Evolution of the control plane for future content distribution services

Ghida Ibrahim

ParisTech
INSTITUT DES SCIENCES ET TECHNOLOGIES
PARIS INSTITUTE OF TECHNOLOGY

TELECOM ParisTech

2014-ENST-0032

EDITE - ED 130

Doctorat ParisTech

T H È S E

pour obtenir le grade de docteur délivré par

TELECOM ParisTech

Spécialité « Informatique et Réseaux »

présentée et soutenue publiquement par

Ghida IBRAHIM

le 18 Juin 2014

Evolution du plan de commande pour les futurs services de distribution de contenus

Directeur de thèse : **Daniel KOFMAN**
Co-encadrement de la thèse : **Alexandra ANSIAUX**

Jury

Mme Annie GRAVEY, Professeur, Telecom Bretagne	Rapporteur
M. Olivier FESTOR, Professeur, INRIA Nancy	Rapporteur
M. Dario ROSSI, Professeur, Telecom ParisTech	Président du Jury
M. Stefano SECCI, Maitre de Conférences, Université Pierre et Marie Curie	Examinateur
M. Diego PERINO, Ingénieur, Alcatel Lucent-Bell Labs France	Examinateur
M. James KRETCHMAR, Ingénieur, Akamai Technologies	Examinateur
M. Daniel KOFMAN, Professeur, Telecom ParisTech	Directeur de Thèse
Mme Alexandra ANSIAUX, Ingénieur, Orange Labs	Encadrant
M. Youssef CHADLI, Ingénieur, Orange Labs	Invité

T H È S E

TELECOM ParisTech
école de l'Institut Mines-Télécom - membre de ParisTech
46 rue Barrault 75013 Paris - (+33) 1 45 81 77 77 - www.telecom-paristech.fr

Acknowledgment

The past three years have been an incredible experience in my life, the kind of experience that is at once tough and amazing and that, in each single day, questions our willingness to learn, innovate and overcome challenges no matter how important they are.

During this journey, I had the opportunity to meet a lot of people from different cultural and professional horizons from whom I learned a lot, from both technical and scientific perspectives. But I have also met people that inspired me, supported me and challenged me repetitively. All these people have, in one way or another, pushed me forward in life and career terms and forced me to cross my own, inner border. Therefore, I thank them for being part of my journey. I dedicate to them, supervisors, managers, colleagues, friends and anonymous, this thesis, the fruit of more than three years of hard work.

This thesis is also dedicated to my family : my mom, my dad, my two sisters and my younger brother. When it comes to you, the word "Thank you" is never enough. Yet, it is the only word that I can find for expressing my gratitude. Thank you for so strongly believing in me even in times where I no longer believed in myself. If I am here today, it is mainly thanks to you. I love you.

Abstract

Content distribution services are evolving fast in various directions. One of them is the federation of content delivery network (CDN) providers with different capacities, footprints and cost models. In the context of this thesis, we introduce a CDN federation solution that is based on a centralized control architecture. This solution allows taking static decisions of federation establishment and provisioning and dynamic decisions of federation real-time control. Federation establishment consists of choosing the CDNs involved in a given federation and specifying the market, in terms of content providers (CPs), of the formed federation. It also consists of choosing the business model of the federation and agreeing on a common strategy of revenue sharing. Federation provisioning consists of deciding where, inside the federation, CPs contents are placed and how CPs load is shared. Federation dynamic control consists of directing, in real-time, incoming users' requests towards different CDNs based on the provisioning phase outputs and on a dynamic vision of different CDNs state.

We introduce an optimization model in order to address the static aspect of federation computation and provisioning. This model aims to maximize the joint gain of the federation while taking into account CPs service requirements and CDN providers constraints in terms of capacity and economic fairness. The model that we introduce is used in order to investigate three use cases of interest for the CDN industry. The first use case addresses the federation of Telco-CDNs. The second use case addresses the federation of pure-play CDNs. The third use case addresses the federation of Telcos and pure-play CDNs. For each of the use cases, we assess the economic gains achieved by different categories of CDN players through federating. We demonstrate that, when the market demand in high, CDN providers always have an interest to federate. In particular, some CDN providers can double their economic gains through federating.

We also address the dynamic phase of federation control. In particular, we focus on the control of peak events due to unexpected traffic surges within a federation of CDNs. In this context, we introduce three frameworks that point to two different approaches for events control within a CDN federation. The first approach privilege an intra-CDN control of peak events consisting of a dynamic adaptation of delivered video resolution at individual CDNs level. The second approach advocate a federated behavior involving different members of the federation. We conduct trace-driven simulations and use two key performance indicators in order to assess the performance of the different frameworks. We demonstrate that, when a joint approach for events control is adopted within a federation of CDNs, the proposed federation is better resilient to peak events. This translates into a higher hit ratio of the federation and a better video resolution witnessed by end users.

Our work on CDN federation leads us to focus on the role that the Telco can play in a CDN federation context. In this context, we identify three major value-added services that can be proposed by the Telco to a federation of CDNs or to individual Over The Top (OTTs) players. We suggest enhancements of the Telco control infrastructure in order to enable the proposed services.

French Abstract

Les services de distribution de contenus évoluent rapidement. Un axe majeur d'évolution concerne la collaboration de fournisseurs de réseaux CDN avec différents capacités, couvertures et modèles de coût. Dans le contexte de cette thèse, nous nous focalisons sur une forme particulière de collaboration liée à la fédération de fournisseurs de réseaux CDN mettant ensemble leurs ressources respectives et agissant en tant qu'entité unique par rapport aux fournisseurs de contenus (Content Providers). En particulier, nous proposons une solution technique basée sur une architecture centralisée qui permet de prendre des décisions statiques d'établissement et de provisionnement de fédérations ainsi que des décisions de contrôle dynamique de fédérations établies. Nous adressons les aspects statiques de prise de décision en introduisant un modèle d'optimisation que nous appliquons à trois scénarios de fédération d'interet pour le marché. Le premier scénario traite la fédération d'opérateurs Telecom possèdant des réseaux CDNs et localisés au sein du meme pays. Le second scénario traite la fédération de fournisseurs de réseaux CDNs qui ne sont pas opérateurs de réseau. Le troisième scenario traite la fédération de fournisseurs de réseaux CDNs qui peuvent ou pas etre des opérateurs Telecom. Nous demontrons que, quand la demande sur le marché de distribution de contenu est élevée, les distributeurs de contenus ont interet, d'un point de vue economique, à fédérer. Notamment, dans le premier scénario, les fournisseurs de réseaux CDNs peuvent aller jusqu'à doubler leur revenus en fédérant par rapport au scénario séparé. Dans le contexte de contrôle dynamique de fédérations, nous nous focalisons sur le contrôle d'évenements de pointe dans une fédération de fournisseurs de réseaux CDNs. Différentes approches de contrôle sont valables à ce niveau. Nous effectuons des simulations basées sur des traces de trafic réelles dans le but de comparer les différentes approches. Nous demontrons que, quand une approche jointe de contrôle d'évènements de pointe est adoptée au sein d'une fédération, la fédération réagit mieux à ces évènements. Ceci se traduit en un moindre volume de sessions rejetées et en une meilleure résolution vidéo ressentie par les internautes. Notre travail sur la fédération nous conduit à se focaliser sur le rôle d'un Telco dans un contexte fédéré. En particulier, nous identifions trois services à valeur ajoutée qu'un Telco peut proposer à une federation de fournisseurs de réseaux CDNs ou à des acteurs « Over the Top » (OTT). Nous proposons des améliorations de l'infrastructure de contrôle Telco existante dans le but de permettre ces différents services.

Synthèse en Français

Contexte Général

Il y'a dix ans, les gens comptaient principalement sur les services conversationnels, y compris les appels téléphoniques et les vidéoconférences pour communiquer et rester en contact les uns avec les autres. L'émergence de portails tels YouTube et DailyMotion permettant aux utilisateurs de générer leur propres contenus et la propagation de plates-formes de médias sociaux tels Facebook et Twitter ont contribué à un changement important des attitudes des internautes. Consulter en ligne de blogs, uploader/downloader des vidéos vers/à partir de YouTube, partager des informations/albums via Facebook et tweeter sont tous devenus une composante importante de notre vie de tous les jours. Plus important encore, ces actions presque intuitives, basées sur la consommation, la production et le partage de contenus qui peuvent etre des pages web, des vidéos ou des albums en ligne, sont devenues notre façon de rester connecté avec le reste du monde.

L'Internet d'aujourd'hui parle en termes de consommation, production et partage de contenus. Cette tendance devrait s'accélérer dans les années à venir. En effet, le trafic vidéo seul, y compris la vidéo à la demande (VoD), le live streaming , puis catch-up TV devrait représenter plus de 90 % de trafic IP mondial en 2017. Un volume plus important de ce traffic va etre généré par des appareils mobiles, y compris les téléphones mobiles et les tablette. Avec l'émergence de concepts tels le Machine to Machine (M2M) et l'Internet des objets (IOT), les objets connectés devraient également etre à l'origine de nouveaux types de contenu avec, probablement, une structure différente que les contenus que nous connaissons aujourd'hui (vidéos, pages web, jeux, etc.) En parallèle, les consommateurs sont moins tolérants à des temps de démarrage élevés ou à des événements de buffering pendant la visualisation de sessions vidéos. Ils sont de même de plus en plus exigeants en termes de la qualité ou résolution des vidéos demandées. Ces aspects de l'évolution de l'internet ont un impact direct sur les différents acteurs impliqués dans la chaine de distribution de contenus de bout en bout.
Ces acteurs peuvent etre classés en quatre grandes catégories. Au sommet de la chaine, nous trouvons les producteurs et fournisseurs de contenus (Content Provider ou CP en anglais) qui fournissent directement ou indirectement, des contenus professionnels ou générés par les internautes au grand public. Parmi ces acteurs, nous trouvons BBC, YouTube, Daily Motion et Netflix. La deuxième catégorie est composée de fournisseurs

de réseaux de distribution de contenu (content delivery network ou CDN en anglais), un CDN étant composé de serveurs distribués dans plusieurs points de présence localisés à la bordure de l'internet. En déléguant la livraison de contenus à CDN tiers, les parties contractantes visent à réduire leurs dépenses (Capital Expenditures), à alléger la charge de leurs serveurs et à améliorer la qualité d'expérience ressentie par les utilisateurs. Un CDN fournit les contenus qu'il détient depuis des serveurs situés à la proximité des internautes et utilise des techniques exclusives d'accélération de contenus. Les fournisseurs de réseau CDN bien connus incluent Akamai qui possède 80 % du marché CDN et offre plus de 25 % du trafic Internet, Level3 et Limelight. Nous nous référons à ces fournisseurs de réseau CDN en tant que 'pure-play' CDNs. La troisième catégorie d'acteurs est formé par les opérateurs de réseau, y compris les opérateurs Telecom et les fournisseurs de services de transit qui permettent au contenu d'etre acheminé vers/depuis les utilisateurs finaux. A la fin de la chaine, viennent les utilisateurs/internautes finaux qui sont à la fois consommateurs et producteurs de contenu.
Il est important de noter que la frontière entre ces catégories n'est pas toujours claire. Par exemple, YouTube et Netflix sont deux fournisseurs de contenus qui ont déployé leurs propres réseaux CDN qui sont dédiés pour leur usage interne. En outre, de nombreux opérateurs de télécommunications (telcos) entrent sur le marché CDN en déployant des last-mile CDNs dans les pays de leur empreinte. Nous nous référons aux opérateurs Télécom possédant une plate-forme CDN comme 'Telco-CDN'.

Les difficultés rencontrées par les acteurs de bout en bout de la chaine sont d'ordres technique et économique. L'explosion de trafic entraine une congestion non seulement à l'intérieur des domaines des opérateurs réseaux, mais aussi au niveau des liens d'interconnexion. Cette congestion augmente les couts du trafic entre opérateurs et dégrade la qualité de service (QoS) assurée par le réseau. Les producteurs et les fournisseurs de contenus visent à trouver un compromis entre la construction de leurs propres plateformes CDNs et la délégation de la distribution de leur contenu à des tiers, les deux scénarios ayant des implications économiques en termes d'investissements CAPEX et d'accors de service (Service Level Agreement) avec des tiers.
Malgré leur large dimensionnement et leur grande empreinte, des pure-play CDNs comme Akamai et Limelight peuvent souffrir de limitations en termes de capacité face à certains types d'événements de pointe. Un exemple bien connu des événements de pointe comprend les scénarios Flash Crowd faisant référence à une hausse rapide et inattendue d'un trafic de fournisseur de contenu donné. En outre, pour des questions liées à la qualité d'expérience ressentie par les utilisateurs, certains pure-play CDNs sont intéressés par localiser leurs contenus au plus près des internautes. D'un autre coté, malgré leurs atouts en termes de proximité des utilisateurs et de contrôle du réseau, les Telco-CDNs souffrent souvent de la nature locale de leur empreinte. En se basant sur ce que nous venons de décrire, nous pouvons conclure que les services de distribution de contenu évoluent d'un système centralisé ou le contenu est monopolisé par un seul acteur (cet acteur est le fournisseur de contenu lui-même ou un pure-play CDN) vers un nouveau système impliquant des distributeurs de contenus répartis et autonomes. Cette tendance est accélérée par

l'émergence de nouveaux acteurs, y compris les opérateurs de télécommunications, dans les marchés de CDN ainsi que par l'émergence de techniques telles la virtualisation et le partage des ressources. En résumé, tout acteur qui possède des ressources de stockage et de streaming dans le réseau ou en amont peut jouer le rôle d'un fournisseur CDN et, par conséquent, devenir une partie d'une plate-forme globale et virtuelle de CDNs.

La collaboration de fournisseurs de réseau CDN

Compte tenu de la forte demande du marché, les fournisseurs de réseau CDN, qu'ils soient bien positionnés sur le marché ou des acteurs émergents, sont incités à collaborer. L'amélioration de la qualité d'expérience des utilisateurs, l'agrégation de capacités et de l'empreinte géographique font parties des incitations techniques pour une telle collaboration. La réduction des cout d'investissement et l'amélioration des revenus font partie des incitations économiques. La collaboration de fournisseurs de réseau CDN peut s'effectuer selon un mode centralisé ou distribué.

Une forme de collaboration consiste à établir des accords mutuels (service level agreement ou SLA) entre deux fournisseurs de réseau CDN. Cette forme de collaboration a été étudiée par le groupe de travail CDN interconnexion (CDNI) groupe de l'IETF.
Une autre forme de collaboration consiste à orchestrer le traffic d'un fournisseur de contenu donné entre différents fournisseurs de CDN, chacun avec son propre modèle d'affaires et qui ne sont pas nécessairement en contact direct. Une telle orchestration peut etre effectuée par le fournisseur de contenu lui-même, par un pure-play CDN ou par un orchestrateur tiers. L'orchestration de fournisseurs de réseau CDN a été abordée dans de nombreux articles scientifiques qui mettent l'accent, entre autres, sur le calcul d'une stratégie optimale de partage de charge entre différentes plateformes CDNs. Des acteurs jouant le role d'orchestrateurs et d'équilibreurs de charge (load balancer en anglais) sont également disponibles sur le marché des CDNs. Un orchestrateur fournit une interface transparente et unique aux fournisseurs de contenus ce qui leur permet de construire en ligne un CDN virtuel et global qui peut etre composé en réalité d'une mutitude de CDNs physiques. La sélection et l'approvisionnement des CDNs ainsi que le routage dynamique de requêtes entrantes entre ces CDN sont entièrement prises en charge par l'orchestrateur CDN dans une transparence totale en ce qui concerne le fournisseur de contenus. Contrairement à un orchestrateur CDN, un équilibreur de charge CDN intervient une fois qu'un fournisseur de contenus a signé des contrats de délégation avec un ou plusieurs fournisseurs de réseau CDN. L'équilibreur de charge est donc responsable du routage, en temps réel, des sessions des utilisateurs vers les plateformes CDNs. Les équilibreurs de charge bien connus incluent Limelight traffic load balancer, Dyn CDN, Conviva et Cedexis.
Une troisième forme de collaboration de fournisseurs de réseau CDN consiste en une fédération ou ces derniers regroupent leurs ressources respectives en termes de capacité et de couverture géographique et agissent en tant qu'un seul fournisseur de CDN global vis à vis des fournisseurs de contenus. La mise en place d'une fédération présente

de nombreux défis techniques et économiques. En effet, les acteurs de ma fédération doivent se mettre d'accord sur un modèle d'affaires commun de la fédération et sur une politique interne de partage du revenu global. En outre, un certain nombre de problèmes techniques liés au contrôle de la fédération doit etre adressé. Il s'agit notamment du placement des contenus, de la partage de charge et la gestion d'événements de pointe au sein d'une fédération établie en plus des questions de confidentialité entre fournisseurs de réseau CDN.
Avec la montée des Telco-CDN, certains constructeurs parmi les principaux acteurs du marché développent des solutions propriétaires pour permettre les opérateurs télecoms possèdant leurs propres réseaux CDN de fédérer. Il s'agit notamment, mais ne sont pas limités à, Cisco, Alcatel-Lucent et Ericsson. A notre connaissance, ces solutions sont encore dans la phase de conception et ne se sont pas encore traduites en une réalité sur le marché CDN.
Malgré son intéret chez les constructeurs informatiques et son attractivité pour les fournisseurs de réseau CDN et les fournisseurs de contenus, la fédération des fournisseurs de CDN n'a pas encore été la substance des efforts dévoués au sein de la communauté scientifique.

Objectifs de la thèse

Dans le cadre de cette thèse, nous abordons la fédération de fournisseurs de réseau CDN autonomes et distinctes, un fournisseur de réseau CDN étant n'importe quel acteur qui possède des ressources de stockage et de streaming libres au niveau réseau ou au niveau applicatif. Nous introduisons une solution de fédération qui permet de prendre deux types de décisions : les décisions statiques et les dynamiques. Les décisions statiques traitent des aspects de création/établissement et d'approvisionnement de fédérations. L'établissement d'une fédération consiste à identifier les fournisseurs de réseau CDN impliqués dans la fédération et le marché, en termes de fournisseurs de contenu conjointement ciblé par ces fournisseurs. L'établissement d'une fédération consiste également à mettre d'accord sur un modèle d'affaires commun de la fédération et sur une stratégie interne de partage des revenus. L'approvisionnement d'une fédération consiste à décider ou, au sein des différentes plateformes CDNs, les contenus des divers fournisseurs de contenus doivent etre placés et comment les futures demandes des utilisateurs doivent etre acheminés. Les décisions dynamiques visent à assurer un routage en temps réel des requêtes des utilisateurs vers les différentes plateformes CDNs de la fédération. Cette décision est basée sur les règles statiques de routage décidées lors de l'approvisionnement et des informations sur l'état, en temps réel, des différents CDNs.
Les décisions statiques sont effectuées sur une base mensuelle et journalière. les décisions dynamiques peuvent se produire à une échelle de temps qui est inférieur à une seconde. Les décisions statiques et dynamiques de contrôle de la fédération peuvent etre réalisées grace à un système de contrôle distribué consistant en une interaction de haut niveau entre les différents CDN. Ces décisions peuvent également etre effectuées grâce à une architecture de contrôle centralisé qui peut etre déployée par l'un des acteurs de la fé-

dération ou par un tiers indépendant. Afin d'alléger la complexité des fournisseurs de réseau CDN, nous supposons une architecture centralisée, basée sur un contrôleur permettant de prendre des décisions de mise en place et d'approvisionnement de fédérations ainsi que des décisions de contrôle dynamique.
Notre travail sur la fédération de fournisseurs de réseau CDN nous a conduit à mettre l'accent sur le role d'un opérateur Télécom à ce niveau. En particulier, nous étudions la valeur ajoutée qu'un opérateur télécom, en tant qu'opérateur de réseau et fournisseur d'un CDN last-mile, peut apporter à une fédération de fournisseurs de réseau CDN ou à des acteur 'Over The Top' indépendants. Nous centrons notre analyse sur le contexte mobile en raison des défis en termes de mobilité et de qualité de service que ce contexte présente.

Nous utilisons la théorie de l'optimisation afin d'aborder les aspects décisionnels statiques dans une fédération de CDN. Nous introduisons un modèle d'optimisation qui vise à établir et à approvisionner des fédérations à un rythme mensuel. Ce modèle vise à maximiser le gain conjoint de tous les acteurs de la fédération, tout en tenant compte des exigences de service des fournisseurs de contenu et des contraintes économiques et de capacité des différents fournisseurs de réseau CDN. Dans ce contexte, nous explicitons la notion d'équité économique au sein d'une fédération de fournisseurs de réseau CDN et nous garantissons qu'une fédération soit économiquement bénéfique à ces divers membres. Nous présentons également une variante du modèle d'optimisation qui permet de ré-approvisionner les fédérations sur une base quotidienne pour prendre en compte la variation de popularité de contenus existants ainsi que l'arrivée de nouveaux contenus.
Nous utilisons le modèle d'optimisation que nous avons introduit afin d'étudier trois scénarios de fédération d'intéret pour le marché CDN. Le premier scénarios est liée à la fédération de Telco-CDNs situés dans le même pays. Le deuxième scénario est lié à la fédération des pure-play CDNs avec des empreintes qui se chevauchent. Le troisième scénario est lié à la fédération de pure-play et de Telco CDNs. Pour chacun des scénarios, nous évaluons les gains économiques obtenus par les différents fournisseurs de réseau CDN par rapport au cas ou ils opèrent de manière séparée. Ces gains permettent de quantifier l'intéret économique de la fédération pour les différentes catégories de fournisseurs de CDNs (Telco Vs pure-play CDNs).

Nous abordons aussi l'aspect dynamique de contrôle de fédérations établies. En particulier, nous expliquons comment les résultats de la phase d'approvisionnement se traduisent par des décisions en temps réel de routage de requêtes au sein d'une fédération de fournisseurs de réseau CDN. Des explosions brusques de traffic sont susceptibles de se produire dans l'Internet. Ces événements influent directement sur les performances des plateformes CDNs. Dans le cadre du contrôle dynamique de fédérations, nous nous concentrons sur le contrôle des événements de pointe au sein d'une fédération de fournisseurs de réseau CDN. Différentes approches sont valables à ce niveau. Une première approche consiste à adapter la politique de routage statique calculée lors de l'approvisionnement d'une fédération. Une seconde approche consiste en un contrôle intra-CDN dans le sens ou chaque fournisseur de réseau CDN adapte la résolution des sessions livrées aux internautes en fonction de l'évolution de la charge de ces divers serveurs. Une dernière approche consiste à combiner les deux premières. Un certain nombre d'indicateurs se rapportant à la qualité d'expérience client sont utilisés pour évaluer les différentes approches. Cela nous permet de quantifier les gains de performance d'une approche de fédérée pour le contrôle des événements dynamiques par rapport à une approche intra-CDN.

L'étude du role d'un opérateur Télécom dans l'écosystème de distribution de contenu en général et en particulier dans un contexte de fédération de fournisseurs de réseau CDN nous permet d'identifier quelques services à valeur ajoutée que le Telco peut proposer dans ces contextes. Ces services incluent, entre autres, l'autorisation de l'accès des abonnés du Telco à des plateformes des tiers, le routage direct de requêtes utilisateurs vers les CDNs adéquat dans un contexte multi-CDNs et l'optimisation de la qualité de service de bout en bout entre les CDNs et les utilisateurs mobiles. Comme la mise en place de ces services nécessite un plan d'un contrôle adéquat, nous proposons des améliorations de l'infrastructure de contrôle existant de l'opérateur de télécommunication. Les améliorations que nous proposons prennent la forme de nouvelles APIs du réseau, de nouvelles entités de contrôle et des adaptations de certaines entités déjà existantes.

Contributions

Les contributions de cette thèse peuvent etre résumées comme suit :

1. **Introduire une architecture de contrôle centralisée permettant l'établissement, l'approvisionnement et le contrôle dynamique des fédérations de fournisseurs de réseau CDN**
   Nous présentons une architecture de contrôle centralisée mise en place par un contrôleur indépendant. Nous définissons les principaux composants de l'architecture, y compris les bases de données, interfaces et modules de prise de décision et les moteurs. Nous spécifions également les données qui doivent etre fournies par les fournisseurs de réseau CDN et par les fournisseurs de contenus au contrôleur à différentes échelles de temps. Une fois les données requises sont spécifiées, nous détaillons le fonctionnement des modules décisionnels statiques et dynamiques. En particulier, nous définissons les échelles de temps au cours de laquelle ces modules fonctionnent et les modèles mathématiques et des algorithmes sous-tendant leur fonctionnement respectif. Enfin, nous précisons les sorties des différents modules de prise de décision et expliquons comment ces sorties se traduisent par des décisions de gestion d'accords économiques, d'ingestion de contenus dans les CDNs et de routage des requêtes au sein de la fédération. Ces contributions sont détaillés dans les chapitres 3, 4 et 6 de la thèse.

2. **évaluer les gains économiques et de performance de la fédération**
   Les gains économiques de la fédération sont évalués par des enquêtes sur des cas

concrets de fédération de fournisseurs de réseau CDN. Pour chacun des cas d'utilisation considéré, nous évaluons les gains, en termes de chiffre d'affaires supplémentaire, réalisé par le fournisseur de réseau CDN concerné par rapport à un scénario non fédéré. Nous démontrons que, quand la demande des fournisseurs de contenus sur le marché des CDNs est élevé, les fournisseurs de réseau CDN ont intéret à fédérer. En particulier, certains fournisseurs de réseau CDN peuvent allers jusqu'à doubler leur chiffre d'affaires via la fédération. Nos contributions dans ce contexte peuvent etre trouvées dans le chapitre 5. Les gains en termes performance de la fédération sont évaluées en comparant les différentes approches de contrôle des événements de pointe au sein d'une fédération de fournisseurs de réseau CDN. Deux indicateurs clés sont utilisés dans ce contexte : le nombre de sessions rejetées et la résolution vidéo moyenne, en Mbps, ressentie par les utilisateurs . Nous démontrons que, en adoptant un comportement fédéré en heures de pointe, les membres d'une fédération peuvent améliorer leur résilience à ces événements. Cela se traduit par un taux d'acceptation de requêtes plus élevé de la fédération et une meilleure résolution vidéo ressentie par les utilisateurs finaux. Nos contributions à ce niveau peuvent etre trouvés dans les sections 6.4 et 6.5.

3. **Proposer des améliorations du plan de contrôle des opérateurs Telecom afin de permettre un nouveau positionnement de ces acteurs dans l'écosystème**
   Suite à une analyse approfondie des atouts du Telco, de son plan de contrôle et de son positionnement général dans l'écosystème, nous proposons des modifications de l'infrastructure de contrôle existant du Telco afin de permettre un certain nombre de services à valeur ajoutée à base de Telco. Il s'agit notamment de l'autorisation de l'accès des abonnés Telcos pour le compte des OTTs, l'acheminement direct des demandes des utilisateurs vers le bon CDN dans un contexte multi-CDNs et une optimisation de bout en bout la qualité de service entre CDN et les abonnés mobiles. Au-delà de la définition des services à valeur ajoutée, nos contributions à ce niveau sont de deux ordres. Premièrement, nous définissons des interfaces réseaux (API) qui permettent une simple souscription des OTTs à ces services. Deuxièmement, nous proposons de nouvelles entités de contrôle et des modifications des entités existantes pour permettre les services que nous proposons. Nous montrons comment les services proposés sont activés par un cas d'utilisation concret. Ces contributions sont détaillées dans le chapitre 7.

Plan de la thèse

La thèse est structurée comme suit.
Dans le chapitre 2, nous effectuons l'état de l'art sur les services de distribution de contenu en général et sur les réseaux de distribution de contenu (CDN) en particulier. Nous analysons les principales tendances d'évolution ayant lieu au à la bordure de l'internet et à l'intérieur du réseau et utilisons ces tendances pour la prévision de l'évolution à court et moyen terme de l'écosystème des services de distribution de contenu. Nous

positionnons le sujet principal de la thèse, qui est la fédération des fournisseurs de réseau CDN autonomes, dans ce contexte.
Dans le chapitre 3, nous introduisons une architecture de contrôle centralisée qui permet l'établissement, l'approvisionnement et le contrôle de fédérations de fournisseurs de réseau CDNs. Nous explicitons les principaux composants de l'architecture, y compris les bases de données, les interfaces et les moteurs de prise de décision ainsi que son fonctionnement global.
Dans le chapitre 4, nous abordons l'aspect statique de mise en place de la fédération et de l'approvisionnement dans une perspective fondée sur la théorie de l'optimisation. Nous introduisons un modèle d'optimisation qui permet l'établissement et l'approvisionnement de fédérations de fournisseurs de réseau CDN sur une base mensuelle. Les exigences service des fournisseurs de contenus sont pris en compte par les contraintes de service du modèle (section 4.3). Les limites en termes de capacité des CDNs appartenant à divers membres de la fédération sont également pris en compte dans le modèle (section 4.4). Dans les sections 4.5 et 4.6, nous introduisons un modèle de partage des revenus qui tient compte des notions d'équité et de rationalité. Dans la section 4.8, nous utilisons une variante de ce modèle pour ré-approvisionner les fédérations établies à la lumière des changements d'exigences des fournisseurs de contenus. Nous explicites les sorties du modèle dans la section 4.8 et discutons de l'efficacité d'une approche centralisée pour le contrôle de fédérations dans la section 4.9.
Dans le chapitre 5, nous utilisons des traces de trafic réelles afin de générer des entrées pour le modèle d'optimisation présenté dans le chapitre 3. Nous appliquons ensuite ce modèle pour trois cas concrets de fédération d'intéret pour le marché. Le premier cas traite de la fédération de Telcos-CDN co-localisés dans le même pays. Le deuxième cas traite de la fédération de pure-play CDNs avec des empreintes qui se chevauchent. Le troisième cas traite de la fédération des Telco-CDNs et de pure-play CDNs. Pour chacun des cas d'utilisation, nous évaluons les gains réalisés par les fournisseurs de réseau CDN via la fédération.
Dans le chapitre 6, nous explicitons le processus de routage de requêtes utilisateurs dans la section 6.1. Dans la section 6.2, nous nous concentrons sur le contrôle des événements de pointe au sein d'une fédération de fournisseur de réseau CDN. Nous explicite l'origine des événements de pointe et introduisons différentes approches de contrôle pour faire face à ces événements. Nous évaluons ensuite les performances des différentes approches en effectuant des simulations basées sur des traces de trafic réelles et nous utilisons un certain nombre d'indicateurs clés de performance.
Dans le chapitre 7, nous abordons le positionnement Telco dans l'écosystème CDN. Dans la section 7.2, nous présentons trois services à valeur ajoutée qui peuvent etre proposées par l'opérateur de télécommunication dans un contexte de fédération de fournisseur de réseau CDN ou individuellement à des OTTs. Dans la section 7.3, nous décrivons le plan de contrôle existant de l'opérateur de télécommunication. Dans la section 7.4, nous proposons des modifications de l'infrastructure de contrôle Telco afin de permettre d'implémenter les services à valeur ajoutée que nous proposons. Les modifications prennent la forme d'APIs de réseau, de nouvelles entités de contrôle et d'adaptations des fonction-

nalités des entités existantes. Nous illustrations via un cas d'utilisation concret comment les services à valeur ajoutée sont activés dans la section 7.5.
Nous concluons la thèse dans le chapitre 8 et donnons un aperçu des futurs travaux de recherche.

