of content (*e.g.* CPU, I/O bus, congestion collapse). Thus, deployment would probably involve the introduction of such considerations into our model of fairness and utility. This could, for example, result in caches actively storing the *same* object in an attempt to share heavy load. Another simplifying assumption is the modelling of delay using hop count (like BGP); whereas this is a useful abstraction, it does not consider the variability introduced by realtime congestion. We consider this an acceptable sacrifice, as introducing such realtime considerations would introduce burdensome overheads. Another point worth highlighting is that we base caching decisions on per-object popularity counts, therefore introducing greater memory overheads than algorithms like Least Recently Used. We emphasise, however, that our heuristic removes all unpopular content, making such counts highly feasible.

IX. RELATED WORK

There are three key related areas of work: collaborative caching, content delivery networks (CDNs) and game theoretical studies of caching. Collaborative in-network caching has been proven as an effective methodology to improve system performance in various contexts [4]–[8], [10], [13], [34], even though edge caching has also been shown to be effective [35]. Previously proposed solutions are either limited by a centralised solver [6], [36] which makes scalability difficult, or limited by distributed heuristics [4], [5], [7], [8], [10], [13], [34], which neither guarantees a global optimum nor Pareto efficiency. FairCache is most related to the latter in that we do not guarantee a global optimum; however, we build on their contributions by introducing the concept of fairness and ensuring Pareto efficiency. Importantly, we also reveal the need for fairness to encourage engagement by cache operators. Recently, fair cache sharing also attracts enough attention in cloud computing and system research, *e.g.* [37] proposes FairRide using blocking to achieve isolation-guarantee and strategy-proofness properties.

The current solution used for Internet-scale content delivery are CDNs. They hold many similarities to ICNs [35], however, unlike our proposal, they are not collaborative entities. Typically, they are operated by distinct companies that deploy independent infrastructures. Some, like Akamai, sell their capacity to third party content providers (arguably a form of collaboration), whilst others build dedicated infrastructures for their own content (*e.g.* Google, Facebook, Netflix). Recent work within the IETF has endeavoured to support inter-CDN cache sharing [38], however, this only provides protocol support, rather than algorithms to decide when, where and how caches should be shared. Hence, our work is orthogonal, and could be applied to CDNs as a recent proposal in [39].

Game theory is an effective tool to analyse the effects of individual behaviours in a complex system; *e.g.* prior work

[40] analyses the fairness achieved in bandwidth allocation by coordinated and uncoordinated rate control over multiple links in peer-to-peer networks. Recent work [3], [36], [41]–[44] applies game theory to study in-network caching. In these papers, the caching problem is modelled as non-cooperative, pure strategic games and the equilibrium is analysed. Unlike us, these work take a system-level utilitarian approach that aims to achieve a global optimum. In contrast, we focus on attaining fairness amongst nodes. More related to us is [3], [36], [42], which look at how selfishness drives nodes to act. These studies show how selfishness impact the equilibrium and efficiency in cache systems (measured by the *Price of Anarchy*). They also show that the global optimum is seldom achieved due to lack of coordination and nodes' inherent selfishness. Again, fairness is overlooked though; we introduce this as an integral requirement of cooperation.

To the best of our knowledge, no prior work in ICN has tried to solve the collaborative caching as a bargaining game and has devised a low-complexity heuristic to embrace both efficiency and fairness.

X. CONCLUSION

To date, studies of collaborative ICN caching have focussed on traditional metrics such as hit rate, assuming that nodes are happy to contribute to achieving a global optimum. In this paper, we have argued that practical situations are unlikely to adhere to this model. Instead, caches operated by separate stakeholders will expect a reasonable level of *fairness*, where they are not penalised for cooperating with others. We began by delineating an optimal solution, which ensures no node attains lower utility by collaborating. To address its high complexity, we have also proposed a heuristic algorithm, FairCache, which we have shown achieves high performance at a fraction of the cost. Unlike past work, FairCache offers Pareto efficiency and proportional fairness, ensuring that *all* nodes are incentivised to collaborate. As well as helping to promote cooperation, our results show that proportional fairness plays a key role in balancing network traffic too. It helps maintain more hits at the edge, rather than globally optimal solutions that centralise hits in the core. We are not prescriptive in how FairCache is deployed and have ensured that it can be used either globally or amongst a subset of collaborating nodes. Hence, our key take-home message is that future collaborative caching designs should cease to assume purely altruistic cooperation and, instead, be explicitly built around the concept of fairness.

There are several lines of potential future work. First, we plan to build a wireline protocol to implement FairCache's design. The most prominent challenge in this regard is implementing a FairCache protocol that is robust against cheating nodes. This is a fascinating area of future work; currently, we assume trusted certified parties that do not lie, however, we wish to expand this to cover more dynamic arrangements in which trust can be formed on-the-fly. Note that this would not necessarily involve significant changes to FairCache — simply augmentary functions, *e.g.* key exchange. Clearly, this should be underpinned by a hardware implementation for exploring

practical feasibility at line rates. There are various other real-world concerns that could also be integrated into FairCache too. For instance, dynamics regarding link availability, congestion and request patterns could be explored. This should extend to integrating new constraints (*e.g.* bandwidth, power, CPU), as well as alternate forms of fairness (*e.g.* bandwidth fairness, user-centric fairness). Lastly, we wish to expand FairCache to consider situations in which caches have external influences (*e.g.* business arrangements) that modify their behaviours. As of yet, little work has considered exogenous incentives that drive caching collaboration. We therefore see this as a fruitful line of study.

APPENDIX

Let $T : \mathbb{R}^n \to \mathbb{R}^n$ be the sorting operator used in [16]. More precisely, that is, we say $\mathbf{U}$ is lexicographically larger than or equal to $\mathbf{U}'$ if we can write $\mathbf{U} \succeq_{\text{lex}} \mathbf{U}' \iff U_i \geq U_i', \forall i \in [0, n)$ after applying T to both $\mathbf{U}$ and $\mathbf{U}'$. $\succ_{\text{lex}}$ can be similarly defined. If we let $\mathbf{U}^*$ be the corresponding utilities achieved by using $(\mathbf{x}^*, \mathbf{y}^*)$ which fulfils a certain well-defined fairness. Let U_i^w denote the minimum utility of v_i of a given solution. Technically, we have the following definitions of the two fairness metrics besides the previously defined Proportional Fairness (PF).

Definition 6. *Egalitarian Fairness (EF):* $(\mathbf{x}^*, \mathbf{y}^*)$ *is EF iff* $\forall(\mathbf{x}, \mathbf{y}) \neq (\mathbf{x}^*, \mathbf{y}^*) \Rightarrow$ *the following two conditions cannot be both true at the same time. (1)* $\exists i$, *s.t.* $U_i \geq U_i^*$*; (2)* $\forall i, j$ *s.t.* $U_i - U_i^w = U_j - U_j^w$.

Definition 7. *Max-Min Fairness (MF):* $(\mathbf{x}^*, \mathbf{y}^*)$ *is MF iff* $\forall(\mathbf{x}, \mathbf{y}) \neq (\mathbf{x}^*, \mathbf{y}^*) \Rightarrow T(\mathbf{U}^*) \succeq_{\text{lex}} T(\mathbf{U})$.

These definitions are followed by the theorems below, both of which have rather straightforward proofs as we will show in the following.

Theorem 3. *In a fair collaborative game* (Ω, u^0)*, the optimal caching strategy* $(\mathbf{x}^*, \mathbf{y}^*)$ *achieves PF.*

Theorem 4. *In a fair collaborative game* (Ω, u^0) *with optimal strategy* $(\mathbf{x}^*, \mathbf{y}^*)$*, EF is sufficient for MF, i.e. EF $\Rightarrow$ MF.*

In our caching games, *PF* is naturally guaranteed by NBS as Theorem 3 states. It is also intuitively easy to understand that if a Pareto efficient solution exists and achieves *EF*, it also achieves *MF* at the same time, as Theorem 4 states.

A. Proof of Theorem 3.

Proof. Because $(\mathbf{x}^*, \mathbf{y}^*)$ is the optimal caching solution, namely $(\mathbf{x}^*, \mathbf{y}^*) = \arg\max_{\mathbf{x},\mathbf{y}} \sum_{v_i \in V} \ln(U_i - u_i^0)$. Let $f(\mathbf{U}) = \sum_{v_i \in V} \ln(U_i - u_i^0)$. For $f(\mathbf{U})$ to reach its maximum, the necessary and sufficient first order condition is $\nabla f^* = 0$.