List of Figures

List of Tables

Contents

Chapter 1

Introduction

1.1 Context and Background

A decade ago, people mainly relied on conversational services including phone calls and video conferences in order to communicate and stay in touch with each others. The rise of user-generated content (UGC) portals including YouTube [127] and DailyMotion [55] and the spread of social media platforms including Facebook [63] and Twitter [121] contributed to a major shift in users' attitudes and behaviors. Consulting online blogs, downloading and uploading videos from/to YouTube, posting on Facebook and tweeting have became a significant component of our every-day life. Most importantly, these almost-intuitive actions, resolving around the consumption, generation and sharing of contents that can be web pages, videos, albums or online posts, have became our way to stay connected with the rest of the world.

Today's internet speaks in terms of content consumption, generation and sharing. This trend is expected to accelerate in the coming years. Indeed, video traffic alone, including Video on Demand (VoD) [122], live streaming [115] and catch-up TV [79], is expected to account for more than 90 % of global IP traffic in 2017 [19] [23]. An over increasing amount of content-related traffic is expected to be generated by non-PC devices including mobile phones and Tablets [19] [20]. With the rise of concepts like Machine to Machine (M2M) [128] and Internet of Things (IOT) [41], connected objects are also expected to be at the origin of new types of content with, probably, a different structure than the contents that we know today (videos, web pages, games etc). In parrallel, consumers are becoming less tolerants to high start times and video freezes and more demanding in terms of video quality or resolution [57]. These aspects of evolution of the internet edge have direct impacts on end to end stakeholders involved in the content provisioning value chain.

Stakeholders may be classified into four main categories. At the top of the chain, are the Content producers and providers (CPs) that directly or indirectly deliver producer-

based or user-generated contents to a large audience. Well-known CPs include BBC [43], YouTube [127], Daily Motion [55] and Netflix [93]. The second category is composed of Content Delivery Network (CDN) providers, a CDN being composed of clusters of edge servers distributed in one or many geographic Points of Presence (PoPs). When delegating their content delivery to third-party CDNs, CPs aim at reducing their capital expenditures (Capex), alleviating their servers load and enhancing the quality of experience (QoE) perceived by end users. CDNs deliver CPs content from the internet edge and use proprietary content acceleration techniques [101]. Well-known CDN providers include Akamai [32] that owns 80% of the CDN market and delivers more than 25% of the internet traffic [64], Level3 [80] and Limelight [82]. We refer to these CDN providers as pure-play CDNs. The third category of stakeholders is formed by Network Operators including users and transit internet service providers (ISPs) [125] which allow content to be routed to/from end users. At the end of the chain, come end users who are at once content consumers and producers.
It is important to note that the edge between these categories is not always clear. For instance, YouTube and Netflix are both CPs that deployed their own CDNs which are dedicated for their internal use [8] [4]. In addition, many Telecom operators (Telcos) are entering the CDN market through implementing their own last-mile CDNs [2] in the countries of their footprint. We refer to Telcos owning a CDN platform as Telco-CDNs.

Challenges faced by stakeholders are of technical and business orders. Traffic explosion is leading to a congestion not only inside network operators domains but also at the level of interconnection links [40]. This congestion is increasing cross traffic costs while degradating network quality of service (QoS). Content producers and providers (CPs) aim at finding a trade-off between building their own distributed, last-mile CDNs and delegating their content distribution to third parties, both scenarios having business implications in terms of Capex investments and Service Level Agreements (SLAs) establishment. Despite their large and distributed clusters, pure-play CDNs like Akamai and Limelight face capacity limitations in some types of peak events. A well-known example of peak events includes Flash Crowds [76] scenarios referring to a rapid and unexpected surge of a given CP traffic. Futhermore, for QoE related issues, pure-play CDNs are interested in locating their content closer to end users [99]. Meanwhile, despite their assets in terms of user proximity and network control, Telco-CDNs often suffer from the local nature of their footprint [53]. Based on the described status-quo, content distribution services are moving from a centralized scheme where content is monopolized by a single player (this player is the CP or a global pure-play CDN) towards a new scheme involving distributed and autonomous content players. This trend is being accelerated by the emergence of new players, including Telcos, in both the CDN and Cloud markets [51] as well as by the advancement of virtualisation and resources sharing techniques [47]. In summary, any player that owns vacant in-network or overlay storage and streaming resources can play the role of a CDN provider and, as a consequence, become a part of a global CDN platform.

1.2 Collaboration of CDNs

Given the high market demand, CDN providers, whether they are well-positioned or emergent, have incentives to collaborate. QoE enhacement, capacity aggregation and footprint extension fall under the technical incentives. Revenue enhancement and resources optimization fall under the economic ones. CDN providers collaboration may occur in a centralized or distributed fashion.

One form of CDN collaboration consists of point to point, bilateral SLAs between two CDN providers. This form of collaboration has been investigated by the IETF [72] CDN interconnection (CDNI) working group [45] [42].
Another form of CDN collaboration consists of orchestrating a given CP traffic among different CDN providers, each with its own business model and not necessarily aware of each others. Orchestration may be done by the CP itself, by an upstream or pure-play CDN or by a third party orchestrator or broker. CDN orchestration has been addressed in many scientific papers that focus, among others, on computing an optimal strategy of load sharing among different CDNs [75] [83]. CDN orchestrators and load balancers are also commercially available in the CDN market. CDN orchestrators provide a transparent and unique interface to CPs for building a virtual and global CDN that is based on a mutitude of physical CDNs. The selection and provisioning of physical CDNs as well as the dynamic routing of incoming requests among these CDNs are entirely handled by the CDN orchestrator in a complete transparency with regards to the CP. Contrarily to a CDN orchestrator, a CDN load balancer intervenes after CPs have established SLAs with many CDNs in order to perform a real-time routing of users' sessions towards these CDNs. Well-known load balancers include Limelight traffic load balancer [94], Dyn CDN manager [59], Conviva [54] and Cedexis [52].
A third form of CDN collaboration consists of a federation of CDN providers that aggregate their respective assets in terms of capacity and geographic footprint and act as a global, unique CDN with regards to CPs. Putting in place a federation is challenging from both technical and economic perspectives. Indeed, the CDN providers in a federation should agree on a common business model of the federation and on an inner policy of revenue sharing. Furthermore, a number of technical issues related to the federation control should be addressed. These include content placement, load balancing and events management within the CDN federation in addition to inter-CDNs privacy issues. With the rise of Telco-CDNs, some IT manufacturers among the key players in the market are developing proprietary solutions for allowing Telco-CDNs to federate. These include, but are not limited to, Cisco [78], Alcatel-Lucent [12] and Ericsson [11]. Up to our knowledge, these solutions are still in the conception phase and have not translated into a reality in the CDN market yet.

Despite its interest among IT manufacturers and its attractiveness for both CDN providers and CPs, the federation of CDN providers has not yet been the substance of dedicated efforts in the research community.

1.3 Thesis Objectives

In the context of this thesis, we address the federation of autonomous and distinct CDN providers, a CDN provider being a player that owns vacant overlay or in-network resources at one or many geographic PoPs. We introduce a federation solution that allows taking two kinds of decision: Static decisions and dynamic ones. Static decisions deal with aspects of federation establishment and provisioning. Establishing a federation consists of identifying the CDN providers involved in the federation and the market, in terms of CPs, jointly targeted by this federation. It also consists of agreeing on a common business model of the federation and on an inner strategy of revenue sharing. Provisioning a federation consists of deciding how, inside a CDN federation, CPs contents should be placed and future users' requests should be routed. Dynamic decisions aim at ensuring a real-time routing of users' requests towards the different CDNs of the federation. Request routing is based on the provisioning phase output and on real-time information of different CDNs state. While static decisions are performed on monthly and daily basis, dynamic decisions can occur at a time scale that is less than one second. Static and dynamic decisions of federation control can be performed through a distributed control scheme consisting in a high level interaction between different CDNs. These decisions can also be performed through a centralized control architecture that can be deployed by one of the federation players or by an independent third party. In order to alleviate CDNs complexity, we assume a centralized, controller-based architecture that allows taking decisions of federation establishment, federation and dynamic control.
Our work on CDN federation lead us to focus on the Telco role. In particular, we investigate the added-value that the Telco, as a last-mile network operator and CDN provider, can bring to a federation of CDNs or to individual Over The Top (OTT) players. We center our analysis on the mobile context due to the many mobility and QoS related challenges that it presents.

We use optimization theory [13] [14] in order to address static decision-making aspects in a federation of CDNs. We introduce an optimization model that allows federation computation and provisioning on a monthly basis. This model aims to maximize the joint gain of the federation while taking into account CPs service requirements and CDN providers economic and capacity constraints. In this context, we explicit the notion of economic fairness within a federation of CDNs. We equally introduce a variant of the optimization model that allows federation re-provisioning on a daily basis.
We use the optimization model that we introduced in order to investigate three use cases of federation of interest for the CDN market. The first use case is related to the federation of Telo-CDNs located in the same country. The second use case is related to the federation of pure-play CDNs with overlapping footprints. The third use case is related to the federation of pure-play CDNs and Telco-CDNs. For each of the use cases, we assess the economic gains achieved by different CDN providers through moving from a separate scenario to a federation. These gains allows quantifying the federation eco-

nomic interest for different categories of CDNs.

We equally address the dynamic aspect of CDN federation control. In particular, we explain how the provisioning phase outputs translate into real-time decisions of request routing within a federation of CDNs. Unexpected traffic surges are likely to happen in the Internet [40]. These events directly impact the performance of CDNs [76]. As part of dynamic federation control, we focus on the control of peak events within a federation of CDNs. We propose different frameworks in this context. While some frameworks privilege an intra-CDN form of control, others advocate a federated control involving different members of the federation. We conduct trace-driven simulations in order to assess the performance of different frameworks. A number of key indicators that relate to users' QoE are used in this context. This allows us to quantify the performance gains of a federated approach for dynamic events control.

When investigating the Telco role in the CDN ecosystem and particularly in a CDN federation context, we identify major added-value services that can be proposed by the Telco to a federation of CDNs, CPs and pure-play CDNs. Since enabling added-value services requires an adequate control plane, we suggest enhancements of the existing control infrastructure of the Telco. The enhancements that we propose take the form of new network APIs, new control entities and adaptations of some already-existing entities.

1.4 Contributions

The contributions of this thesis can be summarized as follows:

1. **Introduce a centralized control architecture that allows CDN federations computation, provisioning and dynamic control**
   We introduce a centralized control architecture implemented by an independent controller. We define the main components of the architecture including databases, interfaces and decision-making modules and engines. We also specify the inputs provided by the CDNs and by the CPs to the controller at different time scales. Once required inputs are specified, we detail the operation of the static and dynamic decision-making modules. In particular, we define the time scales at which these modules operate and the mathematical models and algorithms underlying their respective operation. Finally, we specify the outputs of different decision-making modules and explain how these outputs translate into decisions of SLAs management, content ingestion and request routing to be enforced at the federation level. These contributions are detailed in Chapters 3, 4 and 6 of the thesis.

2. **Assess the economic and performance gains of the federation**
   The economic gains of the federation are assessed through investigating concrete use cases of federation related to the federation of different categories of CDN players. For each of the considered use cases, we assess the gains, in terms of

extra revenue, achieved by involved CDNs through moving from a separate scenario to a federation. We demonstrate that some CDNs can double their revenue through federating. Our contributions in this context can be found in Chapter 5. The performance gains of the federation are assessed through comparing different frameworks for events control within a federation of CDNs. Two key indicators are used in this context: the number of sessions rejected by the CDNs and mean video resolution experienced by end users. We demonstrate that, through adopting a federated behavior upon peak events, the CDNs of a federation can enhance their resilience to these events. This translates into a higher joint hit ratio of the federation and a better video resolution witnessed by end users. Our contributions at this level can be found in sections 6.4 and 6.5.

3. **Suggest enhancements of the Telco control plane in order to allow a new positioning of the Telco in the CDN ecosystem**
   Based on an extensive analysis of Telco assests, the Telco control plane and the Telco positioning in the CDN ecosystem, we suggest modifications of the existing control infrastructure of the Telco in order to allow a number of Telco-based added-value services. These include user authorization on behalf of OTTs, direct routing of users' requests to CDNs in a multi-CDNs context and an optimization of end to end QoS between CDNs and mobile subscribers. Beyond defining the added-value services, our contributions at this level are twofold. First, we define APIs that allow a simple subscription of OTTs to these services. Second, we suggest new control entities and modifications of existing ones for enabling the services that we propose. We show how the proposed services are enabled through a concrete use case. These contributions are detailed in Chapter 7.

1.5 Overview of the dissertation

The dissertation is structured as follows.
In Chapter 2, we describe the status-quo of content distribution services in general and of content delivery networks (CDN) in particular. We analyze the main evolution trends occuring at the internet edge and inside the network and use these trends for forecasting the short to mid term evolution of the ecosystem of content distribution services. We position the main topic of the thesis, that is the federation of autonomous CDN providers, in this context.
In Chapter 3, we introduce a centralized control architecture that allows federation computation, provisioning and control. We explicit the main components of the architecture including databases, interfaces and decision-making engines as well as its overall operation.
In Chapter 4, we address the static aspect of federation establishment and provisioning from an optimization theory-based perspective. We introduce an optimization model that allows federation establishment and provisioning on a monthly basis. CPs service and performance requirements are taken into account through the model service con-

straints (section 4.3). The capacity limitations of CDNs are taken into account through the model capacity constraints (section 4.4). In sections 4.5 and 4.6, we introduce a revenue sharing model that takes into account the rationality and fairness notions. In section 4.8, we use a variant of this model for re-provisioning established federations in the light of changes of CPs requirements. We explicit the model outputs in section 4.8 and discuss the scalability of a centralized approach for federation control in section 4.9.
In Chapter 5, we use real traffic datasets in order to generate inputs for the optimization model introduced in Chapter 3. We then apply this model to the instantiations of three use cases of federation. The first use case addresses the federation of Telco-CDNs colocated in the same country. The second use case addresses the federation of pure-play CDNs with overlapping footprints. The third use case addresses the federation of Telcos and pure-play CDNs. For each of the use cases, we assess the gains achieved by CDN providers through moving from a separate scenario to a federation.
In Chapter 6, we explicit the request routing process in section 6.1. In section 6.2, we focus on the control of peak events within a federation of CDNs. We explicit the origin of peak events and introduce different control frameworks for dealing with these events. We then assess the performance of different frameworks through conducting trace-driven simulations and using a number of key performance indicators.
In Chapter 7, we address the Telco positioning in the CDN ecosystem. In section 7.2, we introduce three added-value services that can be proposed by the Telco in a CDN federation context or to individual OTTs. In section 7.3, we describe the existing control plane of the Telco. In section 7.4, we suggest modifications of the Telco control infrastructure in order to enable the added-value services that we propose. Modifications take the form of network APIs, new control entities and adaptations of the functionalities of existing entities. We use a concrete use case in order to illustrate how added-value services are enabled in section 7.5.
We conclude the thesis in Chapter 8 and give an overview of future research work.

Chapter 2

State of the art of Content Distribution Services

This chapter is dedicated to the description of the Status-quo and the evolution of content distribution services. We start by giving an overview of the content provisioning value chain. We then focus on content delivery networks (CDNs). In particular, we explicit the topology and the inner operation of CDNs. We also give an overview of the CDN market and the CDN pricing schemes. Later, we describe the main evolution trends witnessed both at the internet edge and inside the network. These can reflect a change of users' behaviors or can be due to the rise of new technologies and storage/routing paradigms. In the light of these trends, we forecast the short to mid term evolution of content distribution services. The contributions of this thesis are positioned in this context.

2.1 Overview of the Content provisioning value chain

The value chain of content provisioning is, in a simplified view, composed of three main families of players: content producers and providers (professionals, e.g. BBC [43], YouTube [127], Netflix [93]...), content distribution service providers (Akamai [32], Level3 [80]...) and network service providers.

A CP owns a proprietary infrastructure that, in addition to content and metadata management (including popularity estimation) and storage, enables customer related functionality like authentication, authorization, accounting, billing, profiling, content adaptation. Content popularity assessment and users profiling facilitates an optimized usage of resources. CPs also may offer some value-added services to end users such as portals personalization per end user and provision of intelligent tools for browsing the content catalog.

As the volume of contents and requests exponentially grows, CPs should find a trade-off between investing in their infrastructure thus increasing their costs, and delegating content delivery to third parties. Since the datacenters of most CPs are centralized, the long distance to end users often induces a degradated QoE due to high routing delays

Figure 2.1: Overview of the Content Provisioning Value Chain

and eventual traffic jams on the internet.
Scalable and efficient content delivery requires CDNs. CDNs consist in clusters of servers distributed in one or many geographic locations or PoPs and connected to ISPs networks through Internet exchange points [101] [100]. There is a moving equilibrium between CPs and CDN providers. In fact, CPs usually delegate their content delivery (at least for the most popular content) to CDN providers but this is not always the case. YouTube and Netflix are well-known exceptions [4] [8]. The CDNs deployed by CPs are usually destined for proprietary use and are not open to third parties. Some CPs may rely on more than one CDN in order to ensure a better quality of experience and master costs (before implementing its own CDN, Netflix had agreements with three CDNs [33]). However, CPs may lack good control solutions to facilitate the management of the overall (composed) system.
CDNs, the systems on which content distribution rely, can be classified in two main families: those that use the Internet for the connectivity requirements, denoted by pure-play CDNs, and those based on end-to-end managed networks, denoted by Telco CDNs.
A CDN performs two essential functions. First, it caches content at the edge of the network, closer to end users, to reduce the IP traffic traversing the core network and, ideally, deliver a higher quality experience to viewers. Second, a CDN positions multiservice, multiprotocol content delivery capabilities at the network edge, allowing a dynamic delivery of non-linear and linear contents to a multitude of devices and end users [101] [100]
CDNs rely on proprietary solutions in terms of caching algorithms and web acceleration techniques. Transparency from the end user point of view is achieved through the usage of the fundamental web protocols and general architecture: DNS for content names resolution (content URL) into IP addresses, HTTP [65] for content exchange (GET/POST) between users and CDNs or CPs servers and RTP/RTSP [112] [113] for managing live streaming contents. The core functionality leading to the targeted transparency is redirection: when CPs delegate content distribution to a CDN player, redirection towards the CDN is configured either at DNS [91] [60] or HTTP levels. The overall content delivery solution is also based on W3C components, including HTML, and its evolution toward HTML5, the related enhancement of browsers, as well as the evolution of the

WebRTC approach [124]. Beyond content distribution and delivery, CDNs implement, in a proprietary manner, limited clients management and session control functions that are required prior to content delivery (Authentication, Authorization, Billing...).
Network service providers include Tier1, Tier2 and Tier3 ISPs [15]. Located at the center of the internet, Tier1 ISPs allow CPs contents to be reachable by end users. Tier3 and Tier2 ISPs provide local and regional connectivity to end users through enabling their access to Tier1 networks. Tier2 ISPs pay transit costs to Tier 1 ISPs. Tier3 ISPs, denoted by users' ISPs, pay transit costs to Tier2 and Tier1 ISPs [15] [125]. As Tier3 ISPs, Telcos are usually paid monthly subscriptions by end users for providing them with a quota-based or an unlimited data access. Charging at this level can occur in online or offline modes. Telcos may also be indirectly paid by CPs for granting an unlimited access to these CPs portals to a part or the entirety of their subscribers. This service is denoted by Sponsored data connectivity [18].

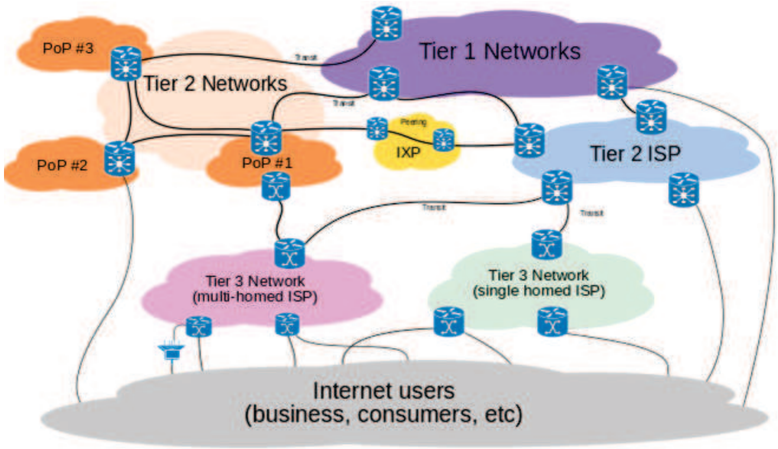

Figure 2.2: Overview of the Internet Players

2.2 Overview of Content Delivery Networks (CDN)

In this section, we describe in details the topology and operation of CDNs. We then give an overview of the CDN market and of the pricing schemes that are widely adopted within the CDN market.

2.2.1 Overview of CDN Topology

CDNs maitain mutiple PoPs with clusters of servers that store copies of (sometimes) identical content, such that users' requests are satisfied by the most appropriate cluster. Figure 2.3 shows an example of a CDN topology. A CDN topology consists of [101]:

- Centralized content servers, also named origin servers, storing the entirety of the content catalogs of the CDN clients (CPs delegating their content delivery to the CDN).

- A set of surrogates or edge servers distributed, in the form of clusters, in one or many geographic locations around the world. CDN clusters are connected, through internet peering points, to regional or local Internet Service Providers (ISPs) [64]. In case of Telco-CDNs, the surrogate servers are connected, in an overlay fashion, to the regional nodes of the Telco backbone.

- Routers and network elements within clusters of surrogates and between clusters and ISPs. These elements allow requests to be routed at an intra CDN level and content to be routed from surrogates to end users.

- Centralized, control level entities handling dynamic and static aspects of CDN control. These include accounting, configuration and request routing modules. The exact functioning of these modules is detailed in the next section.

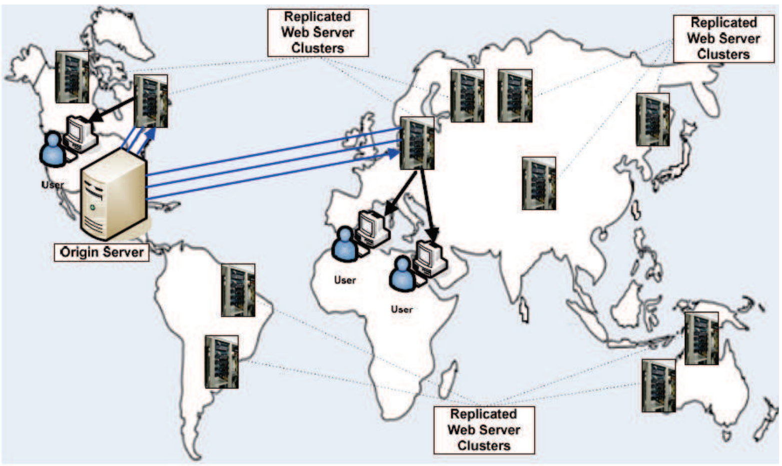

Figure 2.3: Overview of CDN topology

2.2.2 Overview of CDN Operation

The CDN operation consists of a set of functionalities perfomed at different time scales by the CDN. These functionalities are enabled through a centralized control plane where the CDN intelligence is concentrated. Next, we detail the main functionalities supported by the CDN control plane.

Content Ingestion and Outsourcing

Content ingestion takes place after a CDN establishes a Service Level Agreement (SLA) with a Content Provider (CP). It consists of pushing CP contents from the CP domain towards the origin servers of the CDN.
Content outsourcing consists of placing CP contents in the CDN surrogates. The goal is to make these contents as close as possible from requesting end users. Two main outsourcing practices are adopted by CDN providers:

- Push-based Outsourcing: Contents are proactively pushed from the origin servers to the CDN surrogates prior to users' requests arrival. The choice of the contents that should be pushed towards different clusters depends of the content selection strategy adopted within the CDN. One strategy consists in placing all contents in all of the CDN clusters. A more optimal strategy consists in placing contents in the clusters located in the geographic areas where they are the most popular. Once contents are pushed towards a given cluster, the surrogates within this cluster cooperate together in order to reduce content replication and update costs. Several replication strategies have been studied in the litterature [67] [77] [107]. In this context, it is noted that greedy-global heuristic algorithms are the best choice in making the replication decisions between cooperating surrogate servers [77]. After contents are placed in different surrogates, an incoming user's request is directed to the closest surrogate that holds the content. If no close surrogate holds the requested content, the request is directed to the origin server.

- Pull-based Outsourcing: Contrarily to Push-based Outsourcing, contents are not proactively placed in the CDN surrogates. Instead, surrogate servers, which are intially empty, are dynamically filled with contents that are being requested by end users. Each user's request is directed to the nearest surrogate server. If a cache miss occurs (content not in server), this server pulls the content from another surrogate server that holds the content and that can be located in the same cluster or in a higher level cluster. In some cases, the content is pulled from the origin.

Well-known CDN providers including Akamai [32] and Mirror Image [90] use the pull-based approach for outsourcing CPs contents towards their surrogates. Even though this approach presents some complexity with regards to the push-based approacg, it has the advantage of ensuring a better responsiveness to the evolution of contents popularity over the different geographic areas of the CDN footprint.

Intelligent Request Routing

Request routing within a CDN consists of deciding, in real-time, the surrogate server towards which an incoming request should be routed. Request routing is performed by a dedicated module, the request router, which uses information on users' geolocation in order to decide the surrogate towards which an incoming request should be routed. The module can also use information about surrogates load and network state in order to perform request routing.
In most cases, the Request Router routes an incoming request towards the surrogate server that is the closest to the requesting end user. Nevertheless, if the concerned server is highly loaded and/or there is a problem on the network path between this server and the end user, the request is directed to another surrogate server that is chosen based on proximity and availability criteria. Typically, the request is routed towards another surrogate server that is close to the requesting user, is not highly loaded and that can be easily accessed through the network. The request routing functionality is ensured through the use of HTTP [65] and/or DNS [91] redirection mechanisms.

Let us consider the use case of an Orange subscriber that has requested an Apple content and has been redirected to Akamai. Figure 2.4 illustrates a DNS-based approach for request routing within Akamai. Figure 2.5 illustrates a HTTP-based approach for request routing within Akamai.

Reporting, Analytics and Billing

When intercepting users' requests, a CDN surrogate maintains logs concerning the identities of requesting users and the popularity of requested contents. Access and popularity logs are provided, on a daily basis, by different CDN surrogates to a dedicated control entity that is responsible of logs aggregation and analysis. The analysis of aggregated logs allows issuing monthly bills corresponding to prices charged by the CDN to each of its clients. Furthermore, it allows gathering statistics on the traffic of CDN clients. Statistics include the monthly volume of traffic generated by a given CP, traffic distribution across different geographic areas, the time variation of traffic distribution, daily popularity of different contents... In addition to CPs accounting, these statistics can be used in order to adapt the CDN topology in terms of placement and number of surrogate servers. For instance, a CDN may decide to place more surrogate servers in a geographic area where the traffic of one of its clients is significantly increasing. If a push-based approach for content outsourcing is adopted, statistics can also be used in order to update the strategy of content distribution referring to the way contents are assigned to the CDN clusters.