$\forall (\mathbf{x}, \mathbf{y}) \neq (\mathbf{x}^*, \mathbf{y}^*) \Rightarrow \exists \boldsymbol{\lambda} \succ 0$ such that $\lambda_i^{-1} = U_i - u_i^0 > 0$. Then $\forall v_i \in V$ we have

$$\nabla f^* - \boldsymbol{\lambda} \prec 0 \Longrightarrow \frac{\partial f^*}{\partial U_i^*} - \lambda_i < 0$$

$$\Longrightarrow \frac{1}{U_i^* - u_i^0} - \lambda_i < 0$$

$$\Longrightarrow \frac{\lambda_i^{-1}}{U_i^* - u_i^0} - \frac{U_i^* - u_i^0}{U_i^* - u_i^0} < 0$$

Sum over all the $v_i \in V$, we have

$$\sum_{v_i \in V} \frac{(U_i - u_i^0) - (U_i^* - u_i^0)}{U_i^* - u_i^0} < 0 \Longrightarrow \sum_{v_i \in V} \frac{U_i - U_i^*}{U_i^* - u_i^0} < 0$$

By definition 5, strategy $(\mathbf{x}^*, \mathbf{y}^*)$ is proportionally fair. □

B. Proof of Theorem 4.

Proof. Instead of directly proving that "*EF* is sufficient for *MF* in a fair collaborative game (Ω, u^0)", we first prove that "*EF* is sufficient for *KS* fairness". Recall that *KS* optimises the players of the worst utility, hence *KS* fairness refers to the strategy that maximises the minimum utility in a game.

We prove the theorem by contradiction. Let's assume solution $(\mathbf{x}^*, \mathbf{y}^*)$ is egalitarian fair, but not *KS* fair. $\mathbf{U}^*$ is the corresponding utility value.

Let's further assume another solution $(\mathbf{x}', \mathbf{y}') \neq (\mathbf{x}^*, \mathbf{y}^*)$ which achieves *KS* fairness, and $\mathbf{U}'$ is its utility value. In a fair collaborative game, based on the nature of Nash bargaining framework, both $(\mathbf{x}', \mathbf{y}')$ and $(\mathbf{x}^*, \mathbf{y}^*)$ are Pareto optimal.

By definition, *KS* fair solution indicates that

$$\min\{U_i' - U_i^w, ...\} > \min\{U_i^* - U_i^w, ...\}, \quad \forall v_i \in V \quad (21)$$

By definition, egalitarian fair solution indicates that

$$\min\{U_i^* - U_i^w, ...\} = U_i^* - U_i^w = U_j^* - U_j^w, \forall v_i, v_j \in V \quad (22)$$

(21), (22) $\Longrightarrow$

$$U_i' - U_i^w \geq U_i^* - U_i^w, \quad \forall v_i \in V \quad (23)$$

$$U_i' - U_i^w > U_i^* - U_i^w, \quad \exists v_i \in V \quad (24)$$

Inequality (24) contradicts with the fact that $(\mathbf{x}^*, \mathbf{y}^*)$ is Pareto optimal. So the assumption does not hold. $(\mathbf{x}^*, \mathbf{y}^*)$ must be both egalitarian fair and *KS* fair.

Because $(\mathbf{x}^*, \mathbf{y}^*)$ is already a Pareto optimal solution, both *MF* and *KS* are the same. Therefore, in a fair collaborative game (Ω, u^0), *EF* is sufficient for *MF*. I.e., $EF \Rightarrow MF$. □

C. Proof of Lemma 1.

Proof. The proof is trivial. Since U_i in eq. (3) is affine and positive, the non-negative weighted sum of U_i is still affine and positive. All the affine functions are log-concave. So the objective function (4) is concave.

In addition, all (5)(6)(7)(8) and (9) are defined over a set of compact and convex constraints. Therefore, problem (4) is a convex optimisation problem. □

D. Proof of Theorem 1.

Proof. Obviously caching decision space $[0,1] \subset \mathbb{R}_+$ is a nonempty, compact and convex set. Since the objective function (4) is a continuously differentiable concave function, and all the constraints on the variables are affine, Karush-Kuhn-Tucker (KKT) conditions are necessary and sufficient for the existence of an optimal solution.

To derive the optimum of a function with constraints, we first derive the Lagrangian $\mathcal{L}(\cdot)$ of eq. (4). Let $\boldsymbol{\alpha} \succeq 0$, $\boldsymbol{\beta} \succeq 0$, $\boldsymbol{\gamma} \succeq 0$, $\boldsymbol{\delta} \succeq 0$ and $\boldsymbol{\lambda} \succeq 0$ be the KKT multipliers associated with constraints. Their subscripts are self-explained by the corresponding constraints associated with. Then we have

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\delta}) =$$

$$\sum_{v_i \in V} \ln(U_i - u_i^0) - \sum_{v_i \in V} \sum_{v_j \in N_i} \sum_{o_k \in O} \lambda_{i,j,k}(y_{i,j,k} - x_{j,k})$$

$$- \sum_{v_i \in V} \alpha_i \left(\sum_{o_k \in O} x_{i,k} - C_i \right) - \sum_{v_i \in V} \sum_{o_k \in O} \beta_{i,k} \left(\sum_{v_j \in N_i} y_{i,j,k} - 1 \right)$$

$$- \sum_{v_i \in V} \sum_{o_k \in O} \gamma_{i,k}(x_{i,k} - 1) + \sum_{v_i \in V} \sum_{v_j \in N_i} \sum_{o_k \in O} \delta_{i,j,k} y_{i,j,k}$$

Note we dropped constraints $x_{i,j} \geq 0$ and $y_{i,j,k} \leq 1$ in making the Lagrangian because constraints (6) and (7) make them redundant. In the following derivation, we let $\tau_{i,k} = U_i - u_i^0 - w_{i,k} x_{i,k}$ and $\tau_{i,k}' = U_i - u_i^0 - \frac{w_{i,k}}{l_{i,j}} y_{i,j,k}$ for the simplicity of representation. For the objective function to reach its optimum, first order necessary and sufficient conditions are

$$\nabla \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\delta}) = 0$$

$$\Longleftrightarrow \frac{\partial \mathcal{L}}{\partial x_{i,k}} = 0, \forall v_i, v_j \in V, \forall o_k \in O$$

$$\Longleftrightarrow \frac{w_{i,k}}{U_i - u_i^0} + \sum_{v_j \in N_i^+} \lambda_{j,i,k} - \alpha_i - \gamma_{i,k} = 0$$

$$\Longleftrightarrow x_{i,k}^* = \frac{1}{\alpha_i + \gamma_{i,k} - \sum_{v_j \in N_i^+} \lambda_{j,i,k}} - \frac{\tau_{i,k}}{w_{i,k}}$$

with complementary slackness

$$\begin{cases} \lambda_{i,j,k}(y_{i,j,k} - x_{j,k}) = 0, & \forall v_i, v_j \in V, \forall o_k \in O \\ \alpha_i(\sum_{o_k \in O} x_{i,k} - C_i) = 0, & \forall v_i \in V, \forall o_k \in O \\ \beta_{i,k}(\sum_{v_j \in N_i} y_{i,j,k} - 1) = 0, & \forall v_i \in V, \forall o_k \in O \quad (25) \\ \gamma_{i,k}(x_{i,k} - 1) = 0, & \forall v_i \in V, \forall o_k \in O \\ \delta_{i,j,k} y_{i,j,k} = 0, & \forall v_i, v_j \in V, \forall o_k \in O \end{cases}$$

Similarly, we can derive the optimal $y_{i,j,k}^*$ as $x_{i,k}^*$. The optimal caching strategy $(\mathbf{x}^*, \mathbf{y}^*)$ of the network can be derived by solving the equation system (25) for all the nodes. □

E. Proof of Theorem 2.

Proof. To prove convergence, we first prove the gradient of the dual function is bounded by a constant K, namely the dual function $d(\boldsymbol{\lambda})$ is K-Lipschitz continuous. Second, we show that given the diminishing step size, the Euclidean distance between the optimum $d(\boldsymbol{\lambda}^*)$ and the best value $d(\boldsymbol{\lambda}^\circ)$ achieved in all previous iterations converges to zero in limit.

Since the primal (2) is strictly convex and all constraints are linear, dual $d(\boldsymbol{\lambda})$ is strictly concave and differentiable.

$$\frac{\partial d(\boldsymbol{\lambda})}{\partial \lambda_{i,j,k}} = y_{i,j,k} - x_{i,k} \implies \left|\frac{\partial d(\boldsymbol{\lambda})}{\partial \lambda_{i,j,k}}\right| \leq 1 \qquad (26)$$

By Mean value theorem, there exists $\mathbf{c} \in (\boldsymbol{\lambda}, \boldsymbol{\lambda}')$ such that

$$d(\boldsymbol{\lambda}) - d(\boldsymbol{\lambda}') = \nabla d(\mathbf{c})^T (\boldsymbol{\lambda} - \boldsymbol{\lambda}') \qquad (27)$$