Servers Control

Servers control consists of taking decisions of surrogates placement, dimensioning and configuration upon the deployment of the CDN. It also consists in updating these decisions in the light of changes of CPs requirements.
The decision of surrogates placement and dimensioning is made by a dedicated control module based on a prior knowledge of the future market of the CDN and of the requirements of potential clients. Requirements include the geographic areas where the demand of the CDN clients originate, the volume of the clients demand, demand distribution across geographic areas, the time variation of clients traffic... If the demand of the CDN clients evolve or the CDN wants to expand its market through targeting new CPs, new decisions of surrogates placement and dimensioning are made by the module. The optimal placement of surrogate servers within a given CDN domain has been the subject of several studies in the litterature. In particular, [56], [106], [103], [35] introduce different placement strategies based on different criteria including the Capex of the CDN and the distance to end users. It is important to note that the placement and dimensioning of a given CDN highly depends of the category of users targeted by this CDN (fixed, mobile users, both...).
Servers configuration consists of deciding the caching policy used in the CDN surrogates referring to the rules underlying contents storage in these surrogates. Servers configuration is particularly important when the Pull-based approach for outsourcing is adopted

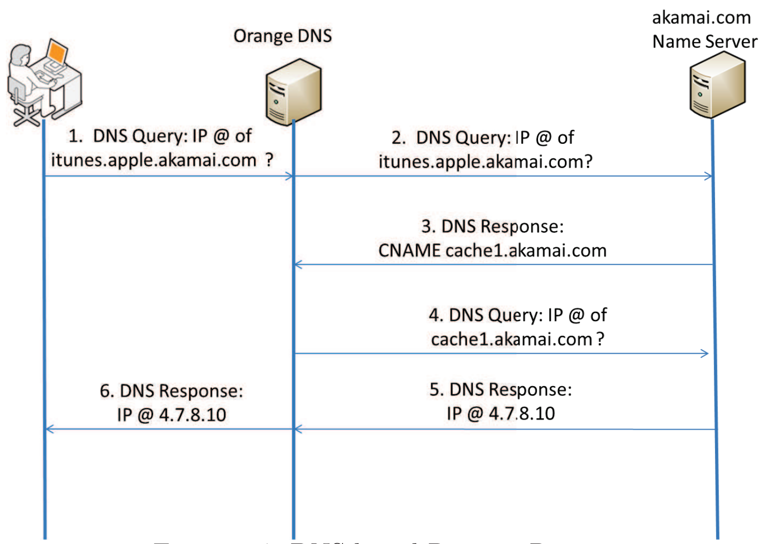

Figure 2.4: DNS-based Request Routing

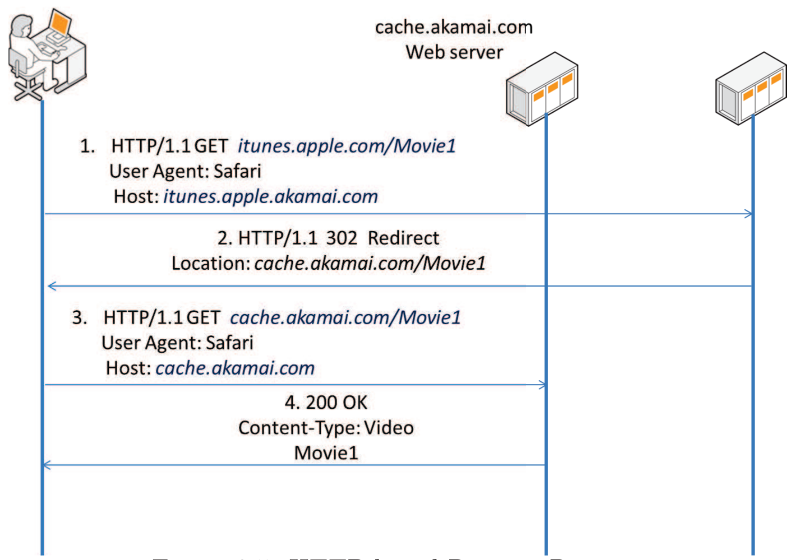

Figure 2.5: HTTP-based Request Routing

within the CDN. Widely deployed caching algorithms include the Least Recently Used (LRU) and Least Frequently Used (LFU) algorithms but other algorithms have also been investigated in the litterature [105] [120]. Servers configuration is made by a dedicated configuration module.
Servers configuration also consists of purging CDN servers upon need. Namely, when a CDN is notified of the invalidity of a content stored at its level (obsolete or illegal content), the configuration module operates immediately in order to remove this content from all the servers where it is present.

Routing Optimization

Routing optimization aims at accelerating content delivery from the CDN to end users as well as at optimizing bandwidth usage at an intra-CDN level. Many mechanisms for routing optimization are used by different CDN providers. Next, we list some of the mechanisms used by Akamai [117] [101]:

- Hierarchical Routing: The CDN topology consists of a hierarchical structure with clusters connected to Tier3 ISPs being local clusters, clusters connected to Tier2 ISP being Regional clusters and clusters connected to Tier1 ISPs being international clusters. An incoming request is first directed to a local cluster. If the content is not found in the local cluster, the closest regional cluster is solicited and so on. Eventually, if none of the clusters in the hierarchical structure of the CDN holds the content, the content is pulled from the origin server. Hierarchical routing allows accelerating content delivery in the sense that each request is fulfilled by the closest cluster that holds the requested content. In addition, when content is pulled from a non-local cluster or from the origin sever, it is stored in all clusters along the way thus shortening the path crossed by the content upon future requests.

- Requests Aggregation and Responses Duplication: This mechanism is applicable when many requests for the same content are intercepted in a quasi-simultaneous fashion by a local cluster and when the requested content is not present at this cluster level. In this context, the hierarchical process of requests routing is performed. Instead of performing hierarchical routing for each of the requests, all requests for the same content are aggregated and the content is pulled once from a regional or international cluster or from the origin server. When the content arrives at the local cluster level, it is duplicated and delivered in a unicast mode to each of the requesting users. Requests aggregation has the advantage of preventing the same content from being routed many times at an intra-CDN level (between the CDN clusters) thus optimizing bandwidth use within the CDN.

- TCP Optimization: When a content is being pulled by a CDN server from another server, a TCP connection is established between the two. TCP optimization consists in aggregating TCP connections corresponding to different, quasi-simultaneous requests into one TCP session in order to reduce content delivery time and to optimize bandwidth usage at an intra-CDN level. Let us assume that many requests are intercepted almost at the same time by a given surrogate server. If a cache miss occurs for more than one request, the surrogate is going to pull *missed* contents from other surrogate servers or eventually from the origin server. If contents corresponding to different requests are pulled from the same server, a single TCP connection is established between the server pulling the contents and the selected server. Since the initiation of a TCP session generates delays (three way handshake [10]), aggregating many requests into one TCP session presents gains with regards to the case where each request is mapped to a separate session. Gains increase with the number of aggregared requests. Furthermore, TCP optimization allows placing different contents in the body of a single TCP message thus eliminating unnecessary traffic due to TCP headers [10].

- Content Prefetching: Web pages often point to many URLs referring to other web pages or other types of content such as video. When requesting a web page, an end user is likely to later request one or many of the objects that are embedded in this page. When a CDN surrogate intercepts a request for a web page that is not stored at its level, it initiates a TCP connection with another surrogate server or with the origin server in order to fetch the missed web page. The surrogate server also identifies the contents embedded in the web page and pre-fetches these contents from the selected surrogate/origin server. URLs pre-fetching allows accelerating content delivery through pushing contents closer to end users prior to the time when they are requested by these users. Second, it allows using a single TCP session for fetching many contents thus reducing the download time and optimizing the use of bandwidth resources at an intra-CDN level.

2.2.3 Overview of the CDN Market

Since first commercial CDNs were launched in 1996 [58], the CDN market has witnessed an exponential growth worldwide [58]. The evolution, since 2010, of the CDN market is shown in Figure 2.6. As can be noticed, CDN providers are mostly present in North America, Western Europe and the Asia Pacific region. Today, the CDN market handles around 35 % of the global internet traffic [19] [23]. This share is expected to increase in the upcoming years. In particular, more than 55 % of the worldwide traffic is expected to be delivered by CDNs in 2017 [19].

The CDN market is composed of two main categories of players: those that use the internet for their connectivity requirements, denoted by pure-play CDNs, and those relying on end to end managed networks denoted by Telco-CDNs.

In its early days, the CDN market was limited to a reduced set of pure-play CDNs with a high capacity and a global reach or footprint. These include Akamai [32], Limelight Networks [82] and CDNetworks [50]. Akamai, for instance, owns more than 127 000 servers distributed in 1150 clusters and in 87 countries worldwide [32] [51]. Even though these players still own a large share of the CDN market (Akamai alone owns today 80 % of the CDN market share [19]), recent years have witnessed the rise of specialist and regionally-focused CDNs joining the ranks of pure-play heavyweights. For instance, BitGravity (founded in 2006, owned by and integrated with Tata Communications) is a CDN provider that is specialized in the delivery of on-demand video and High Definition (HD) live streaming over India and Australia [46]. In order to cope with evolving users' trends, giant pure-play CDNs like Akamai enhanced their basic functionalities (listed in section 2.2.2) for supporting the delivery of mutiple types of contents (video, live, web pages...) to a multitude of devices (fixed and mobile users) [58]. Value-added and customized services are also being proposed by many pure-play CDNs to their customers among Content Providers and businesses. These include online security, bitrate adaptation and portals personalization. On the other hand, the CDN market has also witnessed

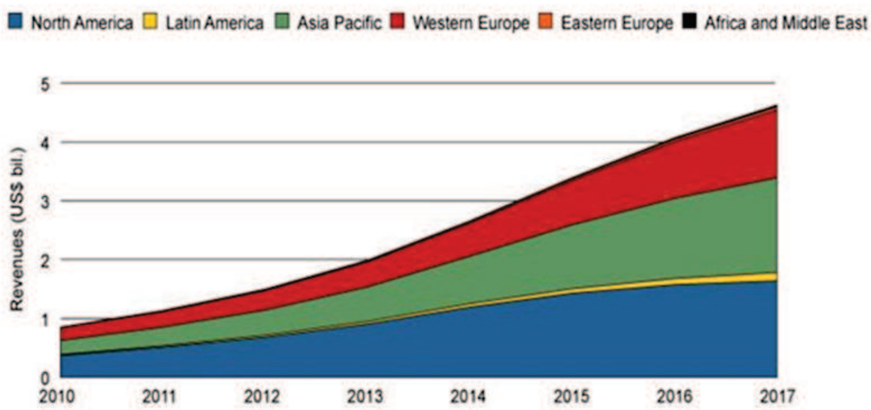

Figure 2.6: Overview of CDN Market Evolution [58]

the rise of a new category of players: the Telco-CDNs. The number of Telco-CDNs is growing through years and has reached 120 players in 2012 [58]. The majority of Telecom operators dedicate today their CDNs for their internal usage [2] [53]. However, more Telco-CDNs include the ability to manage over-the-top (OTT) Internet traffic. The monetization of the CDN service can be through a Business to Client (B2C) model (Retail CDNs [53]) or a Business to Business (B2B) model (wholesale CDNs [53]). In the B2B model, a Telco-CDN can propose a last-mile content delivery service to CPs and to pure-play CDNs. Telco-CDNs attractiveness in this context is limited to local and regional CPs. AT & T is an example of a Telco-CDN that has partnered, through a B2B model, with local CPs in the US [2]. Given their leadership position in the market, pure-play CDNs has not found so far enough incentives to delegate their content delivery to Telco-CDNs. Nevertheless, this trend is evolving as will be observed in the next section.

The benefits of CDNs for Telecom operators include improved traffic management and associated cost savings on IP transit, an enhanced QoE for their subscribers and the potential to deliver new services and generate new revenues. When compared to pure-play CDNs, Telecom operators own many assets of interest for content owners and providers [53]. End users' proximity, ownership of the networks that connect users to online contents and the ability to control and preserve throughput through many devices are the main assets. Telco-CDNs present a major drawback related to the national and local nature of their footprint. The collaboration of Telco-CDNs and of Telco and pure-play CDNs may be useful for overcoming this drawback.

2.2.4 Overview of CDN pricing schemes

CDN charging often occurs on a monthly basis. CDN providers charge CPs according to the traffic that these CPs generate and to the geographic regions where this traffic is generated. A charging region is often associated with a large country like the US or India or with a continent like Europe. The charging policy adopted by a CDN over a charging region depends of the amount of competition in this region. In particular, CDN providers charge high prices in regions where they monopolize content delivery. On the other hand, low prices per GigaByte (GB) of delivered content are charged in regions where the overlapping of CDN providers is important. Factors influencing the charging policy of CDN providers also include: The bandwidth cost, the variation in traffic distribution, the size of content catalog, the dimensioning of the CDN and the security requirements of the CP. [70] discusses how these different elements are likely to influence the CDN charging policy.
Different charging models exist in the CDN market. While some CDN providers provide online information about their charging schemes, the charging schemes of others remain unknown. Based on the charging information available online and on our discussions with well-known CDN providers, we can claim that two main trends, in terms of charging models, are observed in the CDN market.
The first charging model, denoted by commitment-based model, is adopted by market leaders like Akamai [32] and some Telco-CDNs like Orange [98]. It consists in the following: When partnering with a CDN, a CP provides to the CDN, on a monthly basis, its service requirements referring to information about its global demand, demand distribution across charging regions, traffic structure over different countries (time evolution of traffic)... The CDN procedes per charging region. In order to assess the monthly price charged to the CP over a given region, the CDN multiplies the amount, in GB, of CP traffic that is expected to be generated in one month over this region by a flat, linear per GB fee that depends on the region (amount of CDN competition or footprint overlapping in this region) and on the amount of expected traffic. The higher the traffic volume is, the lower is the per GB fee. For instance, an anonymous CDN charges, in average, 0.15 $ per GByte of delivered traffic to CPs that generate up to 10 TBytes of traffic per month over the US. The same CDN charges 0.03 $ per GByte of delivered traffic to CPs which traffic volumes exceeds 5 PBytes per month over the US. On the other hand, if we consider a region like Australia where the CDN competition is less fierce, the prices per GByte corresponding to different amounts of monthly traffic are higher than those charged in the US. The CDN uses the traffic forecasts of its clients in order to calculate the minimum fee that they should pay at the end of a billing cycle (a

Table 2.1: CloudFront Pricing Model [36]

	US	Europe	Asia	South America	Australia
First 10 TB/month	$ 0.120	$ 0.120	$ 0.190	$ 0.250	$ 0.190
Next 40 TB/month	$ 0.080	$ 0.080	$ 0.140	$ 0.200	$ 0.140
Next 100 TB/month	$ 0.060	$ 0.060	$ 0.120	$ 0.180	$ 0.120
Next 350 TB/month	$ 0.040	$ 0.040	$ 0.100	$ 0.160	$ 0.100
Next 524 TB/month	$ 0.030	$ 0.030	$ 0.080	$ 0.140	$ 0.095
Next 4 PB/month	$ 0.025	$ 0.025	$ 0.070	$ 0.130	$ 0.090
Over 5 PB/month	$ 0.020	$ 0.020	$ 0.060	$ 0.125	$ 0.080

billing cycle being a month). If the traffic generated by a given CP during a billing cycle is less than the forecasts, this CP still pays the pre-calculated minimum fee. If, instead, real traffic exceeds the expectations, the CDN charges an extra fee to the concerned CP for all traffic going beyond the requirements that it provided.

The second charging model, referred to as a pay as you go model, was introduced by rising pure-play CDNs like Amazon CloudFront [37] and MaxCDN [87] in order to compete with well-positioned CDNs like Akamai [32]. Contrarily to the first model, no minimum fee is charged to CPs. Instead, CPs exactly pay for the real traffic they generate during a billing cycle. In this context, each CDN maps its footprint to one or many charging regions. A non-linear, volume-based charging policy is adopted by the CDN over each of its charging regions. Typically, the CDN pre-defines traffic segments corresponding to different volumes of traffic generated per month and per CP. It then associates different fees per GB to different segments. Amazon CloudFront charging [36] model illustrates this approach. This model is shown in Table 2.1.

2.3 Evolution Trends

In this section, we discuss the evolution trends observed at the Internet edge and inside the network .

2.3.1 Transformation of the Internet Edge

Growing IP Video Traffic

Consumers' demand for video is rapidly growing over years. Internet video traffic has reached 50 % of the global IP traffic in 2012 [19]. According to [19], video traffic is expected to form 90 % of the worldwide internet traffic in 2017. It is also expected that, every second, the equivalent of a million minutes of video content will cross the internet [19]. This growth is mainly due to the growing number of connected devices and users and to the fact that more videos with a higher quality and a longer duration are available online. Beyond traditional CPs like Netflix [93], YouTube [127] and DailyMotion [55], the entry of players like Amazon and Apple in the online video market has also contributed to IP video growth.

Higher QoE constraints

Fixed and mobile end users are expecting a better quality of experience (QoE) when consuming web content, especially videos. In the context of video content, end users are particularly sensitive to three metrics: the time to first Byte, the mean bitrate and the buffering ratio. Different metrics are detailed in [57]. The time to first Byte refers to the delay between clicking on a video and visualizing it on the screen. The mean bitrate refers to the average resolution or definition of the video stream. The buffering ratio gives indications about the occurence of buffering events during an ongoing video session. These metrics have a direct impact of users engagement both during and after an ongoing video session.
Among the three metrics, [57] shows that the third metric has the highest impact on users' engagement. In particular, a 1 % increase in buffering reduces engagement by 3 minutes. As more users tend to consume premium video, video resolution is also expected to have a more important impact on users' engagement in next years. Finally, dependently of the video length, a high start time is likely to push end users to abondon the content provider's site [57].

Extended Mobility Requirements

As most consumers become mobile-oriented [20], more internet traffic is generated by mobile devices including smartphones and Tablets. In 2013, 20 % of the internet traffic was generated by non-PC devices. This percentage is expected to reach 50 % in 2017 [20] [19].
Ensuring a seamless content delivery in a mobile context is particularly challenging. Geographic mobility is a well-known challenge. As more mobile devices are able to connect, simultaneously or in a non-simultaneous fashion, to many types of access networks (3G, 4G, Wifi...) , dealing with access mobility becomes also crucial. Beyond geographic and access mobility, cellular and non-cellular accesses can be subject to important bandwidth fluctuations during very short time scales. When access mobility, geographic mobility or bandwidth fluctuations occur during an ongoing video session, the QoE witnessed by end users is likely to be subject to important degradations if adequate control mechanisms are not used. These mechanisms aim at ensuring a seamless flows re-routing and resoution adaptation during an ongoing video sessio. Well-known redirection mechanisms include network level approaches for mobility control like Mobile IP (MIP) [102]. Mechanisms for video resolution adaptation are limited to HTTP adaptive streaming in its industrial implementations (Microsoft Smooth Streaming [89], Apple Live Streaming [39]) and standardized form (MPEG DASH [116]).
Existing mechanisms for mobility control may not be enough for dealing with new forms of mobility. For instance, vertical mobility, referring to a change of user's home or service address, cannot be handled through MIP. Similarly, HTTP adaptive streaming is based on measurements performed by the terminal and not on a real vision of the network state. It also presents the drawback of being ractive since resolution adaptation is performed after QoE degradations are witnessed at the end user side. More sophisticated

control mechanisms involving new players of the content provisioning value chain may be required at this level.

Users-based Content Generation

With the spread of social networks [63] [121], users generate and publish an always increasing amount of content. Deployment of optical fibers for fixed access and introduction of Long Term Evolution (LTE) [95] and LTE-advanced [29] accessed is accelerating the pace mainly through increasing bandwidth availability for both uplink and downlink directions.

Emergence of new Types of Content

New types of content are emerging. The most relevant ones are those generated by systems like M2M [128] and IOT [41]. Those systems will generate an exponentially growing amount of content with an extremely high diversity and with very different structure and requirements than the mainly video centered state of things. Dealing with these contents require the deploying new types of caching, processing and bandwidth resources both at the edge of the network and in the Internet backbone.

2.3.2 Network transformation

The core of the internet is witnessing an important transformation. Indeed, technologies like Network functions virtualization (NFV) [47], Software Defined Networks (SDN) [88] and Information Centric Networks (ICN) [30] have emerged, although they are today at different levels of maturity.

Network Functions Virtualisation (NFV) and Software Defined Networks (SDN)

NFV is a concept that consists of realising network, control and application level functions which today reside in proprietary nodes on commodity IT servers by using modern virtualisation and cloud technologies. Since 2013, NFV is an Industry Specification Group (ISG) in ETSI [62]. The main objective of NFV is to build virtualised environment through general purpose hardware in order to realize today's nodes and functions as much as possible as pure software [47]. The NFV ISG group counts OTTs and network operators among its members. The activities of the group have been so far limited to the definition of the NFV architecture and to the specification of relevant use cases for NFV.

The NFV end to end architecture consists of three main components [47]: the NFVI (Network Functions Virtualisation Infrastructure), the Virtual Network Function (VNFs) and the NFV Management and Orchestration entity (NFV M & O). The NFVI refers to the physical infrastructure on top of which virtual compute, storage and network entities are instantiated. A VNF refers to a network, control or application level function or node that utilises virtualised resources of the NFVI. Typically, a VNF may use

virtual machines instantiated on top of many physical appliances and connected via a virtual network. The NFV M & O handles the orchestration between VNFs and virtualised resources in the NFVI. This includes instantiating the virtual networks required for ensuring the communication among distributed virtualised resources associated to the same virtual function.
Virtualisation can occur at the user plane, the VNF being for instance a forwarding engine. It can occur at the control plane, the VNF being for instance a (SDN) controller or at the application level, the VNF being for instance a CDN server. Different use cases for NFV are detailed in [73] . One of the use cases addressess CDN virtualisation. CDN virtualisation consists in transforming a CDN node into a VNF that can be implemented on top of distributed, general-purpose servers placed in the CDN domain. As a consequence, many virtual CDN nodes can coexist on top of the same physical server. Similarly, one virtual CDN node can be dynamically multiplexed on top of many physical appliances. CDN virtualisation has the advantages of opening CDNs to a wider portfolio of contents and content providers, reducing CDN providers Capex and ensuring a better resilience to overload events (due to the statistical multiplexing effect). The main drawback of CDN virtualisation is related to the increase of the Operational Expenditures (Opex) of CDN providres

SDN is a concept that is strongly pursued by the Open Networking Foundation since 2011. The term "Software Defined Networks" deals with network programmability. This concept is based on two main ideas: the abstraction of the physical network into a set of virtualised networks which can be used for dedicated purposes, the programmability of the virtualised networks through the use of a controller software [88]. SDN, in its OpenFlow version, was implemented by OTTs like Google in order to optimize flows routing between their own datacenters [81]. Nevertheless, beyond some PoC (Proof of Concept) iniatives launched by some OTTs and network operators internally, SDN is not yet widely deployed in the Internet.
SDN can be seen as an enabler of NFV. In fact, through enabling a separate control of network equipments, SDN allows instantiating many virtual functions on top of a single equipemnt. Examples of functions include flows routing and in-network storage of data packets.

Information Centric Networks (ICN)

The Internet speaks today in terms of content retrieval and generation rather than host to host communication. This motivated the development of future, disruptive internet architectures based on named data objects (NDOs) rather than physical addresses. The approach of these architectures is called Information Centric Networking (ICN).
The ICN architectures leverage in-network storage for caching, multi-party communication through replication and interaction models decoupling senders and receivers [30]. The goal is to achieve a worldwide, scalable and high performace platform for content distribution outside of dedicated systems like P2P [109] and CDNs.
Since the introduction of the ICN concept in Standford in 1999 [30], many ICN ar-

chitectures have been proposed in the litterature. Data Oriented Network Architecture (DONA) [38], Content Centric Networking (CCN) [74], Publish-Subscribe Internet Routing Paradigm (PSIRP) [31] and Network of Information (NetInf)[30] are the most known ones.

Despite their differences, different ICN architectures share many design components. To begin, they are all based on NDOs (Named Data Object) referring to web pages, videos or any type of object (full object or packet) that can be published by a source, stored in the network nodes and retrieved upon demand by any receiver/requester. A NDO is decorelated from a specific domain or physical address and is identified through a unique, universal name. Dependently of the ICN architecture, the name of a NDO can follow a hierarchical or a flat structure. Similarly, NDO names can be celf-certifying meaning that the object hash is in the name or can point to a trusted third party that is able to certify the name. The routing process consists of two phases: Routing of NDO requests and routing of NDO back to the requester. In some ICN architectures (DONA and PSIRP), a Name Resolution Service (NRS) is used in order to bind the NDOs to locators pointing to storage locations in the network. In this context, NRS can be seen as a network-level equivalent of the DNS [91] (DNS) used in traditional host to host communication. Once one or many locations are identified, the request message is routed to one or many of these locations and the corresponding NDO is routed back to the requester. In other ICN architectures like CCN, the NRS service is bypassed and the request message is directly routed by the requester to one or mutiple data sources based on the NDO name. Storage for caching NDOs is an integral part of the ICN service. All nodes are expected to have caches inluding nodes in operator networks infrastructure, home gateways and eventually mobile terminals. Requests for NDOs can be satisfied by any node holding copy in its cache. ICN hence enables a generic and dynamic in-network caching that can be applied to all types of contents including UGC contents.

The advocates of ICN list a number of advantages of this approach with respect to the current state of things, mainly referring to CDNs. When compared to CDNs, ICN leverages the use of in-network storage resources. Furthermore, it eliminates the overhead induced by some CDN functionalities such as DNS and HTTP based request routing, servers dimensioning and placement. By unbinding the name of an object from its geographic location, ICN prevents delays induced by DNS lookups and problems caused by a sudden change of object location or domain. From a security point of view, ICN ensures name-data integrity and origin verification of NDOs. Since ICN is not based on end to end connections, a moving client can continue to issue requests for NDOs (we assume a packet-level granularity) from new accesses. NDOs will be delivered from new sources and routed back to new accesses. Handover and Mobility scenarios are hence suported by default.

On the other hand, the ICN approach is still in the research stage and is facing many challenges. While ICN is intended to provide a scalable content distribution, the number of NDOs and their potential unhierarchical nature makes the aggregation of routing information in today's routers and in a future Name Resolution System (NRS) problematic. From a security point of view, the lack of end to end connections makes it difficult to tie a request to a particular person. From a legal point of view, ubiquitous caching may not sound very appealing to content owners and providers who prefer to know where their contents, including confidential ones, are stored. Finally and most importantly, the business model behind ICN is not clear. The ICN approach requires important investments in the network in terms of deploying new ICN-enabled network nodes or adapting the functionality and dimensioning of existing ones. Network operators should hence be able to monetize these investments meaning that in-network storage should be controlled. On the other hand, ICN should be positioned with regards to overlay and controlled CDNs. In this context, two scenarios are possible. Either ICN is proposed as an alernative for CDNs or in-network storage is seen as a complementary for CDN-based storage. Given the status-quo of the content provisioning value chain and the level of maturity of the ICN approach, the second scenario is likely to make more sense in short to mid term horizons.

2.4 Impact on CDN Stakeholders

Whether occuring at the edge of the Internet or inside the network, the trends listed in the above section directly impact the content provisioning value chain and, more specifically, CDN providers and network operators. Next, we detail the major impacts and forecast the short to mid term evolution of content distribution services.

2.4.1 Toward a closer and lower content distribution

One way of dealing with users' higher QoE constraints consists of placing content closer to them. Content providers and pure-play CDNs are already aware of the impact of distance on the delivery of sensitive contents like videos. This awareness is expected to widen in the coming years [85].
In fact, many OTTs including CPs and pure-play CDNs are bypassing the internet backbone through directly connecting their servers to Tier3 ISP networks. YouTube [127] and Netflix [93] are part of this trend and the recent agreement between Netflix and Comcast in the US goes in this direction [6]. Similarly, Akamai is today within a distance of one autonomous system of 90 % of worldwide internet users [64]. It is important to mention that OTTs like Akamai pay local ISPs for providing them with a direct paid peering connection [64], others like Netflix claim that they should not pay.

In order to bring their contents closer to end users, some regional CPs delegate their content delivery to Telco-CDNs. In addition, a new type of B2B partnerships between pure-play CDNs and Telco-CDNs is emerging. Given Telcos lack of CDN-related expertise, some pure-play CDNs propose to help Telcos in developping and operating their own CDN solution. Licensed and managed CDNs are good illustrations of this scenario

[53]. On the counterpart, pure-play CDNs request the termination of their clients traffic through the Telco-CDNs. In this context, a B2B agreement has recently been established between Akamai and Orange. According to this agreement, Akamai provides software licenses to Orange while using Orange CDN for terminating some its clients traffic over Orange broadband network in France [99]. EdgeCast [97] also provides software licenses to AT & T while using its CDN for terminating its clients traffic in the US [1].

On the other hand, the concept of in-network caching enabled through the virtualisation of netwok equipments is gaining more popularity among ISPs in general and Telcos in particular. Beyond the disruptive ICN approach, architectures like Next generation PoPs (NGPoP) [68] are being seriously investigated by Telcos even though they are not yet deployed. NGPoPs consist in forwarding engines and in-network caches co-located in points of presence in the Telco access network (e.g. at the e-NodeB level).

2.4.2 Content Delivery Networks re-defined

The raise of virtualisation and in-network caching techniques is motivating the re-definition of Content Delivery Networks (CDNs). According to the current state of things, a CDN consists of clusters of dedicated servers distributed, in an overlay fashion, in one or many geographic locations worldwide. The evolution of CDNs is likely to happen in two directions.

The first direction is related to the virtualisation of CDNs [73]. When CDNs become virtualised, dedicated servers supporting a restricted set of contents, formats and protocols are replaced by common purpose hardware that can be easily virtualised. As a consequence, any CDN will be able to deliver any kind of content independently of the related requirements and of the owning CP. CDN virtualisation have benefits for the CDN providers themselves but also for the CDN market in general. In paticular, through moving from dedicated CDNs to virtualised ones, the technological barrier for CDN collaboration no longer exists.