By Cauchy–Schwarz inequality, let $n = |O| \times |V|^2$, we have

$$\|d(\boldsymbol{\lambda}) - d(\boldsymbol{\lambda}')\|_2 = \|\nabla d(\mathbf{c})^T (\boldsymbol{\lambda} - \boldsymbol{\lambda}')\|_2 \qquad (28)$$
$$\leq \|\nabla d(\mathbf{c})\|_2 \|\boldsymbol{\lambda} - \boldsymbol{\lambda}'\|_2 \qquad (29)$$
$$\leq \sqrt{n} \|\boldsymbol{\lambda} - \boldsymbol{\lambda}'\|_2 \qquad (30)$$

$\|\cdot\|_2$ above denotes the Euclidean norm. Therefore, $d(\boldsymbol{\lambda})$ is K-Lipschitz continuous and Lipschitz constant $K = \sqrt{n}$. Let $\boldsymbol{\lambda}^*$ denote the maximiser of dual function $d(\boldsymbol{\lambda})$, then

$$\|\boldsymbol{\lambda}^{(t+1)} - \boldsymbol{\lambda}^*\|_2^2 = \|(\boldsymbol{\lambda}^{(t)} + \xi_k \mathbf{h}^{(t)})_+ - \boldsymbol{\lambda}^*\|_2^2 \qquad (31)$$
$$\leq \|\boldsymbol{\lambda}^{(t)} + \xi_k \mathbf{h}^{(t)} - \boldsymbol{\lambda}^*\|_2^2 \qquad (32)$$
$$= \|\boldsymbol{\lambda}^{(t)} - \boldsymbol{\lambda}^*\|_2^2 + 2\xi_k \mathbf{h}^{(t)T}(\boldsymbol{\lambda}^{(t)} - \boldsymbol{\lambda}^*) + \xi_k^2 \|\mathbf{h}^{(t)}\|_2^2 \quad (33)$$
$$\leq \|\boldsymbol{\lambda}^{(t)} - \boldsymbol{\lambda}^*\|_2^2 + 2\xi_k (d(\boldsymbol{\lambda}^{(t)}) - d(\boldsymbol{\lambda}^*)) + \xi_k^2 \|\mathbf{h}^{(t)}\|_2^2 \quad (34)$$

Inequality (32) comes from the fact that projection of a point onto $\mathbb{R}_+^{|O||V|^2}$ makes it closer to the optimal point in $\mathbb{R}_+^{|O||V|^2}$. Apply inequality (34) recursively, we have

$$\|\boldsymbol{\lambda}^{(t+1)} - \boldsymbol{\lambda}^*\|_2^2 \leq$$
$$\|\boldsymbol{\lambda}^{(1)} - \boldsymbol{\lambda}^*\|_2^2 + 2\sum_{i=1}^{k} \xi_i (d(\boldsymbol{\lambda}^{(i)}) - d(\boldsymbol{\lambda}^*)) + \sum_{i=1}^{k} \xi_i^2 \|\mathbf{h}^{(i)}\|_2^2$$

Because $\|\boldsymbol{\lambda}^{(t+1)} - \boldsymbol{\lambda}^*\|_2^2 \geq 0$ and $\sum_{i=1}^{k} \xi_i > 0$, and let $d(\boldsymbol{\lambda}^{\circ}) = \max_{0 \leq i < k} d(\boldsymbol{\lambda}^{(i)})$, then

$$2\sum_{i=1}^{k} \xi_i (d(\boldsymbol{\lambda}^*) - d(\boldsymbol{\lambda}^{\circ})) \leq \|\boldsymbol{\lambda}^{(1)} - \boldsymbol{\lambda}^*\|_2^2 + \sum_{i=1}^{k} \xi_i^2 \|\mathbf{h}^{(i)}\|_2^2$$
$$\implies d(\boldsymbol{\lambda}^*) - d(\boldsymbol{\lambda}^{\circ}) \leq \frac{\|\boldsymbol{\lambda}^{(1)} - \boldsymbol{\lambda}^*\|_2^2 + \sum_{i=1}^{k} \xi_i^2 \|\mathbf{h}^{(i)}\|_2^2}{2\sum_{i=1}^{k} \xi_i}$$
$$\implies d(\boldsymbol{\lambda}^*) - d(\boldsymbol{\lambda}^{\circ}) \leq \frac{\|\boldsymbol{\lambda}^{(1)} - \boldsymbol{\lambda}^*\|_2^2 + K^2 \sum_{i=1}^{k} \xi_i^2}{2\sum_{i=1}^{k} \xi_i}$$

$d(\boldsymbol{\lambda}^*) - d(\boldsymbol{\lambda}^{\circ}) \to 0$ if we choose a diminishing step size which lets $\xi_i \to 0$ and $\sum_1^{\infty} \xi_i = \infty$, then $\frac{\sum_1^{\infty} \xi_i^2}{\sum_1^{\infty} \xi_i} = 0$. (e.g. we can let $\xi_i = \frac{\xi_0}{i}$, then $\sum_1^{\infty} \xi_i = \infty$ and $\sum_1^{\infty} \xi_i^2 = \frac{\pi^2}{6}$.) Since the duality gap is zero, eventually the primal problem will converge to its optimum when its dual problem converges. $\square$

REFERENCES

[1] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," in *Proceedings of the 5th ACM Conext.* New York, NY, USA: ACM, 2009, pp. 1–12.

[2] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A data-oriented (and beyond) network architecture," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 181–192, 2007.

[3] V. Pacifici and G. Dán, "Selfish content replication on graphs," in *Proceedings of the 23rd International Teletraffic Congress*, ser. ITC '11. ITCP, 2011, pp. 119–126. [Online]. Available: http://dl.acm.org/citation.cfm?id=2043468.2043488

[4] E. Rosensweig and J. Kurose, "Breadcrumbs: Efficient, best-effort content location in cache networks," in *INFOCOM 2009, IEEE*, April 2009, pp. 2631–2635.

[5] J. Dai, Z. Hu, B. Li, J. Liu, and B. Li, "Collaborative hierarchical caching with dynamic request routing for massive content distribution," in *INFOCOM, 2012 Proceedings IEEE*, March 2012, pp. 2444–2452.

[6] M. D. Dahlin, R. Y. Wang, T. E. Anderson, and D. A. Patterson, "Cooperative caching: Using remote client memory to improve file system performance," in *Proceedings of the 1st USENIX Conference on Operating Systems Design and Implementation*, ser. OSDI '94. Berkeley, CA, USA: USENIX Association, 1994. [Online]. Available: http://dl.acm.org/citation.cfm?id=1267638.1267657

[7] W. K. Chai, D. He, I. Psaras, and G. Pavlou, "Cache "less for more" in information-centric networks," in *Proceedings of the 11th international IFIP TC 6 conference on Networking - Volume Part I*, ser. IFIP'12. Berlin, Heidelberg: Springer-Verlag, 2012, pp. 27–40.

[8] S. Saha, A. Lukyanenko, and A. Yla-Jaaski, "Cooperative caching through routing control in information-centric networks," in *INFOCOM, 2013 Proceedings IEEE*, April 2013, pp. 100–104.

[9] W. Wong, L. Wang, and J. Kangasharju, "Neighborhood Search and Admission Control in Cooperative Caching Networks," in *Proc. of IEEE GLOBECOM*, December 3-7 2012.

[10] I. Psaras, W. K. Chai, and G. Pavlou, "Probabilistic in-network caching for information-centric networks," in *Proceedings of the Second Edition of the ICN Workshop on Information-centric Networking*, ser. ICN '12. New York, NY, USA: ACM, 2012, pp. 55–60. [Online]. Available: http://doi.acm.org/10.1145/2342488.2342501

[11] L. Wang, "Content, topology and cooperation in in-network caching," ser. Series of Publications A. Helsinki, Finland: University of Helsinki, Department of Computer Science, Mar. 2015, pp. 472–484. [Online]. Available: http://urn.fi/URN:ISBN:978-951-51-0825-8

[12] D. Vega, L. Cerda-Alabern, L. Navarro, and R. Meseguer, "Topology patterns of a community network: Guifi.net," in *Wireless and Mobile Computing, Networking and Communications (WiMob), 2012 IEEE 8th International Conference on*, Oct 2012, pp. 612–619.

[13] S. Borst, V. Gupta, and A. Walid, "Distributed caching algorithms for content distribution networks," in *Proceedings of IEEE INFOCOM*. IEEE, 2010, pp. 1–9.

[14] L. Wang, S. Bayhan, and J. Kangasharju, "Effects of cooperation policy and network topology on performance of in-network caching," *IEEE Communications Letters*, vol. 18, no. 4, pp. 680–683, April 2014.

[15] ——, "Optimal chunking and partial caching in information-centric networks," *Computer Communications*, vol. 61, pp. 48 – 57, 2015. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0140366414003740

[16] D. Bertsimas, V. F. Farias, and N. Trichakis, "The price of fairness," *Operations research*, vol. 59, no. 1, pp. 17–31, 2011.