The second direction is related to the concept of in-network storage. In fact, even though some CDNs are located close to end users (Telco-CDNs), CDN nodes still belong to an overlay layer and are not co-located with network equipments. Nevertheless, network equipments including routers and access nodes own vacant storage and computing resources that can be used, when these equipments are virtualised, for purposes of content storage and delivery. An in-network CDN is hence likely to emerge even though it may not completely replace the traditional overlay CDN. This new form of CDN is likely to exist in a standalone fashion (a network operator that uniquely uses its in-network resources for content caching and delivery) or to co-exist with an already-deployed overlay CDN.
In summary, a future CDN is likely to consist of general purpose, easily virtualised overlay or in-network appliances placed in one or many geographic points of presence.

2.4.3 Collaboration of CDNs

As users request more contents and become more demanding in terms of content quality, well-positioned CDNs including Akamai [32] and Limelight Networks [82] may face capacity limitations during certain types of events. Flash crowds [40], referring to an unexpected surge of a given CP's traffic, are among the events that can potentially lead to CDN overload [76]. Furthermore, more OTTs are aware of the impact of the last-mile on the QoE witnessed by end users. On the other hand, Telco-CDNs are failing to establish B2B agreements with global CPs due to the local nature of their footprint. Finally, many pure-play CDNs do not own local clusters in some geographic regions where the demand on some global CPs websites is becoming more significant.

Even though the CDN market is competitive, CDN providers, whether they are well-positioned or emergent, have incentives to collaborate. QoE enhacement, capacity aggregation and footprint extension fall under the technical incentives. Revenue enhancement and resources optimization fall under the economic ones. CDN providers collaboration may occur in a centralized or distributed fashion.

One form of CDN collaboration consists of point to point, bilateral SLAs between two CDN providers. This form of collaboration has been investigated by the IETF [72] CDNI working group [102]. In its deliverables [45] [42], the IETF CDNI working group defines use cases for interconnecting two CDNs. Proposed use cases include, among others, capacity aggregation, footprint extension and QoE enhancement. Recent contributions of IETF CDNI focus on defining standards for interconnecting CDN providers using different CDN technologies. The CDN interconnection as defined by IETF CDNI consists of a simple selection of one or many downstream or local CDN(s) made, in a centralized and dynamic manner, by an upstream or global CDN.

Another form of CDN collaboration consists of orchestrating a given CP traffic among different CDN providers, each with its own business model and not necessarily aware of each others. Orchestration may be done by the CP itself, by an upstream or pure-play CDN or by a third party orchestrator or broker. CDN orchestration has been addressed in many scientific papers that focus, among others, on computing an optimal strategy of load sharing among different CDNs . [75] addresses the orchestration of massively distributed CDNs, a CDN being composed of in-network capacities located within a given ISP domain. Each of the CDNs is controlled by a centralized tracker. Orchestration is performed in a distributed fashion through the exchange of signalling messages between the trackers of different CDNs. The decision-making logic, implemented through an optimization model, focuses on performance enhancement and minimizing cross traffic costs. [83] addresses the orchestration of a reduced set of overlay CDNs with a global reach and a high capacity. The enabling architecture relies on a centralized optimizer. The optimization model aims at minimzing the price charged to the content provider (CP) while providing a good level of performance, in terms of perceived QoE, to consumers. CDN orchestrators and load balancers are also commercially available in the

CDN market.
CDN orchestrators provide a transparent and unique interface to CPs for building a virtual and global CDN that is based on a mutitude of physical CDNs. The selection and provisioning of physical CDNs as well as the dynamic routing of incoming requests among these CDNs are entirely handled by the CDN orchestrator in a complete transparency with regards to the CP. Proprietary, unknown load balancing rules are used in this context. OnApp [7] [86] and EdgeCast [97] are examples of CDN providers that implemented online platforms in order to allow, on the one hand, other CDN providers to sell their vacant capacity and, on the other hand, CPs to build online their virtual content delivery platforms by using this capacity.
Contrarily to a CDN orchestrator, a CDN load balancer intervenes after CPs have established SLAs with many CDNs in order to perform a real-time routing of users' sessions towards these CDNs. Well-known load balancers include Limelight traffic load balancer [94], Dyn CDN manager [59], Conviva [54] and Cedexis [52]. Many load balancers use pre-defined, statically computed rules including geographic location, time of the day, CDN performance, and weighted allocation for splitting their clients traffic among different CDNs. While rules like geographic location and time of the day are quite intuitive, the weighted allocation rule requires computing the shares of load (percentage) assigned to different CDNs. Unknown, proprietary algorithms are used in this context. Client-based load balancers including [54] and [126] monitor, in real-time, different CDNs performance and perform, based on it, a dynamic selection of the target CDN at the beginning of a video session and, sometimes, during video playback.

A third form of CDN collaboration consists of a federation of CDN providers that aggregate their respective assets in terms of capacity and geographic footprint and act as a global, unique CDN with regards to CPs. Putting in place a federation is challenging from both technical and economic perspectives. Indeed, the CDN providers in a federation should agree on a common business model of the federation and on an inner policy of revenue sharing. Furthermore, a number of technical issues related to the federation control should be addressed. These include content placement, load balancing and events management within the CDN federation in addition to inter-CDNs privacy issues.
With the rise of Telco-CDNs, some IT manufacturers among the key players in the market are developing proprietary solutions for allowing Telco-CDNs to federate. In this context, a federation solution named *cisco3* was introduced by Cisco in order to allow carrier CDNs including British Telecom [118], Deutsche Telecom [119] and Orange [98] to federate. Other manufacturers including Ericsson [61] and Alcatel Lucent [34] have joined this trend through proposing their own, proprietary solutions for CDN federation [12] [11]. Up to our knowledge, these solutions are still in the conception phase and have not translated into a reality in the CDN market yet.
The federation of distributed facilities has been already addressed in the litterature. [**?**] addresses the federation of distributed PlanetLab [104] facilities aiming to enhance their joint revenue which depends of the number of experiments accepted by the federation. The federation of CDNs is different from the federation of computing facilities because

it involves new types of content-specific requirements (defined in the CPs SLAs) and constraintts. For instance, the business model of a provider of computing facilities is different from the business model of a CDN provider).

2.4.4 Joint CDN-Network Operation

As the amount of video traffic generated by mobile devices significantly grows [20], mobile video is becoming a major challenge for both CPs and mobile operators. This is mainly due to the bursty and data-heavy character of video traffic as well as to the particular requirements of the mobile access (these requirements are detailed in section 2.3).
Telcos can capitalize on their knowledge of users' access context and of the mobile network conditions as well as on their ability to control the access network in order to propose a mobile-CDN service to CPs and pure-play CDNs. Beyond ensuring a last-mile delivery of videos to mobile users, a mobile Telco-CDN aims at controlling the path to users' devices in order to prevent QoE degradations due to bandwidth fluctuations in the mobile context.
More Telcos are aware of their priviliged position in a mobile context. Therefore, they are investigating the deployment of mobile equivalents for their fixed broadband CDNs [108]. Mobile Telco-CDNs can be based on in house solutions or on partnerships between Telcos and pure-play CDNs (Licensed, Managed CDNs [53]). On the other hand, pure-play CDNs are aware of the strong position of the Telco in the mobile context. Therefore, some pure-play CDNs are tending to establish QoS SLAs with Telcos in order to ensure a priviliged routing of their content towards the Telco mobile network.

2.5 Summary: Motivating our contributions

Despite its rising interest among IT manufacturers, the federation of CDN providers is one scheme of collaboration that has not been deeply investigated in the litterature. Federation aims at transforming autonomous and distinct CDNs into a unique and global CDN. The so-formed CDN has a significant competitive advantage in the market. Indeed, through aggregating their individual capacities and footprints, CDN providers are able to target more CPs and businesses thus enhancing their gains with respect to the separate case.
Contrarily to other schemes of collaboration that aim at maximing the gain of an upstream CDN, a CP or a third party orchestrator, the federation of CDN providers aims at maximizing the joint gain of all of the federation players. In addition, a fair share of the federation gain is assigned to each of the players. The notion of economic fairness is particularly interesting in a federation context and can translate into many strategies of revenue sharing within the federation. This notion is not taken into account in other collaboration schemes.
If we analyze the status-quo of content distribution services, we observe that, due to cost, capacity and QoE related reasons, more CPs are interested in delivering their content through third party infrastructures at the Internet edge. As the volume of content-related

traffic significantly increases, the CDN market demand moves in the same direction.
If we analyze the status-quo of the CDN ecosystem, we observe two main trends: first, pure-play CDNs are in a fierce competition which often leads to lower prices and lower margins achieved by these players. Second, Telco-CDNs are, on the one hand, not visible with regards to global CPs and, on the other hand, not competitive with regards to pure-play CDNs. Whether well-positioned or emergent, different players in the CDN market are likely to face limitations in terms of capacity, footprint or QoE with respect to the market demand.
Given high CDN market demand, federation of pure-play CDNs, of Telco-CDNs and of both categories of CDN players is a real opportunity for different CDN providers. Indeed, the federation of distinct CDN providers ensures their respective welfare while providing to CPs a unified access to a wide pool of distributed resources at a reasonable price that is inline with the CDN market trends.

In the remainder of this dissertation, we tackle the federation of autonomous and distinct CDN providers. we introduce a CDN federation solution that allows taking static decisions of federation establishment and provisioning and dynamic decisions of federation control. We enable static decision-making through proposing an optimization model that takes into account CPs service requirements and CDN providers capacity, footprint and economic constraints. We apply the model to different use cases of federation of interest for the CDN industry and we quantify, in economic terms, the interest of federation for different categories of CDN players. We enable dynamic decision-making through proposing algorithms for real-time request routing and events control within a federation of CDNs. We conduct trace-driven simulations in order to assess different strategies of dynamic decision-making. We quantify the gains, in terms of performance, of a federated approach for events control.
Through the different contributions of this thesis, we aim to justify the interest of the federation, both in economic and performance terms, for different categories of players in the CDN market. We equally aim at providing an overall technical solution that allows this particular form of CDN collaboration to translate into a reality in the CDN market in the coming years.

Chapter 3

CDN Federation: Control Architecture

In this chapter, we address the federation of distinct CDN providers, a CDN provider being a player that owns storage and streaming capacities at one or many geographic locations. We consider a system composed, on the one hand, of Content providers (CPs) with different service requirements and, on the other hand, of CDN providers with capacity and footprint constraints. CDN providers include traditional pure-play and Telco CDNs but the defintition widens to include Cloud providers and providers of in-network and virtualised resources like carriers.
A federation refers to a number of CDN providers that put together their resources in order to jointly target a market formed by a number of content providers (CPs). Putting in place a federation requires agreeing on a common business model of the federation and on an inter-players revenue sharing strategy. The federation of heterogeneous content players involves different steps. The first step addresses federation establishment or the selection of the players involved in a given federation and of the market targeted by this federation. The second step addresses federation provisioning which consists in computing, statically, an optimal strategy of load balancing within the federation. Optimality refers to performance, capacity and economic fairness criteria. The third step addresses the dynamic aspect of federation control consisting in a real-time routing of incoming requests to adequate CDNs. This step is based on the provisioning phase output and on a dynamic vision of different CDNs state.

We assume a control architecture based on a centralized controller. In order to take into account the privacy concerns of CDNs, this architecture can typically deployed by by an independent third party like a clearing house. The controller performs a set of functions that include: the selection of the CDN providers involved in a federation, the establishment of SLAs with the selected players, the computation of static strategies for content placement, dynamic requests routing and monitoring of peak events. Figure 3.1 gives an overview of the functions that can be performed by the controller. The cen-

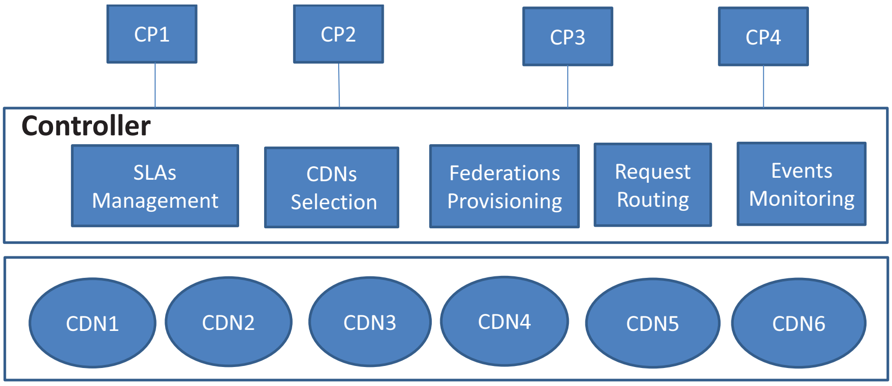

Figure 3.1: Sketch of the Control Architecture

tralized approach aims at removing the complexity induced by federation establishment, provisioning and dynamic control from the CDNs and to place it in a separate control plane. At this stage, a similarity with the SDN principles, as defined by ONF [66], could be mentioned. Indeed, the controller plays an intermediate role between the data plane composed of a set of distinct CDNs and the applications referring to a set of CPs with different service requirements. Contrarily to CDN orchestrators [97] [7], the role of the controller is limited to executing policies that the CDN providers agree on.

The architecture of the controller is shown in Figure 3.2. The architecture components intervene at three phases: federation establishment or provisioning, federation provisioning and federation control.

Let us introduce some notation that will be used in the following sections. All notations related with service requirements and with CDN characteristics are summarized in Table 3.1 and Table 3.2, respectively. We denote by:

- $\mathcal{G}$ the set of CPs under consideration, means those that issue service requests.
- $\mathcal{D}$ the set of CDNs willing to share part of their infrastructure.
- m a given CP in $\mathcal{G}$
- SR_m the service request issued by m
- $\mathcal{F}$ the footprint of $\mathcal{G}$, means the union of the footprints of each m

The footprint of a service request is defined as the set of consumers for which the corresponding content will be available.
The controller maintains in a *CPs Repository* the list of SR_m for all m in $\mathcal{G}$ and it is able to identify those that have to be treated at the next decision time (the new requests). The list corresponds to the market demand. Each request SR_m is composed of the following information (see Table 3.1 for notations): the catalog of content (we asume that all objects in the catalog are stream content with an intrinsic fix or average bitrate imposed by the coder), the mean size in Gbytes of a an object in the catalog, the mean duration in seconds of an object, the mean bitrate in Mbps of an object, the footprint

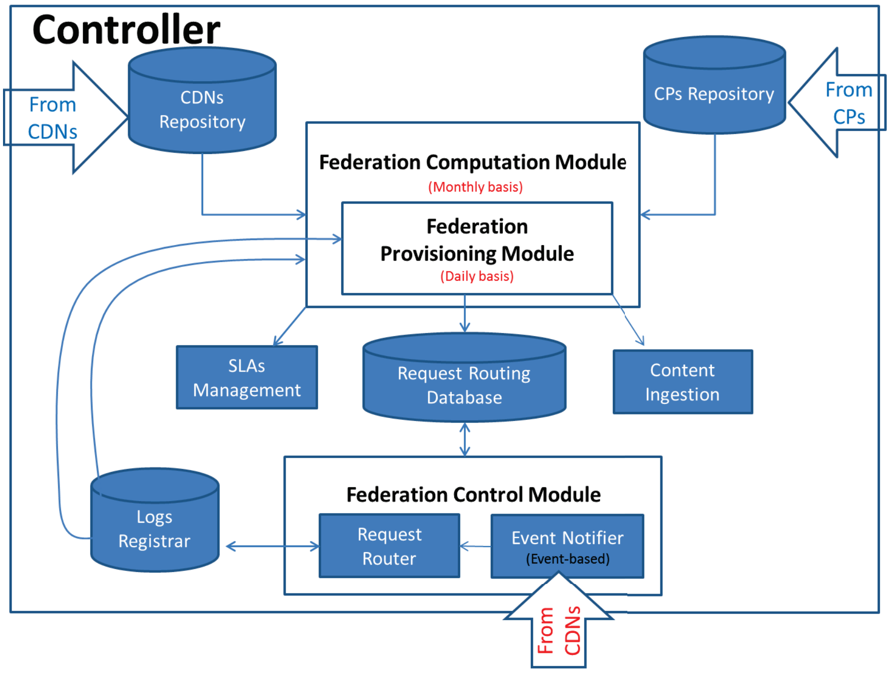

Figure 3.2: Controller Inner Structure

referring to a set of geographic and/or administrative zones targeted by the CP (a geographic zone refers to a country, an administrative zone refers to an Autonomous System (AS)), the number of content requests issued by consumers during a month, the variation of traffic distribution referring to the daily evolution of CP traffic, the distribution of the consumers' requests among the various zones of the footprint and the popularity of each object in the catalog over each zone of the footprint. It is important to note that popularity is expressed in terms of the ratio of requests targeting a well-defined content with respect to the overall volume of requests targeting the CP. It is measured over a daily time scale.
Each CDN in $\mathcal{D}$ provides to the controller a set of data (see Table 3.2 for notations) that is kept in a *CDNs Repository*, including the list of its clusters (points of presence), the footprint, storage and streaming capacity of each cluster and the charging policy per zone. A cluster foootprint is formed by the geographic zones (e.g. country, state) and/or administrative zones (e.g. Autonomous systems) that are within less than two ASs of the cluster servers. Each CDN also provides generic information about its cost model (charging model). We assume that the nodes/servers of a CDN provider are virtualised in the sense that they can be used for hosting and delivering any type of content. As a consequence, a CDN does not communnicate particular preferences in terms of content format or content delivery technique (HTTP, RTP...).

We denote by $\mathcal{S}$ the subset of $\mathcal{D}$ composed of CDNs that have footprint with a non-empty intersecton with $\mathcal{F}$. We refer to $\mathcal{S}$ as the set of candidates to target $\mathcal{G}$ demand

Table 3.1: CP Service Requirements

m	CP Id
K_m	Nb of sessions/requests served by m in 1 month
$\mathcal{C}_m$	Content Catalog of m
$\mathcal{F}_m$	Footprint of m
s_m	Mean size in GByte of a content object $\in \mathcal{C}_m$
d_m	Mean duration in seconds of a content object $\in \mathcal{C}_m$
b_m	Mean bitrate in Mbps of a content object $\in \mathcal{C}_m$
D_m^i	% of m demand originating from zone i $\in \mathcal{F}_m$
V_l^i	Popularity of content l $\in \mathcal{C}_m$ over i $\in \mathcal{F}_m$
$Peak_m^t$ (Mbps)	Peak amount of simultaneous traffic generated by m at an hour t (t $\in$ [1-24])

Table 3.2: CDN provider Inputs

j	CDN Id
$\mathcal{P}_j$	List of clusters $\in$ j
SC_p	Storage capacity in GByte of p $\in \mathcal{P}_j$
PC_p	Streaming capacity in Mbps of p $\in \mathcal{P}_j$
$\mathcal{F}_p$	Footprint of p
Dis_p^i	Distance (Nb of intermediate ASs) between p $\in \mathcal{P}_j$ and i $\in \mathcal{F}_p$

over $\mathcal{F}$.

3.1 Federation Computation

We assume that this phase is performed periodically by the *Federation Computation Module* of the controller at each billing cycle; a billing cycle being equal, for example, to one month.

Given a set $\mathcal{G}$ of requesting CPs and a set $\mathcal{S}$ of candidate CDNs, the aim of this phase is to identify one or many federations $\mathcal{S}'$ of CDNs in $\mathcal{S}$ targeting one or many CPs $\in \mathcal{G}$. A CP $m \in \mathcal{G}$ is said to be targeted by a federation $\mathcal{S}' \in \mathcal{S}$ if each j $\in \mathcal{S}'$ holds at least one l $\in \mathcal{C}_m$.
If no federation is able to address a given service request, we consider that the request is rejected.
Once a federation is established, the *SLA Management* engine of the controller handles the establishement of B2B SLAs between the members of the members of the federation and between the federation and its client(s).

If new content providers and/or CDN providers issue new federation or service requests during a billing cycle, the controller correspondingly updates CPs and/or CDNs Repositories. Nevertheless, these updates are only taken into account upon federation recomputation at the end of an ongoing billing cycle.

3.2 Federation Provisioning

The provisioning of a federation $\mathcal{S}'$ associated to one or many service requests SR_m is performed on a daily basis by the *Federation Provisioning Module* of the controller. It consists in deciding where, inside the federation, the contents of the CPs targeted by $\mathcal{S}'$ should be placed and how these future users' requests should a priori be distributed among different j $\in \mathcal{S}'$. The provisioning of different federations is performed along with federations computation at the beginning of each billing cycle. Once established, federations are re-provisioned on a daily basis during a billing cycle mainly to take into account content catalog updates (in case of user-generated content portals) and daily changes of contents popularity. The outputs of the provisioning phase, in terms of request routing decisions within established federations, are placed in the *Request Routing Database* of the controller. Outputs, in terms of decisions of content placement are communicated to *Content Ingestion* engine.

3.3 Federation Dynamic Control

The dynamic phase of federation control takes place between two phases of federation provisioning. It is performed by the *Federation Control Module* of the controller. This phase aims at ensuring a real-time routing of users' requests to adequate CDNs based on the provisioning phase outputs and on a dynamic vision of different CDNs state.
The *Federation Control Module* is composed of two engines: the *Request Router* and the *Event Notifier*. The *Request Router* directs users' requests to the appropriate CDNs based on the rules stored in the *Request Routing Database*. The *Event Notifier* assesses different CDNs performance through monitoring their workload. Upon overload, the *Event Notifier* updates the rules stored in the *Request Routing Database*.
Request routing is ensured through the use of international, standardized DNS-based or HTTP-based redirection mechanisms. Furthermore, the controller maintains request logs and uses them as inputs for re-provisioning established federations each day and re-computing federations each month.

Chapter 4

CDN Federation: Static Decision-making

In this chapter, we use optimization theory in order to address the static aspects of federation computation and provisioning. We introduce an optimization model that allows computing and provisioning federations at the beginning of a billing cycle. We present the model variables, the objective function and the constraints. In particular, we explain the notion of economic fairness referring to *fair* strategies of revenue sharing within a CDN federation. We explicit the model outputs in terms of established federations, revenue per CDN and strategies of content placement and load balancing. We introduce a variant of the model that is used for re-provisioning established federations on a daily basis. Finally, we assess the model performance and scalability.

4.1 Variables

The model includes two categories of variables: the decision-making variables and the provisioning variables. The main variables of the model are the followings:

- z_j is a binary variable which indicates whether a given j $\in \mathcal{S}$ is part of at least one CDN federation.

- γ_m is a binary variable which indicates whether the service request of a given m $\in \mathcal{G}$ is accepted or not.

- x_{jp}^{ml} is a binary variable which indicates whether a content l $\in \mathcal{C}_m$ is placed or not in cluster p $\in \mathcal{P}_j$. We refer to this variable as content placement coefficient.

- y_{jp}^{iml} is a binary variable that indicates whether a request originating from a zone i $\in \mathcal{F}$ and targeting a content l $\in \mathcal{C}_m$ should be redirected to a cluster p $\in \mathcal{P}_j$. We refer to this variable as request routing coefficient.

4.2 Objective Function: the point of view of CDNs

The model objective is to maximize the global gain of the system composed by all $j \in \mathcal{S}$. The objective function is considered as the difference between a global revenue function and a global cost function.

4.2.1 Revenue Function

The revenue of $\mathcal{S}$ is the sum of the revenues earned by different j $\in \mathcal{S}$ through fulfilling the demand of different m $\in \mathcal{G}$ whose service requests have been accepted. In this context, all j $\in \mathcal{S}$ agree on a common business model according to which a joint price is charged by $\mathcal{S}$ to each m $\in \mathcal{G}$. We denote by $Rev_{\mathcal{S}}^{mi}$ the price charged by $\mathcal{S}$ to a given m for handling its monthly demand over a zone i of its footprint. $Rev_{\mathcal{S}}^{mi}$ depends of the business model jointly adopted by all j in $\mathcal{S}$ and, in particular, of the volume in GBytes of m monthly demand over i. The revenue function can be expressed as follows:

$$Rev_{\mathcal{S}}(\gamma_m) = \sum_{m \in \mathcal{G}} \gamma_m \times \sum_{i \in \mathcal{F}} Rev_{\mathcal{S}}^{mi} \tag{4.1}$$

Equation (4.1) presents the revenue jointly earned by different j $\in \mathcal{S}$ at the end of a billing cycle if the monthly demand of the CPs targeted by different federations in $\mathcal{S}$ (m $\in \mathcal{G}$ / $\gamma_m = 1$) is inline with the requirements they provided at the beginning of the billing cycle (mainly referring to K_m parameters).

4.2.2 Cost Function

The cost of $\mathcal{S}$ is the sum of the costs induced by content storage in different $j \in \mathcal{S}$ and content delivery from different $j \in \mathcal{S}$ to different $i \in \mathcal{F}$.

The monthly cost induced by content storage in a given $j \in \mathcal{S}$ is a function of the content placement strategy. It is calculated as follows:

$$Cost_1^j(x_{jp}^{ml}) = \sum_{p \in \mathcal{P}_j} \sum_{m \in \mathcal{G}} \sum_{l \in \mathcal{C}_m} C_{stor} \times s_m \times x_{jp}^{ml} \tag{4.2}$$

Where C_{stor} is the mean operational cost required for storing 1 GByte of content in a server for one month (according to [84], C_{stor} is equal to 0.05 $ /GByte/month).

If j is a pure-play CDN (and is not a network operator), it should establish a peering agreement with one or many ISPs in order to reach its different footprint zones. We assume a paid peering agreement meaning that j pays a monthly cost to ISPs that provide it with a direct connectivity to their networks. When a given j in $\mathcal{S}$ contracts an ISP in order to reach a given i $\in \mathcal{F}$, the cost paid by j to this ISP is usually proportional to the maximal amount of traffic simultaneously routed between j and i. We denote by T_{pi}^t the peak volume of traffic routed by a cluster p $\in \mathcal{P}_j$ to a zone i $\in \mathcal{F}$ at an hour t of the day (t $\in$ [1-24]). T_{pi}^t depends of the traffic structure of different m $\in \mathcal{G}$ and, more particularly, of the peak amount of traffic reached by different m $\in \mathcal{G}$ at t. T_{pi}^t also depends of the policy of load sharing within $\mathcal{S}$ (y_{jp}^{iml} coefficients). In order to estimate T_{pi}^{t}, we assume an homogeneous structure of traffic among different i $\in \mathcal{F}$ (same shape of traffic daily evolution over different zones) and an homogeneous popularity of different contents during different hours of the day (The number of requests for a given content may evolve from one hour to another but the percentage of this content demand with regards to the overall demand is the same). T_{pi}^{t} can be approximated as follows:

$$T_{pi}^{t}(y_{jp}^{iml}) = \sum_{m \in \mathcal{G}} Peak_{m}^{t} \times D_{m}^{i} \sum_{l \in \mathcal{C}_m} V_{l}^{i} \times y_{jp}^{iml} \tag{4.3}$$

The monthly cost paid by a given j $\mathcal{S}$ for routing its content to a given zone i $\in \mathcal{F}$ can be approximated as follows:

$$Cost_{2}^{ji} = \sum_{p \in \mathcal{P}_j} A_{jp}^{i} \times \sum_{t \in [1-24]} T_{pi}^{t} \times \lambda_{jp}^{it} \times C_{band} \tag{4.4}$$

Where:
C_{band} is the mean cost for routing 1 Mbps through an intermediate AS during one month (According to [3], $C_{band} = 0.63\$Mbps$)
A_{jp}^{i} is a binary parameter that indicates whether j is a pure-play CDN or not. A_{jp}^{i} can be expressed as follows:

$$A_{jp}^{i}(Dis_{p}^{i}) = \begin{cases} 0 \; Dis_{p}^{i} = 0 \\ \\ 1, \quad Otherwise \\ \end{cases} \tag{4.5}$$

λ_{jp}^{it} is a binary variable that satisfies the following constraints:

$$\sum_{t \in [1-24]} \lambda_{jp}^{it} = 1 \ , \forall (j \in \mathcal{S}, p \in \mathcal{P}_j, i \in \mathcal{F}) \tag{4.6}$$

$$\sum_{t \in [1-24]} \lambda_{jp}^{it} \times T_{pi}^{t} \geqslant T_{pi}^{t_1} \ , \forall (j \in \mathcal{S}, p \in \mathcal{P}_j, i \in \mathcal{F}, t_1 \in [1-24]) \tag{4.7}$$

A linearization of the cost function is required in order for the model to be solved by CPLEX [5]. The linearization process is detailed in Appendix A.
It is important to note that equations (4.2) and (4.4) allow calculating the costs which are expected to be paid by different j $\in \mathcal{S}$ at the end of a billing cycle. Nevertheless, a gap may exist between the expected costs and the real ones. For instance, an unexpected surge of one or more CPs traffic over one or many of their footprint zones is likely to lead to higher content delivery costs for the CDNs handling this CP traffic. Similarly, an heterogeneous structure of CPs traffic over their respective footprint zones may lead to delivery costs for CDNs that are not inline with the expectations. A more accurate estimation of CDNs costs can be hence performed through monitoring real CDN traffic during a billing cycle.