[17] J. Nash, John F., "The bargaining problem," *Econometrica*, vol. 18, no. 2, pp. pp. 155–162, 1950. [Online]. Available: http://www.jstor.org/stable/1907266

[18] E. Kalai and M. Smorodinsky, "Other solutions to nash's bargaining problem," *Econometrica*, vol. 43, no. 3, pp. 513–18, 1975. [Online]. Available: http://EconPapers.repec.org/RePEc:ecm:emetrp:v:43:y:1975:i:3:p:513-18

[19] F. P. Kelly, A. K. Maulloo, and D. K. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research society*, vol. 49, no. 3, pp. 237–252, 1998.

[20] H. Boche and M. Schubert, "Nash bargaining and proportional fairness for wireless systems," *Networking, IEEE/ACM Transactions on*, vol. 17, no. 5, pp. 1453–1466, 2009.

[21] A. Muthoo, *Bargaining theory with applications.* Cambridge University Press, 1999.

[22] L. .Wang, G. Tyson, J. Kangasharju, and J. Crowcroft, "Faircache technical report: Introducing fairness to icn caching." [Online]. Available: https://www.dropbox.com/s/jni7txgofqvugl5/report.pdf?dl=0

[23] S. Boyd and L. Vandenberghe, *Convex optimization.* Cambridge university press, 2004.

[24] L. Wang, S. Bayhan, J. Ott, J. Kangasharju, A. Sathiaseelan, and J. Crowcroft, "Pro-diluvian: Understanding scoped-flooding for content discovery in information-centric networking," in *Proceedings of the 2Nd International Conference on Information-Centric Networking*, ser. ICN '15. New York, NY, USA: ACM, 2015, pp. 9–18. [Online]. Available: http://doi.acm.org/10.1145/2810156.2810162

[25] A. Dabirmoghaddam, M. M. Barijough, and J. Garcia-Luna-Aceves, "Understanding optimal caching and opportunistic caching at "the edge" of information-centric networks," in *Proceedings of the 1st*

International Conference on Information-centric Networking, ser. ICN '14. New York, NY, USA: ACM, 2014, pp. 47–56. [Online]. Available: http://doi.acm.org/10.1145/2660129.2660143

[26] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, and S. Moon, "I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, ser. IMC '07. New York, NY, USA: ACM, 2007, pp. 1–14. [Online]. Available: http://doi.acm.org/10.1145/1298306.1298309

[27] X. Cheng, C. Dale, and J. Liu, "Statistics and social network of youtube videos," in *Quality of Service, 2008. IWQoS 2008. 16th International Workshop on*, june 2008, pp. 229 –238.

[28] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP topologies with rocketfuel," in *Proceedings of ACM SIGCOMM*, 2002.

[29] L. Wang, A. Sathiaseelan, J. Crowcroft, and J. Kangasharju, "Litelab: Efficient large-scale network experiments," in *2016 13th IEEE Annual Consumer Communications Networking Conference (CCNC)*, Jan 2016, pp. 60–66.

[30] B. Huffaker, M. Fomenkov, and k. claffy, "Internet Topology Data Comparison," Cooperative Association for Internet Data Analysis (CAIDA), Tech. Rep., May 2012.

[31] S. Bornholdt, H. G. Schuster, and J. Wiley, *Handbook of graphs and networks*. Wiley Online Library, 2003, vol. 2.

[32] M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, S. Uhlig, and G. Tyson, "Ripki: The tragic story of rpki deployment in the web ecosystem," in *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*, ser. HotNets-XIV. New York, NY, USA: ACM, 2015, pp. 11:1–11:7. [Online]. Available: http://doi.acm.org/10.1145/2834050.2834102

[33] D. Perino and M. Varvello, "A reality check for content centric networking," in *Proceedings of the Workshop on Information-centric networking*, 2011.

[34] L. Fan, P. Cao, J. Almeida, and A. Z. Broder, "Summary cache: A scalable wide-area web cache sharing protocol," *IEEE/ACM Trans. Netw.*, vol. 8, no. 3, pp. 281–293, Jun. 2000. [Online]. Available: http://dx.doi.org/10.1109/90.851975

[35] S. K. Fayazbakhsh, Y. Lin, A. Tootoonchian, A. Ghodsi, T. Koponen, B. Maggs, K. Ng, V. Sekar, and S. Shenker, "Less pain, most of the gain: Incrementally deployable icn," in *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, ser. SIGCOMM '13. New York, NY, USA: ACM, 2013, pp. 147–158.