4.2.3 Objective Function

Given the above Revenue and Cost fonctions, the objective function to be maximized is the following:

$$Rev_{\mathcal{S}} - \sum_{j\in\ \mathcal{S}} Cost_1^j - \sum_{j\in\ \mathcal{S}} \sum_{i\in\ \mathcal{F}} Cost_2^{ji} \tag{4.8}$$

4.3 Service Constraints

The content distribution and load balancing policies computed through the optimization model should take into account many service constraints. First, a content $l \in \mathcal{C}_m$ should not be placed in a given $p \in \mathcal{P}_j$ if there is no demand on l originating from p footprint. These constraints translate as follows:

$$x_{jp}^{ml} \leqslant \lceil \sum_{i\in\ \mathcal{F}_p} V_l^i \rceil \ , \forall (j \in\ \mathcal{S}, p \in\ \mathcal{P}_j, m \in\ \mathcal{G}, l \in\ \mathcal{C}_m) \tag{4.9}$$

Secondly, we correlate content placement and request routing. In fact, we assume that a request originating from a given $i \in \mathcal{F}$ and targeting a given $l \in \mathcal{C}_m$ should not a priori be routed to a cluster that does not host l. This translates as follows:

$$y_{jp}^{iml} \ \leqslant \ x_{jp}^{ml} \ , \forall (i \in\ \mathcal{F}, j \in\ \mathcal{S}, p \in\ \mathcal{P}_j, m \in\ \mathcal{G}, l \in\ \mathcal{C}_m) \tag{4.10}$$

When accepting the service request of a given $m \in \mathcal{G}$, the controller commits to completely alleviating m load. In other terms, any request coming from any $i \in \mathcal{F}_m$ and targeting any $l \in \mathcal{C}_m$ should be routed to a given $j \in \mathcal{S}$. Since the load balancing strategy computed at this level is static, we assume that all requests originating from the same i $\in \mathcal{F}$ and targeting the same l $\in \mathcal{C}_m$ should a priori be routed to the same j in $\mathcal{S}$. These constraints translate as follows:

$$\sum_{j\in\ \mathcal{S}} \sum_{p\in\ \mathcal{P}_j} y_{jp}^{iml} = \gamma_m \times \lceil V_l^i \rceil \ , \forall (i \in\ \mathcal{F}, m \in\ \mathcal{G}, l \in\ \mathcal{C}_m) \tag{4.11}$$

When accepting the service request of a given m in $\mathcal{G}$, the controller also commits to deliver the traffic of m from CDNs that are at most at a distance of two ASs of the requesting users. In order to enforce this condition, we introduce a binary parameter I_{jp}^i that indicates whether a cluster p $\in \mathcal{P}_j$ is or not within a distance of 2 ASs of a zone i $\in \mathcal{F}$. I_{jp}^i can be expressed as follows:

$$I_{jp}^i = \begin{cases} 1 \ , & Dis_p^i \leqslant 2 \\ 0, & Otherwise \end{cases} \tag{4.12}$$

The subsequent constraints are the followings:

$$y_{jp}^{iml} \leqslant I_{jp}^i \ , \forall (j \in\ \mathcal{S}, p \in\ \mathcal{P}_j, i \in\ \mathcal{F}, m \in\ \mathcal{G}, l \in\ \mathcal{C}_m) \tag{4.13}$$

4.4 Capacity Constraints

Given the fact that CDNs are subject to capacity limitations, taking into account capacity-related constraints is an important part of federation provisioning. In particular, the storage and streaming capacity of each j $\in \mathcal{S}$ should be considered when decisions of content placement and load balancing are made.
The volume of contents stored at a given cluster level should not exceed this cluster storage capacity. This can be expressed as follows:

$$\sum_{m\in \mathcal{G}} \sum_{l\in \mathcal{C}_m} x_{jp}^{ml} \times s_m \leqslant SC_p \ , \forall(j \in \mathcal{S}, p \in \mathcal{P}_j) \tag{4.14}$$

Similarly, the maximal amount of traffic simultaneously delivered by a CDN cluster should not exceed this cluster capacity referring to the sum of streaming capacities of the surrogates in the cluster. We denote by w_p^t the workload, in Mbps, of a given p $\in \mathcal{P}_j$ at an hour t of the day (t in [1-24]). w_p^t depends of the peak amount of traffic delivered by p to different i $\in \mathcal{F}_p$ during t. w_p^t can be expressed as a function of the peak traffic reached by different m $\in \mathcal{G}$ at t ($Peak_m^t$ parameters) and of the load balancing policy adopted within $\mathcal{S}$ (y_{jp}^{iml} coefficients). If we assume an homogeneous structure of traffic among different i $\in \mathcal{F}$ and an homogeneous popularity of contents over time, w_p^t can be expressed as follows:

$$w_p^t(y_{jp}^{iml}) = \sum_{m\in \mathcal{G}} Peak_m^t \sum_{i\in \mathcal{F}} D_m^i \sum_{l\in \mathcal{C}_m} V_l^i \times y_{jp}^{iml} \tag{4.15}$$

Constraints on clusters streaming capacity hence translate as follows:

$$w_p^t \leqslant PC_p \ , \forall(j \in \mathcal{S}, p \in \mathcal{P}_j, t \in [1-24]) \tag{4.16}$$

Additional bandwidth-related constraints exist on the paths between CDN clusters and end users as well as between the different clusters of a given CDN provider. Nevertheless, we cannot include these constraints in our model since a large percentage of the traffic circulating on these paths is a background traffic that is hard to predict. We hence assume that bandwidth provisioning on the content path is performed in real-time by CDNs and network operators .

4.5 Revenue Model

The global revenue jointly earned by all j in $\mathcal{S}$ is given in equation (4.1). We refer to this revenue as $\mathcal{S}$ global revenue.
We denote by $Rev_{\mathcal{S}}^j$ the federation revenue of a given j $\in \mathcal{S}$. $Rev_{\mathcal{S}}^j$ refers to the revenue earned by j through contributing to different federations in $\mathcal{S}$ targeting different $m \in \mathcal{G}$. we denote by α_j^{mi} the share of a given m demand originating from a zone i $\in \mathcal{F}$ and assigned to a given j $\in \mathcal{S}$. α_j^{mi} can be expressed as follows:

$$\alpha_j^{mi} = \sum_{l\in \mathcal{C}_m} V_l^i \sum_{p\in \mathcal{P}_j} y_{jp}^{iml} \tag{4.17}$$

The strategy of revenue sharing that we adopt consists in assigning to each j in $\mathcal{S}$ a revenue proportional to its contribution to different federations in $\mathcal{S}$. $Rev_{\mathcal{S}}^{j}$ can be hence expressed as follows:

$$Rev_{\mathcal{S}}^{j} = \sum_{m\in\ \mathcal{G}} \sum_{i\in\ \mathcal{F}} \alpha_{j}^{mi} \times Rev_{\mathcal{S}}^{mi} \tag{4.18}$$

Equation (4.14) allows calculating the federation revenues that are expected to be earned by different j $\in \mathcal{S}$ if the real demand of CPs in $\mathcal{G}$ targeted by different federations $\in \mathcal{S}$ is inline with the requirements they provided at the beginning of the billing cycle.

4.6 Rational Fairness

Different CDNs should find incentives in federating with regard to the case where they operate separately. In other terms, when a given j in $\mathcal{S}$ is assigned the share of the monthly demand of one or many m in $\mathcal{G}$, the subsequent revenue should be strictly higher than *j* separate revenue. Next, we explicit how the separate revenue of a CDN provider can be calculated. We then add constraints to the optimization model in order to enforce a rational form of fairness within computed federations.

4.6.1 Separate Revenue

We denote by Rev_j the separate revenue of a given j $\in \mathcal{S}$. Rev_j is the maximal revenue that can be earned by j through separately fulfilling the entire demand of one or many m $\in \mathcal{G}$ over its own footprint. In other terms, the streaming and storage capacities of j should be positioned with regards to the service requirements, in terms of catalog size and peak demand, of different m in $\mathcal{G}$ over $\mathcal{F}_j$ ($\mathcal{F}_j = \cup_{p\in\ \mathcal{P}_j} \mathcal{F}_p$). j will choose to target the CP(s) which bring him the highest revenue and which (joint) requirements are inline with his capacity constraints.

In order to calculate the separate revenue of a given j $\in \mathcal{S}$, we introduce a new optimization model that takes as inputs j capacity and footprint information as listed in Table 3.2 and the service requirements of different m $\in \mathcal{G}$ as listed in Table 3.1. The new model includes one binary variable β_j^m that indicates whether the service request of a given m in $\mathcal{G}$ over $\mathcal{F}_j$ is accepted or not by j. We denote by Rev_j^{mi} the price charged by j to a CP m when handling its monthly demand over a zone i in $\mathcal{F}_j$. Rev_j^{mi} depends of j charging model and, in particular, of the amount of m monthly demand over i. Given these elements, the separate revenue Rev_j of a given j in $\mathcal{S}$ can be expressed as follows:

$$Rev_j = \sum_{m\in\ \mathcal{G}} \beta_j^m \times \sum_{i\in\ \mathcal{F}_j} Rev_j^{mi} \tag{4.19}$$

The goal is hence to decide, for each j $\in \mathcal{S}$, different m $\in \mathcal{G}$ which service requests should be accepted by j if j wants to maximize its revenue while not overloading its servers. The corresponding optimization problem translates as follows (we assume an

homogeneous distribution of different m in $\mathcal{G}$ daily traffic over different i $\mathcal{F}$):
For each j $\in \mathcal{S}$,
Maximize

$$Rev_j = \sum_{m \in \mathcal{G}} \beta_j^m \times \sum_{i \in \mathcal{F}_j} Rev_j^{mi} \tag{4.20}$$

Subject to:

$$\sum_{m \in \mathcal{G}} \beta_j^m \times \sum_{i \in \mathcal{F}_p} D_m^i \times Peak_m^t \leqslant PC_p \ , \forall (p \in \mathcal{P}_j, t \in [1-24]) \tag{4.21}$$

$$\sum_{m \in \mathcal{G}} \beta_j^m \times s_m \sum_{l \in \mathcal{C}_m} \lceil V_l^i \rceil \leqslant SC_p \ , \forall (p \in \mathcal{P}_j, i \in \mathcal{F}_p) \tag{4.22}$$

In order to be part of one or more federations, the federation revenue of a given j in $\mathcal{S}$ should be superior to j separate revenue. The optimization model given by equations (4.20), (4.21) and (4.22) is solved for different j $\in \mathcal{S}$. This allows calculating the separate revenues of different j $\in \mathcal{S}$. Once these revenues are calculated, the following constraints are added to the main optimization model:

$$\sum_{m \in \mathcal{G}} \sum_{i \in \mathcal{F}} \alpha_j^{mi} \times Rev_{\mathcal{S}}^{ji} \geqslant z_j \times Rev_j \ , \forall j \in \mathcal{S} \tag{4.23}$$

Complementary constraints concerning the decision-making variables are added to the main model. If, due to the lack of economic incentives, the request of a given j is rejected, no content should be stored at j level. This can be expressed as follows:

$$x_{jp}^{ml} \leqslant z_j \forall (j \in \mathcal{S}, p \in \mathcal{P}_j, m \in \mathcal{G}, l \in \mathcal{C}_m) \tag{4.24}$$

When $\mathcal{S}$ is over-dimensioned with regard to $\mathcal{G}$ demand, the addition of the rationality constraints to the model can lead to the exclusion of one or more j $\in \mathcal{S}$ from the established federations ($z_j = 0$). This is particularly the case when a subset of CDN providers in $\mathcal{S}$ can target the demand of all CPs in $\mathcal{G}$. In contrast, if $\mathcal{S}$ is under-dimensioned with regard to $\mathcal{G}$ demand, each j $\in \mathcal{S}$ will be part of at least one federation. Each j $\in \mathcal{S}$ will also be given enough incentives to federate.

4.7 Computing Optimization and Outputs

We use a CPLEX solver [5] in order to solve the exact model for a well-defined dataset. We run the CPLEX solver on top of a remote Linux machine with a 2.4 GHZ Intel Xeon processor and 1.5 GB of RAM.
When solving the previous model for a well-defined dataset, we calculate the values of the decision-making and provisioning variables that lead to an optimal global gain of $\mathcal{S}$. According to the revenue sharing strategy that we consider, the federation revenue of each j $\in \mathcal{S}$ is function of one of the provisioning variables (Equation 3.14). Solving the model allows calculating the federation revenue of each j $\in \mathcal{S}$, referring, in case of

commitment-based charging (explained in section 2.2.4), to the minimal revenue earned by j at the end of the billing cycle. Solving the model also allows knowing different CDN federations to be established in $\mathcal{S}$ and the policies of content placement and load balacing to be adopted within these federations. Next, we explicit how the values of different variables allows identifying different federations $\mathcal{S}'$ of CDNs established in $\mathcal{S}$. We also explicit how these variables translate into static policies of content distribution and load balancing inside $\mathcal{S}$.

4.7.1 Established CDN Federations

In order to identify the CDN federations established in $\mathcal{S}$, the *Federation Computation Module* of the controller proceeds as follows:

- For each m $\in \mathcal{G}/\ \gamma_m = 1$, we identify the set $\mathcal{S}_m$ of CDNs in $\mathcal{S}$ sharing m demand. $\mathcal{S}_m$ is defined as follows:

$$\mathcal{S}_m = j \in \ \mathcal{S}/ \sum_{l \in \ \mathcal{C}_m} \sum p \in \ \mathcal{P}_j \sum_{i \in \ \mathcal{F}_p} y_{jp}^{iml} \neq 0 \tag{4.25}$$

- Once the sets of j $\in \mathcal{S}$ associated to different m $\in \mathcal{G}\ /\ \gamma_m = 1$ are identified, these sets are compared. If $\exists(m_1, m_2 \neq m_1) \in \ \mathcal{G}/\mathcal{S}_{m_1} = \mathcal{S}_{m_2}$, then all j $\in \mathcal{S}_{m_1} \cup \mathcal{S}_{m_2}$ belong to a single federation targeting m_1 and m_2. The same reasoning applies if the matching exists between more than two sets. If all sets are the same, a single federation targeting all m in $\mathcal{G}$ with an accepted service request will be formed. This is typically the case when the same footprint is targeted by different m $\in \mathcal{G}$. If all sets are different ($\nexists(m_1, m_2 \neq m_1) \in \ \mathcal{G}/mathcal S_{m_1} = \mathcal{S}_{m_2}$), each set will form a federation targeting a single m in $\mathcal{G}$. It is possible that a given j $\in \mathcal{S}$ takes part of two or more federations. The footprint of a given federation is the union of the footprints of its clients.

4.7.2 Content Distribution Strategy

The content placement coefficients (x_{jp}^{ml}) computed through the optimization model define the strategy of content placement inside $\mathcal{S}$. Once different federations are computed and B2B SLAs are established between federations and targeted CPs, the *Federation Provisioning Module* starts by filtering non-null coefficients from all computed content placement coefficients. These coefficients are used in order to map the contents of different m $\in \mathcal{G}/\ \gamma_m = 1$ to different j $\in \mathcal{S}/\ z_j = 1$.
The *Federation Provisioning Module* procedes per CDN. For each j $\in \mathcal{S}\ /z_j = 1$, the module identifies the contents that are mapped to a given j. The *Content Ingestion* engine then intervenes in order to route these contents from different CPs domains towards the origin servers of j. Once contents are pushed towards j, the controller does not have any visibility on these contents placement within j domain. Dependently of its outsourcing approach, j can either decide to pro-actively push ingested contents towards different p $\in \mathcal{P}_j$ or to pull, on demand, these contents from its origin servers towards its surrogates

(The pull and push approach are explained in details in section 2.2.2).

4.7.3 Requests Routing Strategy

The *Request Router* engine of the controller handles requests routing during the dynamic phase of federation control. The *Request Router* queries the *Request Routing Database* in order to decide towards which CDN a request identified by an origin zone and a target URL should be directed.
This requires translating the y_{jp}^{iml} coefficients computed through the optimization model into a format that can be supported by the database. The database can be seen as a three columns matrix. The first column refers to a geographic or administrative zone, the second column refers to a content URL and the third colum refers a CDN identifier or name. Each line of the matrix corresponds to a routing decision concerning a request identified by an origin zone and a target content. In order to fill the *Request Routing Database*, the *Federation Provisioning Module* of the controller filters non-null coefficients from the request routing coefficients computed through the optimization model (y_{jp}^{iml} coefficients). For each of the filtered coefficients, the *Federation Provisioning Engine* tracks three parameters: the zone Id i, the content id (m,l) and the CDN id j. The tracked parameters are used in order to progressively fill the *Request Routing Database*.

4.8 Federation Re-provisioning

The optimization model introduced in the previous sections allows computing and provisioning federations at the beginning of a billing cycle. When solved, this model allow taking two types of decisions: strategic decisions such as acceptation/rejection of CPs/CDNs requests and revenue distribution among different $j \in \mathcal{S}$ and provisioning decisions on how CPs contents should be placed and how CPs load should be shared within $\mathcal{S}$.
While strategic decisions are expected to remain unchanged during a billing cycle (one month), provisioning decisions should be updated on a daily basis. Indeed, these decisions are partly based on statistics that can change from one day to another. Statistics on contents popularity fall in this category. According to equation (4.18), the federation revenues assigned to different CDNs are proportional to the amount of CPs load that they handle. A change of contents popularity can hence translate into an uneven load distribution inside $\mathcal{S}$. As a consequence, contents placement should be re-computed on a daily basis so that the revenue of each $j \in \mathcal{S}$ remains proportional to its load.
The lack of daily provisioning is also likely to lead to a more complicated and costly phase of dynamic control. Namely, some overload events can be prevented through daily provisioning. Finally, if among the CPs targeted by $\mathcal{S}$, there are ones that rely on user-generated content, daily re-provisioning is key for taking into account the requirements, in terms of volume and popularity, of the newly-generated contents.

In order to re-provision established federations, we use a variant of the optimization model that was introduced in the previous sections. Solved on a daily basis, the new model presents the following modifications with regards to the initial model:

- The decision-making variables γ_m and z_j are fixed for the billing cycle duration and are assigned the same values as at the beginning of a billing cyle. This translates in additional constraints on these variables to be added to the initial model
- The content catalogs $\mathcal{C}_m$ of different m in $\mathcal{G}/\ \gamma_m = 1$ are updated in order to take into account contents that have been generated in one day
- The popularity V_l^i of different l in $\mathcal{C}_m\ /\ \gamma_m = 1$ is also adapted based on daily logs collected during the dynamic phase
- The rationality constraints are omitted and replaced by the following constraints ($Rev_{\mathcal{S}}^j$ are the federation revenues of different j $\in \mathcal{S}$ as calculated at the beginning of a billing cycle):

$$Rev_{\mathcal{S}}^j = \sum_{m\in\ \mathcal{G}} \sum_{i\in\ \mathcal{F}} V_l^i \times \sum_{p\in\ \mathcal{P}_j} y_{jp}^{iml} \times Rev_{\mathcal{S}}^{mi} \tag{4.26}$$

4.9 Model Performance and Scalability

We track the resolution time required for solving the instantiations of different use cases of federation involving different numbers of CDNs and CPs. We notice that the resolution time increases with the number of CPs and CDNs. For use cases including up to four CDN providers and three Content providers, the exact resolution time does not exceed 15 minutes. The achieved delays are compliant with the operational requirements of the control architecture (provisioning should be performed on a daily scale).

Resolution time also depends of the overlapping, in terms of footprint, of different CDNs in $\mathcal{S}$. In particular, the higher is the number of j $\in \mathcal{S}$ that cover each i $\in \mathcal{F}$, the higher is the resolution time. According to [2], The number of wholesale CDNs in the market does not exceed 30 providers. If we assume that a footprint zone is formed by a single country, we figure out that the number of overlapping CDNs is relatively limited. In fact, the most important overlapping is observed in the US where the number of overlapping CDNs is six (Verizon, AT& T, Akamai, EdgeCast [97], Amazon [36] and Level3).

In summary, given the status-quo of the CDN market, the resolution time required for solving the optimization model remains acceptable thus ensuring the scalability of a centralized approach for CDN federation.

Chapter 5

Use Cases of Federation

In this Chapter, we use the optimization model introduced in the previous Chapter in order to analyze three use cases of CDNs federation. The first use case addresses the federation of Telco-CDNs located in the same country. The second use case addresses the federation of pure-play CDNs with overlapping footprints. The third use case addresses the federation of Telco-CDNs and pure-play CDNs. We use a traffic dataset in order to gather information on the requirements of three CPs over a country C. This information is used as inputs for the optimization model for each of the use cases that we consider. Given this dataset, we solve the model and demonstrate the interest of federation in each of the considered use cases. We also assess, for each of the use cases, the gains achieved by CDN providers through moving from a separate scenario into a federation.

5.1 Traffic Dataset

In order to assess the interest of federation in each of the considered use cases, we consider a federation market formed by a group $\mathcal{G}$ of CPs including FranceTelevision (FT), DailyMotion (DM) and YouTube (YT). We perform traffic measurements in Orange network in order to obtain the service requirements of the different CPs over a footprint formed by Orange fixed subscribers in France (anonymous data). In particular, we analyze a traffic trace including mixed web traffic collected through probes placed in five BAS (Border access servers) placed in Orange France network.

We analyze the traffic trace in order to gather statistics on different CPs demand (Nb of requests per day, popularity, mean content size...) over Orange AS. We extend the obtained statistics to an anonymous country C. The service requirements of different CPs over C are shown in Table 5.1.
We conduct trace-driven simulations in order to know the structure of YT, DM and FT traffic over Orange AS. Traffic structure refers to the daily, per hour evolution of these CPs traffic over Orange AS (Autonomous System). In this context, we use a simulator consisting in a Software developed in the context of the OCEAN european project [84].

The simulator takes two kinds of inputs: traffic inputs and parameters. Traffic

Table 5.1: CPs Service Requirements

	FranceTelevision	DailyMotion	YouTube
Nb of Requests/day/C	130000	80000	400000
Nb of Requested Objects	2500	8000	2500
Mean object size (GB)	0.04	0.02	0.0125
Content Catalog size (GB)	100	160	40
Mean object duration (Seconds)	160	160	100
Mean object bitrate (Mbps)	2	1	1
Peak Demand/C (Mbps)	600	200	600

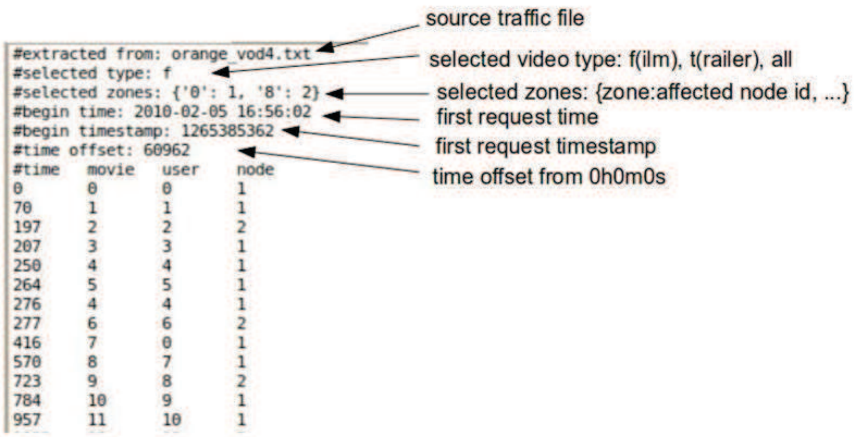

Figure 5.1: Format of the Traffic file

inputs consist in a set of video requests that are either based on pre-defined models or correspond to a real traffic trace. These inputs are placed in a file according to the column row format shown in Figure 5.1. Each row corresponds to a video request. The first column corresponds to the time when a video is requested, the second colum correponds to the movie Id, the third colum correponds to the user Id and the fourth column correponds to the node Id referring to the network node towards which the request should a priori be routed. The parameters of the simulator can be classified into two categories: the simulation parameters and the network parameters. The network parameters are specified in a separate file where the structure of the network is described. This includes the names and sizes (storage capacities) of the network nodes, their caching policies (LRU, LFU...) and their collaboration scheme if applicable. The simulation parameters include, among others, the time step of the simulation, the output interval and default values for video bitrates and durations. We consider a network formed by three nodes, each node being assimilated to a CDN. We filter YT, DM and FT traffic from the original traffic trace. We use the root URL of requested contents in order to decide the target node towards which a request should be directed. We assume that requests targeting YouTube URLs are directed to cdn_1, requests targeting DailyMotion

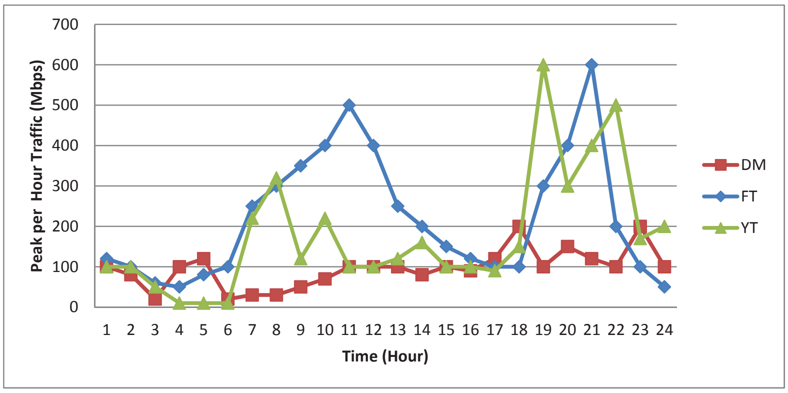

Figure 5.2: Structure of YT, DM and FT Traffic over C

URLs are directed to cdn_2 and requests targeting France Television URLs are directed to cdn_3. We translate the traffic trace into the format that is supported by the simulator (a user mac/IP address is translated into a user or zone Id, a movie URL is translated into a movie Id, the node targeted by each of the requests is specified). the resulting trace is used as a traffic input for the simulator. We use a small time step (1 second) in order to track, each second, the load of different CDNs referring to the amount of traffic simultaneously generated by a CDN in one second. We then replace the load values (in Mbps) tracked at different CDNs level during an hour t (t $\in$ $[1-24]$) by the maximal value reached during a time step in t. We hence obtain the daily, per hour evolution of YT, DM and FT traffic over Orange AS. The daily, per hour evolution of YT, DM and FT traffic over C ($Peak_m^t$ parameters) can be obtained through stretching the curves visualizing the evolution of these CPs traffic over Orange AS. The daily, per hour evolution of YT, DM and FT traffic over C is shown in Figure 5.2.

5.2 Use Case 1: Federation of Local Telco-CDNs

As discussed in section 2.3.3, more CPs and pure-play CDNs are aware of the impact of last-mile content delivery on the QoE perceived by end users. Thus, they are keen to establish B2B agreements with Telcos that can terminate their traffic from their own last-mile CDNs. Due to the high competitiveness of the european Telco market, we can notice that, on a country scale, no single Telco owns market monopoly. Indeed, in most countries, the market is shared between many well-positionned Telcos. Even though more Telcos are deploying their own CDN and/or Cloud platforms (e.g. Orange

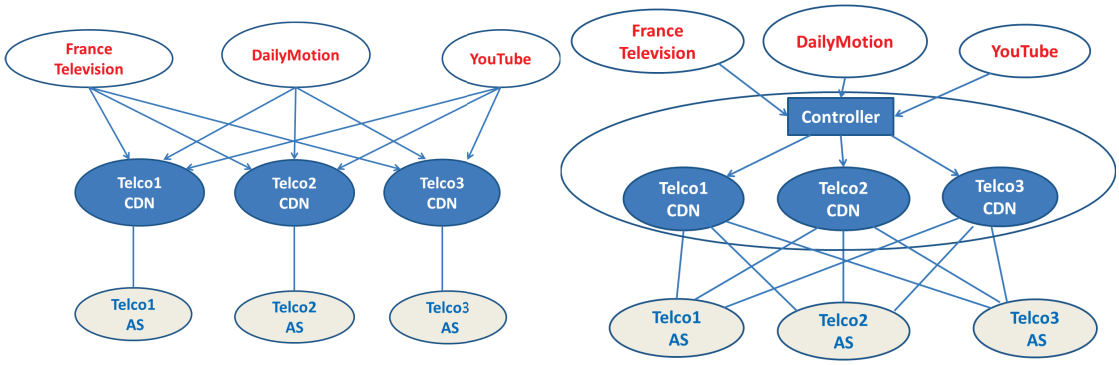

(a) $Scenario_1$ (b) $Scenario_2$

Figure 5.3: $Scenario_1$ Vs $Scenario_2$

in France, DT in Germany...) that are open to Over The Tops (OTTs), Telcos maintain a mode of content delivery according to which the CDN deployed by a given Telco only handles requests issued by this Telco subscribers. As a consequence, if, in a country, a CP or pure-play CDN aims to ensure a last-mile delivery of popular contents to end users, it should establish separate agreements with each of the Telcos in this country. In order for this status-quo to evolve, we suggest that Telcos put together their respective resources and make use of their established peering connections in order to offer a unique, country-scaled CDN service to external customers that can be local or global CPs as well as pure-play CDNs.
Next, we instantiate a use case related to the federation of local Telco-CDNs to which we apply our optimization model. We explicit the different hypothesis that we consider. We also exploit the model outputs in order to quantitatively assess the eventual gains of the federation.

5.2.1 Use Case Presentation and Inputs

The use case that we are interested in is presented in Figure 5.3. Given a group $\mathcal{G}$ of CPs with the requirements explicited in section 5.1, we consider a group $\mathcal{S}$ of CDN providers formed by three Telco-CDNs in C. We denote by $Telco_1$ the first Telco, $Telco_2$ the second Telco and $Telco_3$ the third Telco. We assume that the three Telcos share the Telco market in C. The market shares of the three Telcos are shown in Table 5.2. We assume that the market shares of the three Telcos also correspond to the D_m^i coefficients, referring to the distribution of different CPs demand among these Telcos subscribers. We also assume an homogeneous traffic distribution among all subscribers in C independently of their home ISP. Based on it, the daily, per hour evolution of YT, DM and FT traffic over $Telco_1$, $Telco_2$ and $Telco_3$ ASs are obtained through shrinking the curves in Figure 4.2 by the D_m^i coefficients. The evolution of YT, DM and FT traffic over different operators ASs in C is respetively shown in figures 5.4.a, 5.4.b and 5.4.c .

Since the service requirements that we own are measured over a fixed access, we assume that the three Telco-CDNs are dedicated to fixed users. Random storage and

Table 5.2: Market shares of Telcos in C

$Telco_1$	45 %
$Telco_2$	35 %
$Telco_3$	20 %

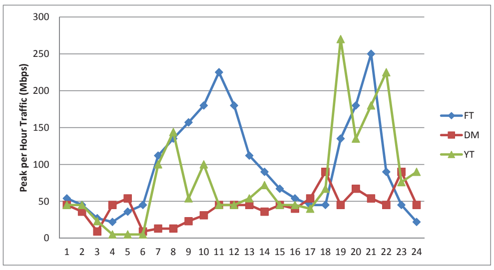

(a) Traffic evolution over $Telco_1$ AS

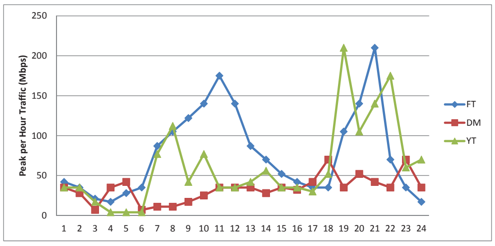

(b) Traffic evolution over $Telco_2$ AS

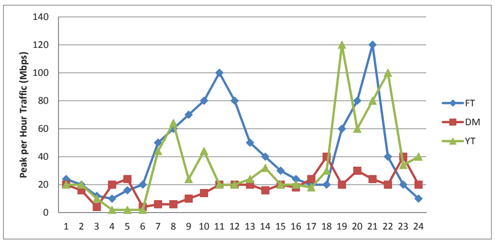

(c) Traffic evolution over $Telco_3$ AS

Figure 5.4: YT, DM, FT Traffic Structure per Telco AS

Table 5.3: CDNs Capacity Inputs

	$Telco_1$ CDN	$Telco_2$ CDN	$Telco_3$ CDN
Storage Capacity (GB)	100	100	100
Streaming Capacity (Mbps)	400	400	400

Table 5.4: Telcos Charging Model

Amount of traffic over C	per unit fee ($ /GB)
[0-1] TB	0.24
[1-10] TB	0.12
[10-50] TB	0.08
[50-150] TB	0.06
[150-500] TB	0.04
$\geqslant$ 500 TB	0.02

streaming capacities are assigned to these CDNs while assuming that the three CDNs are equally dimensioned. The capacity information of $Telco_1$, $Telco_2$ and $Telco_3$ CDNs are shown in Table 5.3. Each Telco-CDN has its own charging/cost model. For simplicity purposes, we assume that the three Telco-CDNs adopt the same commitment-based charging model (commitment-based charging model is explained in details in section 2.2.3). According to this model, the minimum price charged by a Telco-CDN to a given CP depends of the amount of traffic that this CP is expected to generate over this Telco AS in one month. In order to calculate this price, the amount of expected traffic is mutiplied by a flat, per unit fee (price per GB) which is function of this amount. Table 5.4 shows the fees per GB that each of the Telco-CDNs associates to different levels of clients commitments in terms of expected volume of monthly traffic.

In the actual context of content delivery, the footprint of each Telco-CDN is limited to its own Autonomous System (AS). In other terms, the traffic delivered by the CDN of a given Telco can only target this Telco subscribers. Nevertheless, direct peering connexions tie the IP backbones of all local Telcos in C thus enabling content stored at any Telco-CDN level to reach the subscribers of another local Telco. When peering connections allow CDN traffic exchange, the footprint of each of the Telcos is extended from its own AS to all ASs in C. The distances between each of the Telcos and different ASs are shown in Table 5.5. As noticed, any user in C is within less of one AS of the CDN that is going to deliver its content. Using existing peering connections for CDN traffic exchange can lead to traffic asymmetry between Telcos ASs. When this asymmetry occurs, the existing threshold-based peering model can be applied. According to this model, as long as the volume of IP traffic routed by a given Telco to another does not exceed a threshold defined as the double of the volume routed in the opposite direction, the symmetry is maintained and no charging is performed between the two. If this condition is not met, the Telco that is generating more traffic will be charged a fixed fee per Mbps for all the traffic going above the defined threshold (Charged fee is around 2 $ per Mbps).

Given three Telcos with established IP peering connexions, a threshold-based peering

Table 5.5: CDNs Footprint Info: Distance to target ASs

	$Telco_1$ CDN	$Telco_2$ CDN	$Telco_3$ CDN
$Telco_1$ AS	0	1	1
$Telco_2$ AS	1	0	1
$Telco_3$ AS	1	1	0

model and well-defined storage and streaming capacities, we aim at comparing two scenarios: the first scenario (shown in Figure 5.3.a) refers to the case where each Telco in C uses its own CDN in order to separately target the demand on the three CPs contents originating from its fixed subscribers. The second scenario (shown in Figure 5.3.b) refers to the case where any fixed user in C is reachable by any Telco-CDN and where all Telcos put together their CDNs in order to jointly target the demand of all CPs in $\mathcal{G}$.

5.2.2 Evaluation Methodology and Outputs

Our methodolody is as follows: In order to assess the first scenario, we consider three zones defined by the three ASs respectively belonging to each of the three Telcos. We position the requirements of each of the CPs over each of the zones with regards to the CDN capacity of the Telco that operates this zone. In order to assess the second scenario, we consider a footprint defined by all ASs in C. We position the requirements of each of the CPs over C with regards to the joint capacity of the three Telco-CDNs. In the light of this positioning, we identify the CPs that can be targeted by each Telco-CDN in the first scenario and by a federation of Telco-CDNs in the second one. The revenue assigned to each of the Telco-CDNs in both scenarios is calculated as well.

We begin by analyzing the first scenario. In this scenario, a Telco-CDN will only target the CP(s) which content catalog(s) fit inside its CDN and which (joint) peak traffic does exceed its streaming capacity. The size of different CPs catalogs is shown in Table 5.1. The daily evolution of different CPs traffic over $Telco_1$ AS, $Telco_2$ AS and $Telco_3$ AS respectively is shown in Figures (5.4.a), (5.4.b) and (5.4.c). We position the storage capacity of each Telco-CDN (Table 5.3) with respect to the size of different CPs catalogs. We also position the streaming capacity of each Telco-CDN with respect to the peak traffic reached by different CPs over this Telco AS. We figure out that, in the first scenario, each Telco has an interest in fulfilling FranceTelevision demand originating from its own subscribers. Subsequent revenues earned by different Telcos are shown in Table 5.7 (Sc_1 Rev). Calculated revenues correspond to Telco-CDNs separate revenue. Since, in this scenario, CDN traffic does not circulate across peering links, costs are only related to content storage. The values of the costs corresponding to the first scenario are also shown in Table 5.7 (Sc_1 Cost).
Once the first scenario is analyzed, we focus on the second scenario. In this scenario, we consider a three zones footprint formed by the ASs of the three Telco in C. In fact, the three Telcos form a virtual CDN that can reach all ASs in C and which capacity is an

Table 5.6: Prices per CP and zone ($Rev_{\mathcal{S}}^{mi}$, unit: $ /Month)

	$Telco_1$ AS	$Telco_2$ AS	$Telco_3$ AS
YT	4050	3150	2400
DM	1728	1344	1552
FT	4212	3276	2496

aggregation of the individual capacities of different Telcos-CDNs. The so-formed CDN has enough storage capacity to store all CPs catalogs and enough streaming capacity to handle these CPs joint peak traffic (Figure (5.2)). We use CPs and CDNs inputs figuring in Tables 5.1, 5.2, 5.3 and 5.4 as well as the the separate revenues of different Telco-CDNs as inputs for our optimization model. We assume that, in the second scenario, the three Telco-CDNs jointly charge each CP in $\mathcal{G}$ a monthly price per zone (Telco AS). We also assume that the price charged by $\mathcal{S}$ to any CP over a well-defined zone is equal to the price that could have been charged to this CP, in the first scenario, by the Telco that operates this zone. The prices charged by $\mathcal{S}$ to YouTube (YT), DailyMotion (DM) and FranceTelevision (FT) for handling their demand over different ASs in C are shown in Table 5.6.

In the light of these inputs, we solve the optimization model introduced in the previous chapter. We use equation (4.18) in order to calculate the respective revenues of $Telco_1$, $Telco_2$ and $Telco_3$ CDNs. We use the α_j^{mi} coefficients (equation 4.17) in order to calculate the shares of YT, DM and FT monthly demands over C respectively assigned to $Telco_1$, $Telco_2$ and $Telco_3$ CDNs. The cost induced by content storage in each of the CDNs is calculated by using equation (4.2). The threshold based peering model is used in order to calculate the content delivery-related cost of each of the Telcos. In particular, the peak amount of traffic delivered by a Telco-CDN to a foreign AS can be calculated using equation (4.3) (the peak amount of traffic delivered by a given p $\in \mathcal{P}_j$ to a zone/Telco AS i $\in \mathcal{F}$ is equal to $max_{t \in [1-24]} T_{pi}^t$).

Diffferent outputs of the model are shown in Table 5.7.

5.2.3 Discussion

Given the outputs of the optimization model, the main conclusions that we reach are the followings:

1. Through federating, Telco-CDNs are able to jointly address the requirements of a higher number of clients. Indeed, in $Scenario_2$, the demand of the three CPs can be jointly addressed by the three Telco-CDNs while only FranceTelevision would have been targeted in $Scenario_1$. The fact that content should no longer be replicated in all CDNs in order to be reached by all users allows new contents to be stored. In parrallel, balancing CPs load among different CDNs allows gains from statistical multiplexing to be achieved.

Table 5.7: $Scenario_1$ Vs $Scenario_2$ Outputs

	$Scenario_1$	$Scenario_2$
Global Revenue ($ /month)	9984	23808
Global Cost ($ / month)	15	100
$Telco_1$-CDN Revenue ($ / month)	4212	9017
$Telco_1$-CDN Cost ($ /month)	5	4.99
$Telco_2$-CDN Revenue ($ /month)	3276	8269
$Telco_2$-CDN Cost ($ / month)	5	4.99
$Telco_3$-CDN Revenue ($ / month)	2496	6521
$Telco_3$-CDN Cost ($ / month)	5	91
% of YT demand handled by $Telco_1$-CDN	0	41.11
% of YT demand handled by $Telco_2$-CDN	0	46.35
% of YT demand handled by $Telco_3$-CDN	0	12.53
% of DM demand handled by $Telco_1$-CDN	0	70.52
% of DM demand handled by $Telco_2$-CDN	0	21.97
% of DM demand handled by $Telco_3$-CDN	0	7.49
% of FT demand handled by $Telco_1$-CDN	45	25.154
% of FT demand handled by $Telco_2$-CDN	35	31.36
% of FT demand handled by $Telco_3$-CDN	20	43.48

2. In $Scenario_2$, The demand of the three CPs is distributed among the three Telco-CDNs. Furthermore, each Telco-CDN handles CPs demand originating from its own subscribers as well as from the subscribers of other Telcos in C. The distribution of the three CPs demand among the three Telco-CDNs is shown in Table 5.7.

3. From an economic perspective, $Scenario_2$ is beneficial for all Telco-CDNs. The revenues assigned to different Telco-CDNs in $Scenario_2$ are shown in Table 5.7 (Sc_2 Rev). These revenues are strictly superior to Telco-CDNs separate revenues (Sc_1 Rev). The values of content distribution and delivery related costs paid by the Telco-CDNs in $Scenario_2$ are shown in Table 5.7 (Sc_2 Cost). In particular, we notice that the cost paid by $Telco_3$-CDN in $Scenario_2$ is significantly higher than the one it pays in $Scenario_2$. This is due to the fact that the traffic delivered by $Telco_3$-CDN to other Telcos subscribers is significantly higher than the traffic delivered by other Telco-CDNs to $Telco_3$ subscribers. The gains achieved by different Telco-CDNs in $Scenario_1$ and $Scenario_2$ are the difference between the revenues and costs of these Telco-CDNs in each of the scenarios. These gains are shown in Figure 5.5. We notice that, through moving from $Scenario_1$ to $Scenario_2$, the gain of each of Telco-CDN $\in \mathcal{S}$ is multiplied by more than two. The difference between $Scenario_2$ gain and $Scenario_1$ gain allows assessing the federation gain of each of the Telco-CDNs, referring to the federation attractiveness for each of the Telcos.

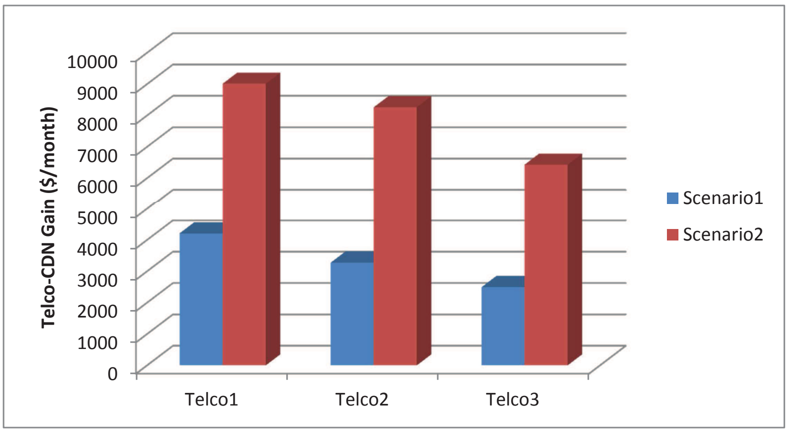

Figure 5.5: $Scenario_1$ Vs $Scenario_2$ Gains

5.3 Use Case 2: Federation of pure-play CDNs

In this section, we address the federation of pure-play CDNs with overlapping footprints. In particular, we consider a country C and two pure-play CDNs with surrogate servers that are connected, through direct peering points, to local ISPs in C. Next, we explicit the different hypothesis that we consider, we solve the optimization model and we assess the economic gains that can be achieved by pure-play CDNs when federating over one country of their common footprint.

5.3.1 Use case Presentation and Inputs

The use case that we are interested in is shown in Figure 5.6. We consider the same group $\mathcal{G}$ of CPs as for use case 1 and a one country footprint formed by C ($\mathcal{F} = C$ and $D_m^C = 1$, $\forall m \in \mathcal{G}$).

Given a group $\mathcal{G}$ of CPs with the service requirements described in section 5.1, we consider a set $\mathcal{S}$ of CDNs formed by two pure-play CDNs. We denote by CDN_1 the first CDN and by CDN_2 the second one. CDN_1 may correspond to Akamai and CDN_2 to EdgeCast. We assume that each of the CDNs owns a cluster of surrogates in C that is directly connected to local ISPs. We denote by p_1 the cluster of CDN_1 and by p_2 the cluster of CDN_2. The capacity information of p_1 and p_2 is shown in Table 5.8. Since CDN_1 and CDN_2 have direct peering agreements with local ISPs in C, p_1 and p_2 are both within a distance of 1 AS of C.

As for use case 1, we assume that CDN_1 and CDN_2 adopt the same commitment-based charging model. According to this model, the minimum price charged by any of

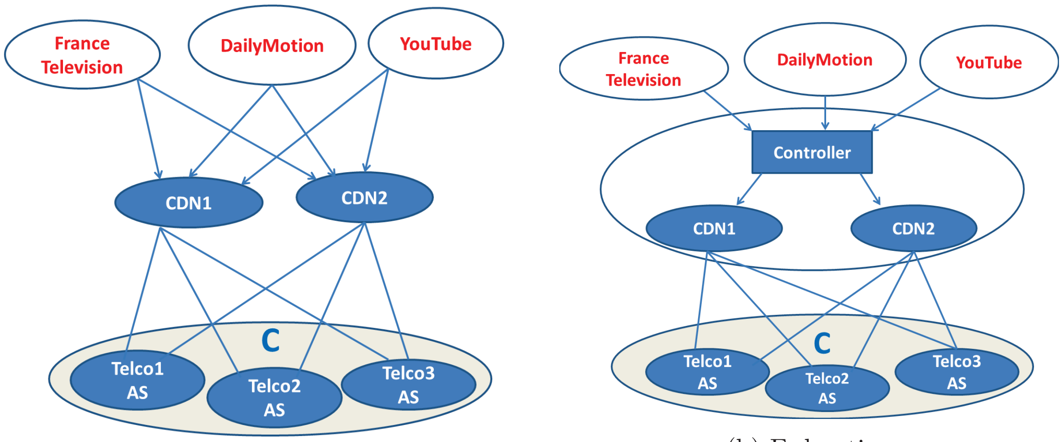

(a) Separate Scenario (b) Federation

Figure 5.6: Separate Scenario Vs Federation

Table 5.8: CDN_1 and CDN_2 Capacity Inputs

	p_1	p_2
Storage Capacity (GB)	50	250
Streaming Capacity (Mbps)	800	400

the CDNs to a given CP depends of the amount of traffic that this CP is expected to generate over C in one month. In order to calculate this price, the amount of expected traffic is mutiplied by a flat, per unit fee (price per GB) which is function of this amount. Table 5.4 shows the fees per GB that CDN_1 and CDN_2 respectively associate to different levels of clients commitments in terms of expected volume of monthly traffic. We assume that the charging model adopted by $\mathcal{S}$ is the same as the one that is separately adopted by CDN_1 and CDN_2. The subsequent prices charged monthly by $\mathcal{S}$ to YT, DM and FT over C are shown in Table 5.9.

5.3.2 Evaluation Methodology and Outputs

Given the use case inputs, we first address the separate scenario. We use the optimization model defined by equations (4.20), (4.21) and (4.22) in order to calculate the separate revenues of CDN_1 and CDN_2. These revenues are shown in Table 5.10. We figure out that, in the separate scenario, CDN_1 has an interest in handling YouTube demand over C while CDN_2 has an interest in handling DailyMotion demand over C.

Table 5.9: prices per CP over C (\$ / Month)

m	France (\$ / month)
YT	9000
DM	3840
FT	6240

Table 5.10: Separate Vs Federation Scenario

	Separate Scenario	Federation
Global Revenue ($ /month)	12840	19080
Global Cost ($ / month)	504	705.6
CDN_1 Revenue ($ / month)	9000	13675
CDN_1 Cost ($ /month)	378	504
CDN_2 Revenue ($ /month)	3840	5405
CDN_2 Cost ($ / month)	126	201.6
% of YT demand handled by CDN_1	100	80
% of YT demand handled by CDN_2	0	20
% of DM demand handled by CDN_1	0	57.18
% of DM demand handled by CDN_2	100	42.81
% of FT demand handled by CDN_1	0	68.74
% of FT demand handled by CDN_2	0	31.45

As a consequence, the storage cost paid by CDN_1 in the separate scenario is the product of YouTube catalog size and C_{stor}, the per GByte storage cost (C_{stor} was introduced in section 4.2.2). Similarly, CDN_2 storage cost is the product of DailyMotion catalog size and C_{stor}. On the other hand, the delivery cost paid by CDN_1 depends of the peak amount of traffic delivered by p_1 to different ISPs in C. If we assume an homogeneous structure of traffic among different ISPs (same shape of traffic evolution is observed at different ISPs level), the delivery cost paid by CDN_1 is the product of the peak amount of traffic reached at p_1 level, referring to YT daily peak over C, and C_{band}, the peering cost per Mbps (C_{band} was introduced in section 4.2.2. Similarly, the delivery cost paid by CDN_2 is the product of DM daily peak over C and C_band. The global costs paid by CDN_1 and CDN_2 in the separate scenario are shown in Table 5.10.

We then address the federation scenario. We assume a one zone footprint formed by C. We use the separate revenues of CDN_1 and CDN_2 together with the data in tables 5.1, 5.8 and 5.9 as inputs for the optimization model presented in Chapter 4. Once the model is solved, We use equation (4.18) in order to calculate the revenues that are respectively assigned to CDN_1 and CDN_2. We use the α_j^{mi} coefficients (equation (4.17)) in order to calculate the shares of YT, DM and FT demands over C respectively assigned to CDN_1 and CDN_2. The cost induced by content storage in CDN_1 and CDN_2 is calculated by using equation (4.2). The content delivery-related cost respectively paid by CDN_1 or CDN_2 is calculated by using equation (4.4).
Diffferent outputs are shown in Table 5.10.

5.3.3 Discussion

Given the outputs of the optimization model, the main conclusions that we reach are the followings:

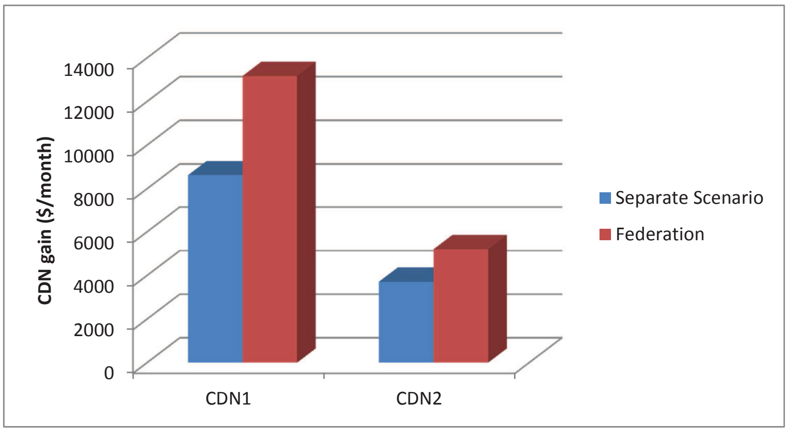

Figure 5.7: Separate Scenario Vs Federation gains

1. Through federating, CDN_1 and CDN_2 are able to jointly address the requirements of an additional client, FranceTelevision, with respect to the separate scenario. This translates in a higher revenue and a higher cost for both CDNs (see Table 5.10 for more details).

2. The demand of each of the CPs in $\mathcal{G}$ is distributed among both CDNs. CDN_1 and CDN_2 hence form a federation targeting the demand of YT, DM and FT over C.

3. The federation scenario is economically beneficial for CDN_1 and CDN_2. The gains achieved by both CDNs respectively in the separate and federation scennarios are shown in Figure 5.7. If we compare the separate and federation gains of each of the CDNs, we figure out that the gains of CDN_1 and CDN_2 are respectively enhanced by 50 % and 40 % with respect to the separate scenario. The difference between the federation and the separate gains of CDN_1 and CDN_2 allows quantifying the federation interest for each of the CDNs.

5.4 Use Case 3: Federation of pure-play and Telco-CDNs

In this section, we address the federation of Telco-CDNs and pure-play CDNs. Contrarily to use case 1, we assume that the footprint of a Telco-CDN is limited to this Telco subscribers. We also assume that an overlapping, in terms of footprint, exists between pure-play CDNs and Telco-CDNs. In particular, we consider a country C, two local Telcos that have CDNs in C and a pure-play CDN with surrogate servers that are connected, through direct peering points, to these Telcos networks.

Next, we explicit the different hypothesis that we consider, we solve the optimization model and we assess the economic gains achieved by Telco-CDNs and by the pure-play CDN through moving to a federation scheme.

5.4.1 Use case Presentation and Inputs

The use case that we are interested in is shown in Figure 5.8. We consider the same group $\mathcal{G}$ of CPs as for use case 1 and 2 and a one country footprint formed by C. The service requirements of different CPs in $\mathcal{G}$ over C are listed in section 5.1. The daily, per hour evolution of different CPs traffic over $\mathcal{C}$ is shown in Figure 5.2.

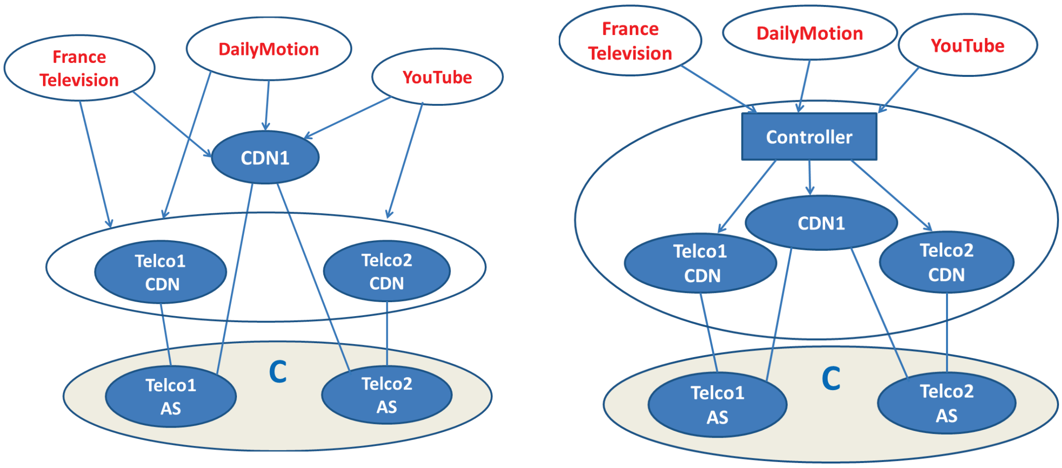

(a) Competition Scenario

(b) Federation Scenario

Figure 5.8: Competition Vs Federation Scenario

We assume that the Telco market in C is equally shared between two Telcos: $Telco_1$ and $Telco_2$. The market shares of the two Telcos also correspond to the D^i_m coefficients, referring to the distribution of different CPs demand among these Telcos subscribers. We also assume an homogeneous traffic structure among different ISPs in C. Based on it, the daily, per hour evolution of YT, DM and FT traffic over $Telco_1$ and $Telco_2$ ASs can be obtained through shrinking the curves in Figure 5.2 by a factor of 0.5.
We assume that $Telco_1$ and $Telco_2$ respectively own last-mile CDNs that are dedicated to their own subscribers. In other terms, $Telco_1$-CDN can only be used for delivering traffic to $Telco_1$ AS ($Telco_1$ AS is formed by the fixed subscribers of $Telco_1$). The same reasoning applies for $Telco_2$. We consider a system $\mathcal{S}$ of CDN providers formed by $Telco_1$-CDN, $Telco_2$-CDN and a pure-play CDN that has a cluster of servers that is directly connected to $Telco_1$ and $Telco_2$ networks. We denote by CDN_1 the latter CDN. We denote by p_1 the cluster of servers owned by CDN_1 in C. We split C into two zones: $Telco_1$ AS and $Telco_2$ AS. Table 5.11 shows the distance, in terms of intermediate ASs, between different CDNs in $\mathcal{S}$ and the two zones of C. Since the AS of a given Telco cannot be part of the footprint of another Telco-CDN, we assume that the distance between a given Telco-CDN and another Telco-AS is equal to three (according to equation (4.13),

Table 5.11: CDNs Footprint Info: Distance to target ASs

	$Telco_1$ CDN	$Telco_2$ CDN	CDN_1
$Telco_1$ AS	0	3	1
$Telco_2$ AS	3	0	1

Table 5.12: CDNs Capacity Inputs

	$Telco_1$ CDN	$Telco_2$ CDN	p_1
Storage Capacity (GB)	100	100	200
Streaming Capacity (Mbps)	400	400	400

a CDN that is within more than two ASs of a given zone cannot be used for delivering traffic to this zone). The capacity information of $Telco_1$ CDN, $Telco_2$ CDN and p_1 are respectively shown in Table 5.12.

As for use case 1 and 2, we assume that all CDNs in $\mathcal{S}$ adopt the same commitment-based charging model. According to this model, the minimum price charged by $Telco_1/Telco_2$ CDN to a given CP depends of the amount of monthly traffic that this CP is expected to generate over $Telco_1/Telco_2$ AS. In order to calculate this price, the amount of expected traffic is mutiplied by a flat, per unit fee (price per GByte) which is function of this amount. The price charged by CDN_1 to a given CP also depends of the amount of the monthly traffic that this CP is expected to generate over C. Table 5.4 shows the fees per GByte that different CDNs in $\mathcal{S}$ associate to different levels of clients commitments in terms of volume of monthly traffic. We assume that the charging model adopted by $\mathcal{S}$ over C is the same as the one separately adopted by different CDNs in $\mathcal{S}$. The subsequent prices charged monthly by $\mathcal{S}$ to YT, DM and FT over C are shown in Table 5.9.

Given these inputs, we aim at comparing two scenarios: the first scenario, denoted by the separate or competition scenario, refers to the case where $Telco_1$-CDN and $Telco_2$ CDN form a single CDN competing with CDN_1 over C. The second scenario, also denoted by the federation scenario, corresponds to the case where Telco and pure-play CDNs in $\mathcal{S}$ act as a single CDN and jointly target the demand of different CPs in $\mathcal{G}$ over C.

5.4.2 Evaluation Methodology and Outputs

Our methodolody is as follows: In order to assess the first scenario, we consider a two zones footprint formed by $Telco_1$ AS and $Telco_2$ AS. We position the requirements of different CPs in $\mathcal{G}$ over each of the zones with regard to the capacity of the Telco-CDN in $\mathcal{S}$ that operates this zone. We also position the overall requirements of different CPs in $\mathcal{G}$ over C with regards to p_1 capacity. This positioning is performed by using the optimization model given by equations (4.20), (4.21) and (4.22). We figure out that, in the competition scenario, CDN_1 has an interest in addressing DM demand over C while $Telco_1$ and $Telco_2$ CDNs have interest in jointly addressing FT demand originating from their respective ASs. The subsequent revenues assigned to different CDNs in $\mathcal{S}$ are shown in Table 5.13. These revenues correspond to the separate revenues of different CDNs in $\mathcal{S}$. The storage costs that are respectively paid by $Telco_1$-CDN and $Telco_2$-CDN are equal to C_{stor}, the cost per GByte, times the size, in GBytes, of FT content catalog. The storage cost paid by CDN_1 is equal to C_{stor} times DM catalog size. On the other hand, since $Telco_1$-CDN is within a distance of 0 of $Telco_1$ AS and $Telco_2$-CDN is within a distance of 0 of $Telco_2$ AS, the content delivery-related cost of both Telco-CDNs is null. If we assume an homogeneous structure of traffic between $Telco_1$ and $Telco_2$ ASs, the content delivery-related cost paid by CDN_1 will be proportional to the daily peak of DM over C. More precisely, this cost will be the product of C_{band}, the cost per Mbps, and the peak amount, in Mbps, of traffic generated by DM over C. The separate costs of different CDNs in $\mathcal{S}$ are shown in Table 5.13.