[36] B.-G. Chun, K. Chaudhuri, H. Wee, M. Barreno, C. H. Papadimitriou, and J. Kubiatowicz, "Selfish caching in distributed systems: a game-theoretic analysis," in *Proceedings of the twenty-third annual ACM symposium on Principles of distributed computing*, ser. PODC '04. New York, NY, USA: ACM, 2004, pp. 21–30. [Online]. Available: http://doi.acm.org/10.1145/1011767.1011771

[37] Q. Pu, H. Li, M. Zaharia, A. Ghodsi, and I. Stoica, "Fairride: Near-optimal, fair cache sharing," in *Proceedings of the 13th Usenix Conference on Networked Systems Design and Implementation*, ser. NSDI'16. Berkeley, CA, USA: USENIX Association, 2016, pp. 393–406. [Online]. Available: http://dl.acm.org/citation.cfm?id=2930611.2930637

[38] L. Peterson and B. Davie and R. van Brandenburg Ed. (2014, August) Framework for Content Distribution Network Interconnection (CDNI). [Online]. Available: https://tools.ietf.org/html/rfc7336

[39] V. Pacifici and G. Dn, "Distributed algorithms for content allocation in interconnected content distribution networks," in *2015 IEEE Conference on Computer Communications (INFOCOM)*, April 2015, pp. 2362–2370.

[40] P. Key, L. Massoulié, and D. Towsley, "Path selection and multipath congestion control," *Commun. ACM*, vol. 54, no. 1, pp. 109–116, Jan. 2011. [Online]. Available: http://doi.acm.org/10.1145/1866739.1866762

[41] V. Pacifici and G. Dan, "Content-peering dynamics of autonomous caches in a content-centric network," in *INFOCOM, 2013 Proceedings IEEE*, April 2013, pp. 1079–1087.

[42] N. Laoutaris, O. Telelis, V. Zissimopoulos, and I. Stavrakakis, "Distributed selfish replication," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 17, no. 12, pp. 1401–1413, Dec 2006.

[43] G. Dan, "Cache-to-cache: Could isps cooperate to decrease peer-to-peer content distribution costs?" *Parallel and Distributed Systems, IEEE Transactions on*, vol. 22, no. 9, pp. 1469–1482, Sept 2011.

[44] M. Dehghan, L. Massoulie, D. Towsley, D. Menasche, and Y. C. Tay, "A utility optimization approach to network cache design," in *IEEE INFOCOM 2016 - The 35th Annual IEEE International Conference on Computer Communications*, April 2016, pp. 1–9.

Liang Wang is a research associate in the Computer Laboratory at University of Cambridge, United Kingdom. In 2003, he received his BEng in Computer Science and Mathematics from Tongji University, Shanghai, China. Later, he received both his MSc and PhD degrees in Computer Science from University of Helsinki, Finland in 2011 and 2015 respectively. Liang's research interests include system and network optimisation, modelling and analysis of complex networks, information-centric networks, and distributed data processing, and so on.

Gareth Tyson is a Lecturer at Queen Mary University of London. He receeived his PhD from Lancaster University, UK in 2010. His research centres on system measurements and design, looking at topics ranging from network operations to social media. He serves as a reviewer and program committee member for a number of prominent journals such as IEEE/ACM ToN, IEEE JSAC, IEEE TPDS, IEEE TNSM ACM TMM, IEEE TC. He recieved the Outstanding Reviewer Award at ICWSM'16.

Jussi Kangasharju received his MSc from Helsinki University of Technology in 1998. He received his Diplome d'Etudes Approfondies (DEA) from the Ecole Superieure des Sciences Informatiques (ESSI) in Sophia Antipolis in 1998. In 2002 he received his PhD from University of Nice Sophia Antipolis/Institut Eurecom. In 2002 he joined Darmstadt University of Technology (TUD), first as post-doctoral researcher, and from 2004 onwards as assistant professor. Since June 2007 Jussi is a professor of computer science at University of Helsinki.

Jon Crowcroft (SM95F04) received the B.S. degree in physics from University of Cambridge in 1979, and the MSc degree in computing and PhD degree from UCL in 1981 and 1993, respectively. He has been a Professor at the University of Cambridge since 2001. He has worked in the area of Internet support for multimedia communications for over 30 years. Prof. Crowcroft is a Fellow of the Royal Society, the Association for Computing Machinery (ACM), the British Computer Society, the IET, and the Royal Academy of Engineering.