We then address the federation scenario. We use the separate revenues of different CDNs in $\mathcal{S}$ together with the data in tables 5.1, 5.9, 5.11 and 5.12 as inputs for the optimization model presented in the previous chapter. Once the model is solved, We use equation (4.18) in order to calculate the revenues that are respectively assigned to different CDNs in $\mathcal{S}$. We use the α_j^{mi} coefficients (equation 4.17) in order to calculate the shares of YT, DM and FT demands over C that are respectively assigned to $Telco_1$-CDN, $Telco_2$-CDN and CDN_1. The cost induced by content storage in different CDNs in $\mathcal{S}$ is calculated by using equation (4.2). Since, in this use case, no CDN traffic is circulating on the peering connection between $Telco_1$ AS and $Telco_2$ AS, the content delivery-related cost paid by $Telco_1$-CDN and $Telco_2$-CDN is null. The content-delivery related cost paid by CDN_1 for traffic delivery from p_1 to C is calculated by using equation (4.4).
Diffferent outputs of the model are shown in Table 5.13.

5.4.3 Discussion

Given the outputs of the optimization model, the main conclusions that we reach are the followings:

1. Through federating, CDN_1, $Telco_1$-CDN and $Telco_2$-CDN are able to jointly address the requirements of an additional client, YouTube, with respect to the separate scenario. This translates in a higher revenue and cost for each of the CDNs (see Table 5.13 for more details).

2. The overall demand of each of the CPs in $\mathcal{G}$ is distributed among the three CDNs. $Telco_1$-CDN, $Telco_2$-CDN and CDN_1 hence form a federation targeting the demand of YT, DM and FT over C.

3. The federation scenario is economically beneficial for the three CDNs. The gains achieved by different CDNs in the separate and federation scenarios are shown in Figure 5.9. We figure out that the federation gains of different CDNs are more important than the ones achieved in a competition scheme. In particular, the gains of $Telco_1$-CDN and $Telco_2$-CDN are enhanced by 60 % with respect to the competition scenario. We also notice that, independently of the considered

Table 5.13: Separate Vs Federation Outputs

	Competition Scenario	Federation
Global Revenue ($ /month)	13200	19080
Global Cost ($ / month)	144	220.72
$Telco_1$-CDN Revenue ($ / month)	4680	7550
$Telco_1$-CDN Cost ($ /month)	5	4.56
$Telco_2$-CDN Revenue ($ /month)	4680	7550
$Telco_2$-CDN Cost ($ / month)	5	4.56
CDN_1 Revenue ($ / month)	3840	3979
CDN_1 Cost ($ / month)	134	211.6
% of YT demand handled by $Telco_1$-CDN	0	45.39
% of YT demand handled by $Telco_2$-CDN	0	45.39
% of YT demand handled by CDN_1	0	9.2034
% of DM demand handled by $Telco_1$-CDN	0	46.01
% of DM demand handled by $Telco_2$-CDN	0	46.01
% of DM demand handled by CDN_1	100	7.965
% of FT demand handled by $Telco_1$-CDN	50	27.197
% of FT demand handled by $Telco_2$-CDN	50	27.197
% of FT demand handled by CDN_1	0	45.6

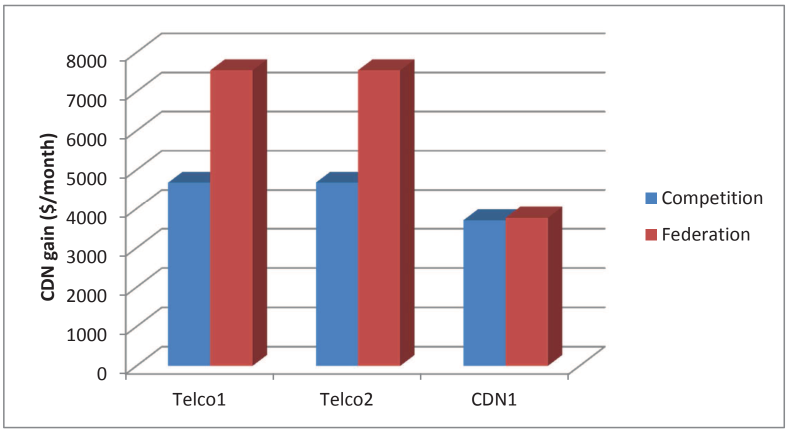

Figure 5.9: Competition Vs Federation gains

scenario (federation or competition), the gains of $Telco_1$-CDN and $Telco_2$-CDN are identical. This is due to the similarity, in terms of capacity, between both Telco-CDNs and, in terms of demand, between both Telco-ASs (same amount of demand originating from both Telcos subscribers). On the other hand, even though the federation scenario is beneficial to CDN_1, the gain of CDN_1 is not significantly enhanced through the federation. This is mainly due to the fact that CDN_1 is the only CDN in $\mathcal{S}$ with a non-null content delivery cost. According to the objective function of the optimization model (see section 4.2 for more details), CDN_1 is less priviliged in terms of load distribution with regards to $Telco_1$-CDN and $Telco_2$-CDN. CDN_1 federation revenue is consequently impacted.

Chapter 6

CDN Federation: Dynamic Decision-making

In Chapter 4, we explained how, based on statistics and logs referring to CPs requirements and on inputs on CDNs capacity and footprint, the controller can compute federations of CDNs on a monthly basis and provision these federations on a daily basis.
In this chapter, we focus on federation dynamic control, that is the phase taking place between two phases of federation provisioning. The dynamic control of a federation aims at ensuring a real-time routing of incoming requests inside the federation based on the provisioning phase output and, eventually, on a dynamic vision of different CDNs state. We first explicit the process of request routing within a federation of CDNs. We then focus on the control of peak events occuring at one or more CDNs level. In particular, we introduce dynamic frameworks for dealing with these events. We conduct trace-driven simulations in order to assess and compare the performance of different frameworks. The number of rejected sessions and the average video resolution experienced by end users are used as indicators in this context.

6.1 Request Routing

Requests routing is the sole operation performed during the dynamic phase of federation control. Let us consider the case of a CP which service request over a well-defined footprint has been accepted by a federation of CDNs. When a user in the CP footprint requests one of this CP contents, many steps are performed before the request is directed to the adequate CDN surrogate. Dependently of the naming scheme adopted by the CP, two schemes of request routing are likely to take place. We denote by $Scheme_1$ the first scenario and by $Scheme_2$ the second one. $Scheme_1$ and $Scheme_2$ are respectively shown in Figures 6.1 and 6.2.

Each content is referred to through a unique URL. The goal of request routing is to map the URL identifying a given content to an IP address corresponding to the surrogate that delivers this content to the user. The URL of any content has a well-defined structure that is proper to the authorative CP [44]. In addition to the domain name,

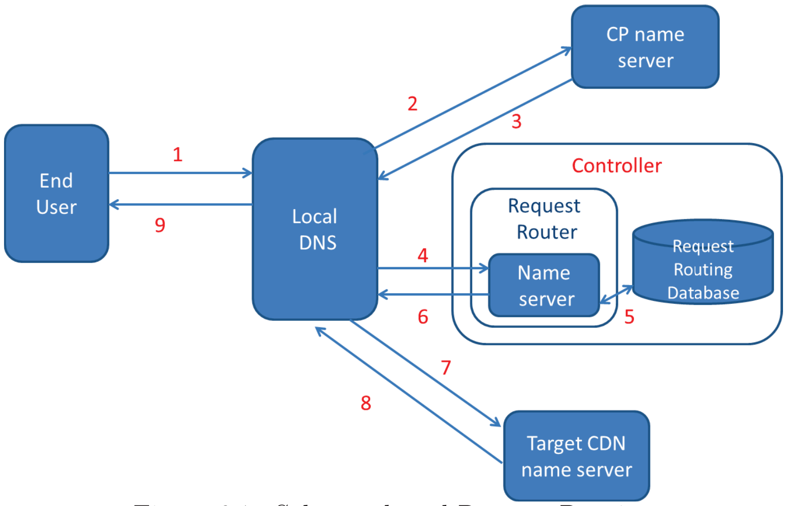

Figure 6.1: $Scheme_1$-based Request Routing

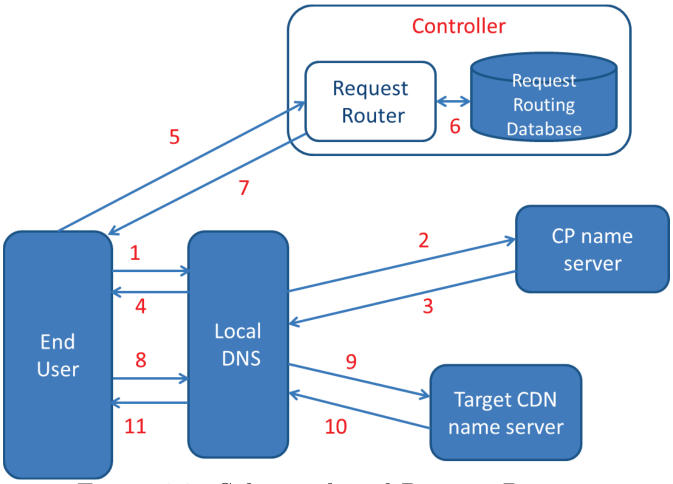

Figure 6.2: $Scheme_2$-based Request Routing

the structure may include a subdirectory and, optionally, a file name [44]. If the domain name in a given content URL is enough for identifying the requested content, requests routing takes place according to $Scheme_1$. If, on the other hand, the identification of a well-defined content requires going through the subdirectory and/or file name fields of its URL, requests routing takes place according to $Scheme_2$. This is mainly due to the fact that DNS only allows a domain name-based request routing.

When request routing takes place according to $Scheme_1$, it consists in the following steps (these steps are shown in Figure 6.1):

1. A DNS query is sent by the end user to its local DNS, that is the Telco DNS.
2. Based on the URL domain, the local DNS sends a DNS query to the CP name

server.

3. The CP name server responds to the query with a DNS response containing a CNAME field that refers to the controller domain.

4. Based on the received response, the local DNS sends a new DNS query to name server of the controller which, we suppose, is embedded in the controller *Request Router* engine (the architecture of the controller is described in Chapter 3).

5. The controller name server uses the IP address of local DNS in order to know the origin zone of the request and the URL in the DNS query in order to identify the target content. Based on it, it queries the *Request Routing Database* in order to identify the CDN towards which the incoming request should be directed.

6. The controller name server responds to the DNS query with a DNS response containing a CNAME field that refers to the target CDN domain.

7. Based on the received response, the local DNS sends a DNS query to the name server of the target CDN.

8. The target CDN selects a surrogate server towards which the request should be routed. The IP address of this surrogate figures in the DNS response sent by the CDN name server to the local DNS. The selection of the adequate surrogate depends of outsourcing policy of the target CDN. In particular, it takes into account elements like the location of the end user (can be estimated through the IP address of the local DNS) and the content URL.

9. The local DNS fowards the surrogate IP address to the end user. The user can now use this IP address in order to send a HTTP request to the surrogate.

When request routing takes place according to $Scheme_2$, it consists in the following steps (these steps are shown in Figure 6.2):

1. A DNS query is sent by the end user to its local DNS, that is the Telco DNS.

2. Based on the URL domain, the local DNS sends a DNS query to the CP name server.

3. The CP name server responds to the query with a DNS response containing the IP address of the controller to the local DNS.

4. The local DNS forwards the received response to the end user.

5. The user sends a HTTP request to the controller.

6. The request is intercepted by the *Request Router* engine that uses the user's IP address in order to identify the origin zone and the target URL in order to identify the target content. Based on it, it queries the *Request Routing Database* in order to identify the CDN towards which the request should be directed.

7. The controller sends a HTTP Redirection response to the end user where it specifies the target CDN as a the location of the content (URL referring to the target CDN domain is written in the location field).

8. Based on the received response, the user resends a DNS query to its local DNS with the URL received from the controller in the HTTP Redirection message.

9. The local DNS forwards the DNS query to the DNS name server of the target CDN.

10. The target CDN selects a surrogate server towards which the request should be routed. The IP address of this surrogate figures in the DNS response sent by the CDN name server to the local DNS. The selection of the adequate surrogate depends of the outsourcing policy of the target CDN. In particular, it takes into account elements like the location of the end user and the content URL.

11. The local DNS fowards the surrogate IP address to the end user. The user can now use this IP address in order to send a HTTP request to the surrogate.

6.2 Control of Peak Events

6.2.1 Context and Motivation

Federation provisioning is performed by using the optimization model introduced in Chapter 4. In particular, capacity provisioning is based on statistics that are either provided by CPs as part of their requirements and/or obtained through daily logs (statistics are summarized in Table 3.2). As seen in equation (4.15), these statistics are used in order to assess the peak amount, in Mbps, of traffic that is expected to be reached at different CDNs levels.

However, the real structure of a given CP traffic may, in some cases, not be inline with the statistics that he provided prior to federation computation ($Peak_m^t$ parameters). In particular, the traffic of video players like YouTube, Netflix and DailyMotion is often subject to flash crowds referring to an unexpected and rapid explosion of users' demand [40]. Flash crowds can be due, for instance, to unpopular or newly-generated contents turning popular. Due to their unexpected character, this kind of events cannot be taken into account upon federation provisioning.

On the other hand, when performing federation provisioning, we assume an homogeneous normalized popularity of contents over different time slots in a day. In particular, the peak amount of traffic at different CDNs level is expressed as a function of contents popularity (V_l^i parameters), this popularity being measured over a daily time scale (equation (4.15)). Nevertheless, the (normalized) popularity of video contents can be subject to important fluctuations in one day. This is particularly what we observe when we analyze the traffic traces of YouTube and DailyMotion collected in Orange France network (information on these traces was given in section 5.1).

Due to the above reasons, the capacity provisioning of established federations may become invalid during the provisioning cycle thus leading to overload scenarios at one or more CDNs level [76]. We refer to these scenarios as peak events. Non-managed peak events can lead to the degradation of users' QoE translating in a rejection of users' requests, high start times, an important degradation of video resolution and significant buffering scenarios.
Next, we propose network-based frameworks for dealing with peak events occuring within a federation of CDNs. the goal is to minimize QoE degradation caused by a low CDN performance. A first solution consists of a dynamic adaptation of static load balancing based on a real-time vision of different CDNs performance. As mentioned in section 1.3.3, some client-based load balancers including Conviva [54] and XDN [126] allow this kind of adaptation to take place. These systems monitor CDNs performance through aggregating, at very short time scale, QoE metrics provided by video players in users' terminals. Beyond its complexity, a users-based approach for monitoring CDNs performance present major drawbacks. First, this approach is reactive in the sense that events are detected after QoE degradations are perceived by an important set of users. Second, this approach does not provide a radical proof of events occurence. For instance, a failure in an ISP network can lead to a similar diagnostic in terms of users' QoE on a country scale.
We introduce a dynamic algorithm that allows adapting pre-defined load balancing rules within a federation of CDNs. Contrarily to industrial solutions, our algorithm assesses CDNs performance through monitoring, in real-time, the workload of CDNs over the countries of their footprint. A second solution consists of an intra-CDN control of peak events. In particular, each CDN dynamically adapts the resolution of its active sessions based on the evolution of its own workload. In this context, multi-bitrate streaming [9] should be supported by the CDNs. A third solution is a mix of the first two in the sense that a federation-based control of peak events is coupled with an intra-CDN control of these events.

We denote by $\mathcal{S}$ a federation of CDNs that is established and provisioned according to the model described in Chapter 4. We denote by $\mathcal{G}$ a set of CPs that form the federation clients. Next, we detail the different frameworks that we propose for controlling peak events inside $\mathcal{S}$. The notations figuring in Table 3.1 and Table 3.2 are used when describing different frameworks. We then conduct trace-driven simulations in order to assess the performance of different frameworks. The amount of sessions rejected bythe CDNs and average video resolution witnessed by end users are used as indicators in this context.

6.2.2 Framework 1: Adaptive Load Balancing-based approach

The controller initially uses the provisioning phase outputs in order to route incoming requests towards different j $\in \mathcal{S}$. Request routing takes place according to $Scheme_1$ or $Scheme_2$. Video sessions accepted by any j $\in \mathcal{S}$ are delivered with a high resolution for their view duration.

A CDN j $\in \mathcal{S}$ is said to be in overload state over a zone i of its footprint if the workload of different p $\in \mathcal{P}_j$ that cover i (are at less than two ASs of i) has reached a critical threshold referring, for instance, to a high percentage of p streaming capacity (e.g. $0.95 \times PC_p$). When a given j $\in \mathcal{S}$ is in overload state over a zone i, it can no longer deliver sessions to users in i. In this context, two scenarios of traffic offload are possible: either traffic offload is performed by j itself for all requests from i or the controller adapts static load balancing so that no requests from i are directed to j. The first scenario requires the existence of signaling interfaces between different j $\in \mathcal{S}$ so that j can be aware of the state and footprint of other CDNs $\in \mathcal{S}$. This scenario induces an extra redirection delay. The second scenario requires a real-time knowledge of the controller of the state of all CDNs $\in \mathcal{S}$ that cover i (CDN in overload state or in steady state over i). This scenario induces an extra control overhead for the controller. In both scenarios, an incoming request from i that is a priori destined to j is (re)-directed (by the controller or by j itself) to a randomly selected CDN $\in \mathcal{S}$ that covers i and that is not in overload state over i. If no CDN $\in \mathcal{S}$ matches these criteria, the request is rejected. This is typically the case when all CDNs $\in \mathcal{S}$ that cover a given zone are overloaded over this zone (Global Overload).
We assume that, when a given j $\in \mathcal{S}$ becomes overloaded over a given i $\in \mathcal{F}_j$, j redirects sessions from i towards other CDNs $\in \mathcal{S}$ that verify the previously-listed criteria. If j remains overloaded for more than a minute over i, it generates an overload event and sends it to the controller. The event includes elements like the time of occurence of the overload, the Id of the overloaded CDN and the Id of the impacted zone(s). The event is intercepted by the *Event Notifier* of the controller and is used in order to adapt static request routing decisions stored in the *Request Routing Database*. Once the event is taken into account by the controller, no future requests from i are directed to j. If, after a given time, j goes back to its steady state over i (workload of different p $\in \mathcal{S}$ is strictly inferior to pre-defined critical thresholds), j sends a cancelation of the overload event to the controller. Static policies of request routing are restored at the controller level when all j $\in \mathcal{S}$ go back to their steady state over different i $\in \mathcal{F}_j$.

6.2.3 Framework 2: Resolution Adaptation-based approach

This framework is enabled if multi-rate streaming [9] is supported by different CDNs in $\mathcal{S}$. Typically, for each m $\in \mathcal{G}$, each l $\in \mathcal{C}_m$ is characterized by many encoding bitrates corresponding to different video qualities or resolutions. In this context, many options are possible. One option consists of using different video streams in order to represent different resolutions of the same video. This option is adopted in the context of adaptive streaming over HTTP [116]. Another option consists of using a single video stream, named SVC (scalable video coding) stream, in order to represent a given video into different versions with different qualities/resolutions [114].
Both options can be used by the CDN in order to dynamically adapt the quality or resolution of an ongoing video session. Compared to the first option, SVC has the advantage of allowing a more optimized use of the CDN resources. This idea is deeply investigated in [111]. Given the advantages of SVC, we assume that differnt m $\in \mathcal{G}$ use SVC streams in order to represent their video contents. We also assume that different j $\in$

$\mathcal{S}$ are able to dynamically adapt the resolution of active video sessions through varying the number of SVC layers uploaded to end users [114]. For simplicity purposes, we suppose that any video content $l \in \mathcal{C}_m$ has two representations corresponding to two resolutions: a high resolution and a low resolution.
The controller always uses pre-computed, static request routing policies in order to direct incoming requests to different CDNs in $\mathcal{S}$. Request routing takes places according to $Scheme_1$ or $Scheme_2$. In the absence of overload scenarios, different CDNs in $\mathcal{S}$ deliver high resolution videos to end users in their respective footprints. When a given $j \in \mathcal{S}$ is in overload state over a zone $i \in \mathcal{F}_j$, it starts by degradating the resolution of a low percentage (e.g. 10 %) of the video sessions that are active over this zone. Sessions that are subject to quality degradation are chosen per seniority order meaning that the most recent sessions are degradated first. Newly arriving sessions are also delivered with a low resolution. If, despite quality degradation, j is again overloaded over i, it degradates the resolution of a higher portion (e.g. 20 %) of sessions that are active over i and so on. If the CDN is witnessing a very intense traffic surge over a given zone, it degradates the resolution of all sessions destined to this zone. If, despite global quality degradation, the CDN is again overloaded over the same zone, it rejects users' requests originating from this zone.
High resolution is progressively restored for sessions impacted by quality degradation when required streaming resources become available at the CDN level. Contrarily to resolution degradation, high resolution is first restored for the oldest sessions. A CDN $j \in \mathcal{S}$ goes back to its steady state over a zone $i \in \mathcal{F}$ when the load of different $p \in \mathcal{P}_j$ that cover i is inferior to a critical threshold (this threshold can be a high percentage of p streaming capacity) and when all sessions that are active over i are being delivered with a high resolution.

6.2.4 Framework 3: Mixed approach

This framework is enabled if multi-rate streaming [9] is supported by different CDNs in $\mathcal{S}$. The controller initially uses pre-computed request routing rules in order to direct incoming requests towards different CDNs $\in \mathcal{S}$. Request Routing takes place according to $Scheme_1$ or $Scheme_2$. In the absence of overload scenarios, sessions are delivered with a high resolution by different CDNs to end users in their respective footprints.
When a given $j \in \mathcal{S}$ is in overload state over a zone i of its footprint, it redirects, in a first time, requests from i to other CDNs $\in \mathcal{S}$ that cover i and are not overloaded over i. If j remains overloaded for more than a minute over i, it generates an event and sends it to the controller. The controller processes the event as for framework 1. Namely, it adapts request routing in order to prevent requests from i to be directed to j. If, due to an intense traffic surge from users in a zone i, all CDNs $\in \mathcal{S}$ that cover i are overloaded, request routing is performed inside $\mathcal{S}$ according to the outputs of the provisioning phase. Nevertheless, when intercepting new requests from i, each $j \in \mathcal{S}$ behaves as described in framework 2. Namely, j starts by degradating the resolution of some of the sesions that are active over i. If j continues to receive new requests from i, it degradates the resolution of a higher number of active sessions and so on. If, on the

other hand, streaming resources become available at j level, j progressively restores high resolution for sessions that are active over i. j is said to be in steady state over i when high resolution is restored for all sessions destined to i and when the workload of different $p \in \mathcal{P}_j$ that cover i drops under pre-defined critical thresholds. j sends a cancelation of the previously-generated overload event to the controller when it goes back to its steady state over i. Static request routing and high resolution video delivery are restored inside $\mathcal{S}$ when different $j \in \mathcal{S}$ go back to their steady states over their respective footprints.

6.2.5 Evaluation and discussion

In this section, we assess and compare the performance of the three control frameworks that we propose. We proceed as follows: we modify the code of the simulator described in section 5.1 in order to enable the three frameworks. The traffic inputs of the simulator are obtained through an Orange VoD traffic trace. The network parameters of the simualtor refer to a federation structure that we define (a number of CDNs with well-defined capacity and footprint). The simulator is expected to generate a number of outputs that allow assessing the performance of different frameworks. These include: time evolution of different CDNs workload, the number of sessions rejected by the CDNs of the federation and the average resolution witnessed by each user which session has been accepted by a CDN of the federation.
Next, we briefly describe the modifications that we introduced to the *OCEAN* simulator in order to enable the three frameworks. In particular, we present the traffic inputs as well as the network and simulation parameters that we use. We conduct trace-driven simulations and explicit the outputs that are reached when each of the frameworks is used as well as when no form of events control is supported inside the federation. We compare the different outputs and reach conclusions concerning the performance of different frameworks.

Simulator Overview

The basic version of the simulator was described in details in section 5.1. This version intends to evaluate and compare the performance of different caching algorithms.
We modify the code of the simulator in order to implement the control frameworks that we propose. In the new version, a network corresponds to a federation of CDNs and a network node is assimilated to a CDN cluster. Extra parameters including the streaming capacity and the footprint of the cluster are added to the network node description in the network file. The traffic file is obtained through converting a real traffic trace into the format that is supported by the simulator. This requires translating video URLs into movie Ids and users' physical addresses (mac address/ IP address) into users Ids referring to geographic/administrative zones. The traffic file should also specify, for each video request in the trace, the CDN towards which this request should be routed. The choice of target CDN is the output of the provisioning model presented in Chapter 4. The inputs of this model include, on the one hand, CDNs capacity and footprint information as figuring in the network file and, on the other hand, the service requirements of the

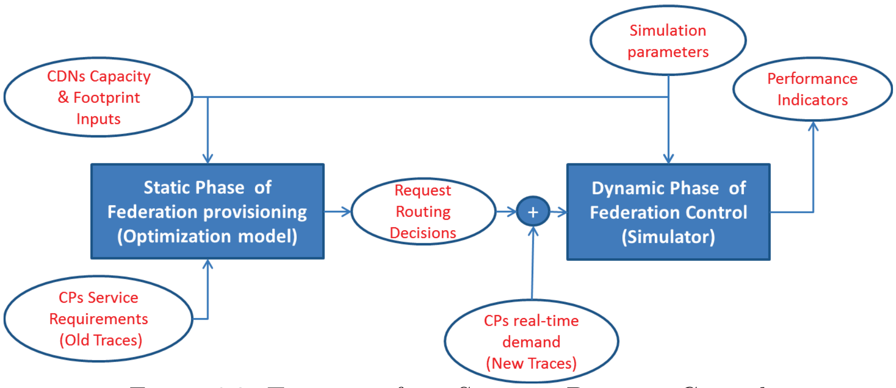

Figure 6.3: Transition from Static to Dynamic Control

Table 6.1: CDNs Capacity Inputs

	cdn_1	cdn_2
Storage Capacity (GB)	150	150
Streaming Capacity (Mbps)	100	50

CP(s) which video requests figure in the traffic trace. These requirements can be obtained through analyzing old traces corresponding to this/these CP(s) traffic as described in section 5.1.
The transition between the provisioning phase, enabled through the optimization model presented in Chapter 4, and the dynamic phase, enabled through the simulator, is shown in Figure 6.3.

Simulation Inputs

Federation Modeling We consider a federation formed by two CDNs, each composed of a single cluster, aiming to jointly address Orange VoD demand originating from fixed broadband users in Orange France network (Orange France AS). We denote by cdn_1 the first CDN and by cdn_2 the second one. Typically, the first CDN is likely to correspond to Orange France CDN and the second CDN is likely to correspond to Orange IBNF CDN. The first CDN is within a distance of 0 of the target footprint (Orange France AS) while the second CDN is within a distance of 1 of this footprint. The capacity information of cdn_1 and cdn_2 is shown in Table 6.1. We refer to these CDNs as $node_1$ and $node_2$ in the network file.

Traffic Dataset The traffic dataset that we use consists of a two days long Orange VoD traffic trace collected through probes placed in five BAS (Border access servers) in Orange France network.
This dataset is divided into two subtraces, a one day long each. The first subtrace is used for gathering statistics on Orange VoD traffic over Orange France AS. This subtrace is analyzed as described in 5.1. Statistics on Orange VoD demand are shown in Table 6.2. The structure of Orange VoD traffic over Orange France AS is shown in Figure 6.4.

Table 6.2: CPs Service Requirements

	Orange VoD
Nb of Requests/day/Orange AS	3400
Nb of Requested Objects	940
Mean object size (GB)	0.15
Content Catalog size (GB)	140
Mean object duration (Seconds)	600
Mean object bitrate (Mbps)	2
Peak Demand/Orange AS (Mbps)	130

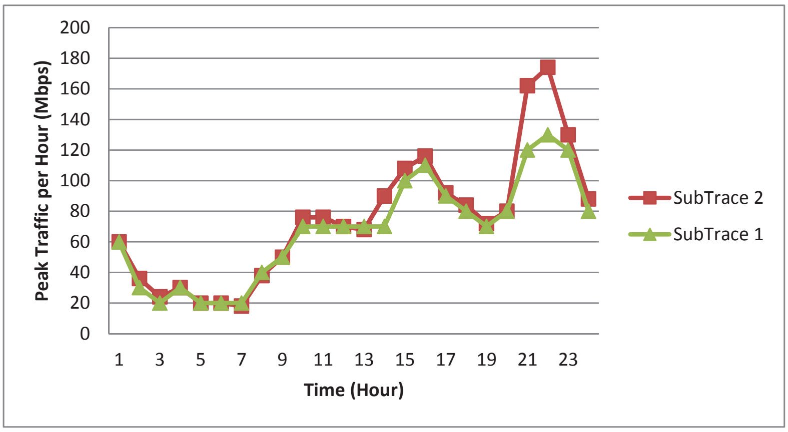

Figure 6.4: Expected Vs Real Evolution of Orange VoD Traffic

Together with cdn_1 and cdn_2 inputs, these statistics are used as inputs for provisioning the federation formed by cdn_1 and cdn_2. We assume that both CDNs adopt the same commitment-based charging model used in different use cases in Chapter 4. Solving the provisioning model allows calculating the request routing coefficients associated to different contents in Orange VoD catalog.

The second subtrace is used together with the computed request routing coefficients in order to generate the traffic inputs of the simulator. Namely, the request routing coefficients are used for deciding to which CDN each VoD request in the second subtrace should be routed.
We analyze the second subtrace. Obtained statistics are shown in Table 6.3. The structure of Orange VoD traffic over Orange France AS is shown in Figure 6.4. As can be noticed, a gap exists between the expected traffic evolution (obtained through the first subtrace) and the real traffic evolution (obtained through the second subtrace).

Table 6.3: CPs Real Demand Profile

	Orange VoD
Nb of Requests/day/Orange AS	3570
Nb of Requested Objects	940
Mean object size (GB)	0.15
Content Catalog size (GB)	140
Mean object duration (Seconds)	600
Mean object bitrate (Mbps)	2
Peak Demand/Orange AS (Mbps)	175

Furthermore, Orange VoD daily peak (175 Mbps) exceeds the (streaming) capacity of cdn_1 (100 Mbps), the (streaming) capacity of cdn_2 (50 Mbps) as well as the joint capacity of cdn_1 and cdn_2 (150 Mbps). Overload events are hence likely to occur at cdn_1 and/or cdn_2 level. Next, we assess how different frameworks react to these events.

Simulation parameters In addition to network and traffic files, simulation parameters include the simulation time step referring to the pace at which incoming requests are processed, the output interval referring to the pace at which measurements are performed at CDNs level as well as the mean duration and bitrate of an Orange video.We consider a time step and an output interval of 1 second. Based on the analyzis of the traffic dataset, we consider that, in average, an Orange video lasts for 600 seconds. Since resolution adaptation is used in two of the frameworks that we propose, we consider that an Orange video is characterized by two bitrates. The first bitrate corresponds to a high resolution of the video and is equal, in average, to 2 Mbps. The second bitrate corresponds to a low resolution of the video and is equal, in average, to 1 Mbps. These parameters are shown in Table 6.4.

Table 6.4: Simulation Parameters

Time Step (Seconds)	1
Output Interval (Seconds)	1
Mean Session Duration (Seconds)	600
High Resolution Bitrate (Mbps)	2
Low Resolution Bitrate (Mbps)	1
Mean object duration (Seconds)	1

Simulation Results

Given the above inputs, we begin by assessing the scenario where no intra-CDN or inter-CDNs control of peak events is performed in the federation. In this case, static rules are used for balancing Orange VoD load among the two CDNs and accepted sessions are delivered with a high resolution for their view duration. We figure out that, in this scenario, 160 requests out of the 3570 requests issued by Orange broadband users in one

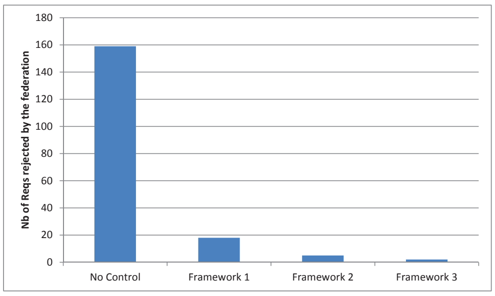

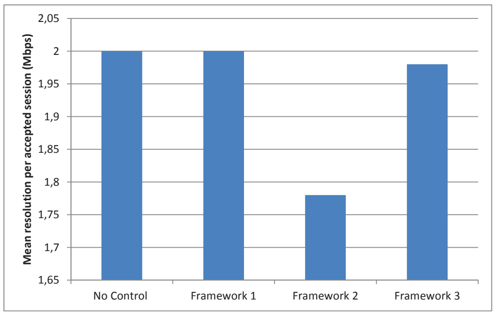

(a) Nb of sessions rejected by cdn_1 & cdn_2 (b) Mean Video Resolution

Figure 6.5: Performance Results

day are rejected by cdn_1 and cdn_2.

When adaptive load balancing is enabled (Framework 1), the number of rejected sessions per day drops to 18. Since no adaptation of video resolution takes place, the average resolution per accepted session is equal to 2 Mbps.
When multi-bitrate streaming is enabled at cdn_1 and cdn_2 level (Framework 2), the number of rejected sessions is equal to 4. In fact, the time gap between the occurence of an overload event and the degradation of sessions resolution causes requests directed in the meantime to the concerned CDN to be rejected. The average resolution per accepted session is calculated as follows: For each of the sessions accepted by cdn_1 or cdn_2, we track the resolution of the session at each of the time steps when it is active. This resolution refers to the encoding bitrate of video packets delivered at a well-defined time step. We sum the obtained values and we divide by the session duration. We hence obtain the average resolution witnessed for a well-defined video session. We sum the average resolution of each of the accepted sessions and we divide by the amount of accepted sessions. Using this methodology, we calculate the average resolution provided by cdn_1 and cdn_2 to an Orange broadband user when Framework 2 is adopted. A value of 1.78 Mbps is obtained. In particular, 35 % of the sessions accepted by cdn_1 and cdn_2 witness a low video resolution for more than one second during their view duration.

When a mixed approach for events control is adopted (Framework 3), the number of sessions rejected by cdn_1 and cdn_2 drops to 2. Furthermore, the average resolution per accepted session is equal to 1.98 Mbps. Only 17 % of the accepted sessions witness a low resolution for at least one second during their view duration. Performance results, in terms of number of rejected sessions and mean resolution per accepted video session, of different control approaches are respectively shown in Figures 6.5.a and 6.5.b.

Discussion

Based on the obtained results, we assess the gains of a federated approach for events control (adaptive load balancing). We notice that, when cdn_1 and cdn_2 do not support multi-bitrate streaming, this approach allows reducing from 160 to 18 the number of sessions rejected by cdn_1 and cdn_2 (comparing Framework 1 to the scenario where no event control is performed). cdn_1 and cdn_2 joint hit ratio (defined as the number of sessions accepted by cdn_1 and cdn_2 divided by the number of sessions received by cdn_1 and cdn_2) thus increases from 95.5 % to 99.5 %Ṡimilarly, when cdn_1 and cdn_2 support multi-bitrate streaming, this approach allows a slight increase of the federation hit ratio (moving from a hit ratio of 99.85 % for Framework 2 to a hit ratio of 99.95 % for Framework 3) and enhances by 12 % the mean resolution experienced by end users (moving from a average video resolution of 1.78 Mbps for Framework 2 to an average resolution of 1.98 Mbps for Framework 3).

In summary, when a federated approach for events control is adopted within a federation of CDNs, the federation is better resilient to peak events. This translates into a higher hit ratio of the federation and a better video resolution witnessed by end users. In addition, coupling an inter-CDNs form of control (adaptive load balancing) with an intra-CDN form of control (intra-CDN resolution adaptation) is required when a federation is subject to intense peak events. A mixed approach for events control is indeed key for keeping the hit ratio of the federation very close to 1 while ensuring to end users a video resolution that is very close to premium video resolution.

Chapter 7

Telco positioning in the CDN ecosystem

In this chapter, we address the positioning of the Telco in the CDN eco-system. In particular, we focus on the roles that the Telco can play and on added-value services that it can provide in a context of CDN federation. Since new Telco services cannot be enabled without an adequate control architecture, we describe the existing control plane of the Telco. We suggest modifications of the Telco control plane and novel network APIs in order to enable the services that we propose. We illustrate the operation of the new Telco control plane via a concrete use case.

7.1 Overview of the Telco assets

The Telco owns many assets that differentiate it from the other players of the content provisioning value chain. These can be summarized as follows:

Proximity The Telco is the first point of contact for end users. Indeed, except for some roaming [16] and Mobile Virtual Network Operator (MVNO) [49] scenarios, the subscribers of a given Telco first pass through this Telco network in order to access the rest of the internet including CPs and pure-play CDNs. As a first point of contact for end users, the Telco owns a good knowledge of users' context and access conditions. The Telco can also maintain content-related user history. This history can be used by CPs for customizing their portals and advertisements. The Telco is also well-placed to take decisions of routing of user-generated traffic or requests throughout the internet. This asset allows decreasing request routing delays in a CDN federation context. This idea will be further investigated in the coming sections.
The Telco is also seen as the last-mile of content delivery in the internet. Indeed, OTTs contents cannot be pushed towards end users further than the Telco domain. Different scenarios of content placement exist within the Telco domain. These include transparent caching, Telco-CDN and in-network caching paradigms (ICN, NGPoP...). Many OTTs

are aware of the impact of last-mile content delivery on the QoE perceived by end users. Therefore, they are pushing their contents towards the Telco domain as described in section 2.4.

Knowledge of user context and information Telcos maintain in their information Systems and network equipment (e.g. Home Subscriber System (HSS [27])) valuable user information including authentication keys, user identities and user service profiles. This information enables them to perform a number of control functions including user authentication, authorization of user access to services and billing on behalf of CPs and CDN providers.
Telcos monitor, in real time, mobile user contexts in terms of geographic location, device type and available access networks, among others. Providing this information to CPs and CDN providers enables the adaptation of both content portals and content resolution to user current contexts. Portals for instance can be adapted to user location and device capabilities. Furthemore, the format and the resolution (encoding bitrate) of a selected content can be also adapted to the constraints, in terms of bandwidth, of user devices and accesses. Some may argue that the Telco role is not primary at this level. In fact, Telco independent approaches like HTTP adaptive streaming [116] already allow this kind of adaptation to take place. Furthermore, geolocation and other capabilities provided by many terminals are likely to provide CPs/ CDN providers with enough information. Thus, there is a moving equilibrium between relying on the Telco for providing context-related data to 3rd parties and counting on the terminal for doing so. Beyond alleviating the complexity from the terminal side, the former approach has two major advantages. Contrarily to the terminal, the Telco has a real vision of its access and bakcbone networks and does not rely on heuristics in order to predict these networks state. The Telco can also triggers actions like bandwidth allocation and paths enforcement inside its network. Finally, Telcos can play an important role in handling vertical mobility referring to scenarios of mobility that involve a change of the service IP address (Home address) of the end user. As a last mile ISP, the Telco can directly track a change of the service IP address. When notified of this change, the CP/CDN can perform actions including flows re-rerouting and content format/resolution adaptation in order to ensure a seamless delivery of contents to end users.

Network Monitoring and Control Through being able to monitor their own networks, Telcos can gather information about both network state and transported traffic flows. Network state monitoring (network topology, links capacities, routers load, QoS metrics, etc) and better understanding of the traffic structure might be helpful for enhancing CDNs performance. From a static topological point of view, Telcos can inform a CDN provider where to distribute, on a national or regional scale, its various surrogates. In addition, Telcos can ensure a short-cut path to OTT contents through providing direct connectivity services to OTTs. Peering agreements established between OTTs and Telcos illustrate this trend. Telcos can also provide local breakouts to OTTs surrogates. This scenario is particularly interesting in the mobile architecture where the first Telco

router providing entry to the IP network is relatively centralized (e.g. a limited number of centralized GGSNs exists in the Orange France backbone).
On the other hand, dynamic monitoring of content flows allows Telcos to gather and aggregate statistics on CPs demand originating from their mobile and fixed subscribers. Statistics include volume of daily traffic, contents popularity, traffic structure etc. These statistics significantly help CDN providers in provisioning their CDNs.
Traffic monitoring allows Telcos to be aware of both source and destination of content flows. Based on it, Telcos can perform a flow-based charging of users on behalf of CPs/ CDN providers. Finally, network monitoring allows the Telcos to track the evolution of bandwidth in their access and core networks. Operations aiming to prevent network bottlenecks and to enhance routing efficiency can be performed.

In addition to network monitoring, Telcos have full control of their network resources and can, potentially, provide bandwidth on demand. Users and network state information, as well as real-time requirements from 3rd parties (e.g. CDNs), can be astronomically processed, based on pre-defined policies, in order to optimize resources allocation. In this context, Telcos can propose QoS related SLAs to CPs and CDN providers thus adding value to these players services. Established SLAs can be enforced through many mechanisms. One mechanism consists in using VPNs. Others involve new architectures like SDN.
Concerning mobility support, Telcos can use their mobility tracking capability in order to re-route content flows toward new user contexts (accesses, devices, IP addresses, etc.). Finally, in-network caching can be seen as an option of network control that can be activated by the Telco upon network monitoring. Typically, the Telco can decide to activate in-network caching for some of the circulating flows when detecting bottleneck events inside its network.

7.2 Overview of potential Telco Roles

In the light of the previously-listed assets, the Telco can play three roles with respect to a federation of CDNs. As a last-mile CDN provider, the Telco can participate through its CDN to a larger federation. The Telco can also play the role of a federation controller. Finally, the Telco can choose not to take part in the federation neither as a member nor as a controller. Nevertheless, it can still propose one or many added-value services to the federation and hence plays the role of an added-value services provider.
Next, we detail the three roles of the Telco and the service(s) and functions associated to these roles. It is important to note that the Telco can play many roles at once. For instance, the Telco can be a member of a CDN federation and its controller. Similarly, the Telco can contribute to the federation through its CDN platform and through added-value services that pure-play CDNs cannot necessarily provide.

7.2.1 CDN provider

As a last-mile CDN provider, the Telco can contribute to a federation through its vacant storage and streaming capacity. The outputs of the federation provisioning phase are used in order to decide which contents should be pro-actively pushed towards the Telco domain and which requests should be routed, in real-time, towards the Telco. As a CDN provider, the Telco usually deliver OTTs contents to its own subscribers. A Telco-CDN can be dedicated to fixed subscribers or to mobile subscribers or can be mutualized between the two. Nevertheless, if local Telcos in a given country federate (this use case is addressed in section 5.2), the Telco footprint is extended to all fixed and mobile subscribers in the country/countries where the Telco is present.
In order to take part of a federation of CDNs, the Telco should provide to the controller different kinds of information at different time scales (information is listed in different sections of Chapter 3). These include information about its footprint referring to its different PoPs, its vacant capacity, its cost model, its access and demand logs and its workload state (overload state Vs steady state).

7.2.2 Federation Controller

The Telco can play the role of a federation controller whether it is a part of a federation or not. In order to play this role, the Telco implements the functional architecture described in Chapter 3. In addition to federation computation, provisioning and dynamic control, the Telco can use user authentication information in order to authorize its subscribers access to OTTs contents hosted in any of the CDNs of the federation. As part of its role as a federation controller, the Telco routes incoming users' requests to different CDNs of the federation. As a first point of contact of end users, the Telco is the first to intercept requests issued by its subscribers. Based on requests URLs, the Telco can identify the requests destined to any of the federation CDNs. The Telco hence uses pre-computed load balancing rules in order to identify the target CDN for a given request. It then immediately directs the request towards this CDN. Limiting the request routing process to the Telco domain significantly decreases the number of DNS and HTTP redirections observed in the two schemes of request routing described in section 6.1. As a consequence, a better experience, in terms of time to first byte, is likely to be witnessed by the Telco subscribers.

7.2.3 Added-value services provider

Beyond being a controlller of the federation or one of its members, the Telco can be an outsider to the federation while still proposing B2B added-value services to the federation. These added-value services are mainly based on the different assets of the Telco listed in the previous section.
Next, we introduce three added-value services that the Telco can propose to a federation of CDNs. Two of the services can also be proposed by the Telco to single OTTs. an OTT or a CDN federation can subscribe to the entirety or to a subset of the added value

services. Furthermore, different added-value services can be proposed by the Telco to third parties in conjection with other services related, for instance, to its role as a CDN provider.

User Authorization

The footprint of an established federation can be divided into many zones, a zone referring to an autonomous system (AS) operated by a Telecom operator. The federation controller can delegate to each Telco operating a zone of the federation footprint the authorization of its subscribers access to different CDNs of the federation. In particular, when a given Telco intercepts a request targeting a client of the federation, it uses the authentication information in user repositories (e.g. HSS in the mobile context) in order to authenticate the requesting user. Based on it, the Telco either authorizes the user's access to the requested content or simply drops the request.

Delegating the authorization function to the Telco has the advantage of relying on information already existing at the Telco side (e.g. HSS) in order to alleviate some of the complexity from the federation side. In particular, if this function was not performed by the Telcos, it should be performed by the controller or by the CDNs that intercept the different requests.

User authorization, as an added-value service, can equally be proposed by a Telco to single OTTs that can be CPs and pure-play CDNs. This is already the case of Vente Privee [123] that delegates to Orange the authorization of Orange users access to its website.

Direct Request Routing

An important aspect of control of a federation of CDNs consists in routing, in real-time, incoming users' requests to the appropriate CDN in the federation. As explained in Chapter 6, request routing is based on outputs of the provisioning phase and on eventual adaptations of these outputs upon detection of peak events.

The speed at which the routing of incoming requests to adequate CDN occurs directly impacts the time to first Byte perceived by end users. Thus, ensuring a better time to first byte consists in minimizing the number of (HTTP and/or DNS) redirections between the first DNS query issued by the user to its local DNS and the final HTTP request between the user and the server that is actually delivering its content.

We divide the federation footprint into many zones, a zone being composed of a Telco AS. As stressed earlier in this section, Telcos can play an important role in fastening the process of request routing within a federation of CDNs through enabling a immediate routing of requests to target CDNs. In this context, the controller should provide information about its routing policy to the Telcos operating different zones of the federation footprint.

In particular, the controller should provide to a given Telco a part of the data contained in its *Request Routing database*. This includes decisions of request routing concerning requests originating from this Telco AS and destined to any of the federation clients.

When receiving this data, the Telco configures its DNS server in order to directly route selected DNS queries (queries for URLs belonging to the federation clients) to the name servers of the adequate CDN. Given the fact that request routing requires an analyzis of the full content URL and that DNS lookup does not always allow this kind of analysis to take place (DNS lookup is based on the domain URL), a potential alternative consists in using a dedicated proxy at the Telco level. Typically, the Telco can direct to this proxy all HTTP requests destined to a federation of CDNs. The proxy analyzes the full content URL and, based on the request routing information provided by the controller, directs the requests of Telco subscribers towards adequate CDNs.

End to End QoS Optimization for mobile users

When consuming video content, end users are particularly sensitive to three QoE metrics: The Time to first Byte, the buffering ratio and video resolution. More information on these metrics and their impact of users' short and long term engagements can be found in section 2.3. More demand on video traffic is originating from mobile devices which raises many challenges related to the geographic and access mobility of users as well as to bandwidth fluctuations in wireless access networks (access networks can be cellular or non-cellular). The specifity of the mobile context was investigated in section 2.3.

Due to the above reasons, ensuring a good QoE for end users, especially mobile ones, requires, beyond the optimization of CDN selection within a federation of CDNs (addressed in Chapter 6), an end to end optimization of the QoS on the path between CDNs and end users. The latter optimization has a double goal. First, it allows taking video resolution into account upon enforcing network control decisions. Decisions include the selection of the content path and the allocation of bandwidth resources at this path level. Second, it allows a dynamic, real-time adaptation of delivered video resolution in the light of the evolution of the content path (change of the user access network, degradation of the network conditions...).

As more video players and pure-play CDNs push video contents closer to end users (this trend is described in section 2.4), the path between the CDNs of a federation and consumers is shrinking. Telcos are controlling an important part, and sometimes the entirety, of this path. This is typically the case when OTTs contents are placed inside the Telco domain. This new strategic position of Telcos make them well-placed to optimize end to end QoS between the CDNs of a given federation and their own subscribers, mainly mobile ones. In order to optimize end to end QoS on content paths, Telcos can primarily count on their assets in terms of network monitoring and control and on their knowledge of the access context of mobile users (available accesses, access conditions...).

As for other added-value services, the federation controller can delegate to different Telcos operating different zones of the federation footprint the optimization of end to end QoS on paths between the CDNs of the federation and their respective subscribers. A given Telco can perform QoS optimization for all sessions initiated by its mobile subscribers and targeting any of the federation clients. When subscribing to this added-value service, a federation of CDNs is likely to enhance its overall performance through ensur-

ing a better QoE to end users in the federation footprint.
Optimization of end to end QoS, as an added-value service, can be equally proposed by a Telco to single OTTs. These include pure-play CDNs and CPs that are either directly peering with the Telco or using the Telco-CDN for terminating their traffic over the Telco AS.

7.3 Overview of the Telco Control Plane

The control plane plays a key role in enabling the different roles of the Telco explicited in the previous section. In fact, an adapted control plane identifies the entities impacted by a well-defined service and orchestrate these entities operation in order to perform this service. A control infrastructure is usually composed of various functional groups and interfaces among them plus a set of APIs that facilitate their usage by other system components.
In this section, we introduce control plane entities and architectures, standardized, in the phase of standardization or implemented by Telcos, that we find the most relevant for enabling the roles of the Telco that we described. In particular, we focus on components and architectures that we find useful for implementing the three B2B added-value services that we propose.

Policy and Charging Control (PCC) Policy Charging and Control (PCC) is a service-aware control architecture defined by 3GPP. It provides network operators with standardized and advanced mechanisms for controlling QoS and charging in the mobile context [22]. PCC mechanisms are applicable to different type of services and applications (conversational, streaming...).
PCC architecture and features have evolved throughout 3GPP specifications. In particular, PCC architecture has evolved in order to enforce policy control and charging over a multitude of domains including General Packet Radio Swithcing (GPRS) [21] and Evolved Packet System (EPS) [28]. PCC has also evolved in order to support new features and functionalities including: Mutiple access technologies, roaming and mobility [22].

A generic sketch of the PCC architecture, as defined in 3GPP Release 11, is shown in Figure 7.1. The main components of the PCC architecture are the followings (detailed information about these components can be found in [22] and [92]):

- Application Function (AF): The AF interacts with applications that require dynamic policy and charging control. It extracts session information and provides this to the policy and charging rules function (PCRF) over the Rx reference point. The AF also can subscribe to certain events that occur at the traffic plane level. These include IP session termination or access technology-type change. When the AF has subscribed to a traffic plane event, the PCRF informs the AF of its occurrence.

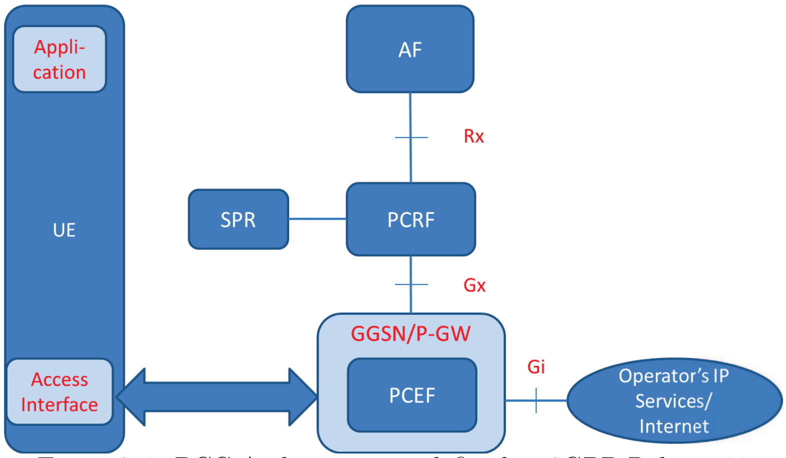

Figure 7.1: PCC Architecture as defined in 3GPP Release 11

- Policy charging and control function (PCRF): The PCRF is the policy engine of PCC. Its main role consists in computing session-level policy decisions that are provided to the Policy Charging and Enforcement Function (PCEF). In order to compute policy decisions, the PCRF makes use of one or many of the following inputs: Application specific inputs received over the Rx reference point and referring to the service requirements of a given application, network input received from the Gx reference point, user subscription data stored in the Subscriber Profile Repository (SPR). The decisions computed by the PCRF may also be based on its inner policy. These decisions allow two kinds of network control: gating control and QoS control. Gating control is the capability to block or to allow IP packets belonging to a certain IP flow. The PCRF could, for example, make gating decisions based on session events (start/stop of service) reported by the AF through the Rx reference point. QoS control allows the PCRF to provide the PCEF with the authorized QoS for a given IP flow. The authorized QoS includes, among others, the authorized QoS class and the authorized bit rates.

- Policy charging and enforcement function (PCEF): The PCEF enforces policy decisions received from the PCRF and also provides the PCRF with user and access-specific information over the Gx reference point. The PCEF can also perform online charging through interacting with the Online Charging System (OCS) [25] and offline charging through reporting usage of resources to the Offline Charging System (OFCS) [24]. This reporting allows the operator to charge third parties (B2B, B2B2C like charging models) and/or end users (B2C charging model) at the origin of this usage.

Opening PCC to third parties As seen in Figure 7.1, the Application Functions (AFs) are interfaced with the PCRF through the standard Rx interface, specified in [22]. This interface is based on the Diameter protocol [48], which is widely used in the telecom

world, but less in the web service ecosystem. So new protocols are currently considered by the industry to interface the PCC with third party service providers. Through opening PCC architecture, network operators can generate new sources of earnings, from the third party service providers (OTTs). In parrallel, third party service providers can benefit from an improved quality of service, to differentiate from other services or monetize themselves this improved QoS towards their own users (in the B2B2C scenario).
Basically, two main architecture possibilities can be considered to open the PCC architecture:

- Direct web interface from the third party service provider to the PCRF
- An interface based on the introduction of a new entity: the Broker

The first alternative already exists in some vendors solutions, and it seems to be a shorter-term-solution. However, this simple solution also has drawbacks, mainly because it makes PCRF more complicated and directly accessible to third parties.
In recent contributions to the Open Mobile Alliance (OMA) [96], many Telcos including Orange advocated the second alternative. In particular, they agreed on the fact that a new entity, identified as a broker, should be introduced between third party application functions and PCC. According to these Telcos, the introduction of an intermediary broker has the following advantages:

1. Allow the mobile operator to offer QoS services capabilities to 3rd-party service providers. These services take the form of B2B and B2B2C offers for a differentiated QoS.
2. Allow the mobile operator offering an easy-to use QoS API for the 3rd-party service providers. This requirement is achieved through a simplified setting of QoS parameters.
3. Allow the mobile operator offering QoS services capabilities that embrace all network assets. Those assets are not limited to PCC and therefore cover other Gi-located network entities.
4. Allow the mobile operator to strengthen direct partnership with 3rd-party service providers and to weaken the risk of disintermediation that is likely to emerge with built-in device solution.

Telcos contributions to the OMA are still in their early stage and did not go, till now, beyond supporting the introduction of a broker between third parties and the PCC. In particular, no concrete services and APIs which make use of the Broker have yet been standardized in the OMA.

Access Network Discovery and Selection Function (ANDSF) The Access Network Discovery and Selection Function (ANDSF) [26] is a control entity that has been optionnally introduced in EPS [28]. When present, this entity enables the UE to manage

in an easiest way the connectivity and mobility in EPS. Since EPS allows user equipment (UE) simultaneous connectivity to a multitude of 3GPP and non-3GPP access networks, the ANDSF can intervene whether the UE is connected to a 3GPP access and to a non-3GPP access or whether it is connected to many non-3GPP accesses at once.
The ANDSF entity has different goals. First, it provides the UE with data about available access networks during the attachment procedure. Data includes access networks type and identifiers. It also enables the UE to choose an access network from several accesses through which a well-defined data flow shall be routed. ANDSF also interferes when handover or reselection procedures take place. In order to select new accesses or connections, the UE relies on some types of information provided to it by the ANDSF entity. For instance, the ANDSF provides the UE with information that allows it to decide whether the mobility is allowed or not and to select the most preferable access type or technology. In addition, The ANDSF information provides the UE with a sufficient idea about new access networks including these networks types and identifiers.

IP Multimedia Subsystem (IMS) IMS has been introduced by 3GPP as a signaling architecture which aims to control the delivery of multimedia services to fixed and mobile end users. IMS decouples the control plane from the underlying data plane. In addition, IMS is agnostic to the access technology and interfaces with application servers entities in order to provide any type of application requiring session control [17] [27].
IMS relies on a signaling protocol called Session Initiation Protocol (SIP) [110]. A SIP message may contain different body types including a Session Description protocol (SDP) [69] body. SDP describes multimedia content sessions and allows a negotiation of the media type and format between the parties involved in the service. Other protocols other than SIP (including Diameter [48] and RTSP [113]) are used at some of IMS interfaces. IMS overall design is illustrated in Figure 7.2.
The core of the IMS architecture consists of the 'Call Session Control Function' bloc that includes three functional groups: the P-CSCF, S-CSCF and I-CSCF (proxy, serving and interconnecting CSCF). IMS application layer includes a number of 'Application Servers' (AS), referring to services implemented on top of IMS. IMS application servers can be or not SIP-based. Entities like HSS and PCRF do not natively belong to IMS. However, these are considered as part of the IMS design since they are accessed, for control purposes, by core IMS entities. In particular, the P-CSCF entity in IMS can provide to the PCRF, over the Rx interface, SDP related information thus playing the role of PCC AF. This information is used by the PCRF in order to perform QoS control at the traffic plane. Beyond being access and service agnostic, IMS provides a bunch of functionalities including integrated user authentication authorization and charging, session mobility control, bearer QoS control, media negotiation and Presence. More information on IMS design and functionalities can be found in [17] [27].
Even though IMS presents many interesting concepts and functionalities, it remains a highly complex architecture which relies on a protocol that is not widely supported in the web services ecosystem. The drawbacks of IMS with regards to the CDN ecosystem are deeply investigated in [71]. Therefore, we do not believe that IMS is the most adequate

control architecture for enabling Telco-based added-value services and APIs to third party players in the CDN ecosystem (pure-play CDNs, CDN federations, CPs...).

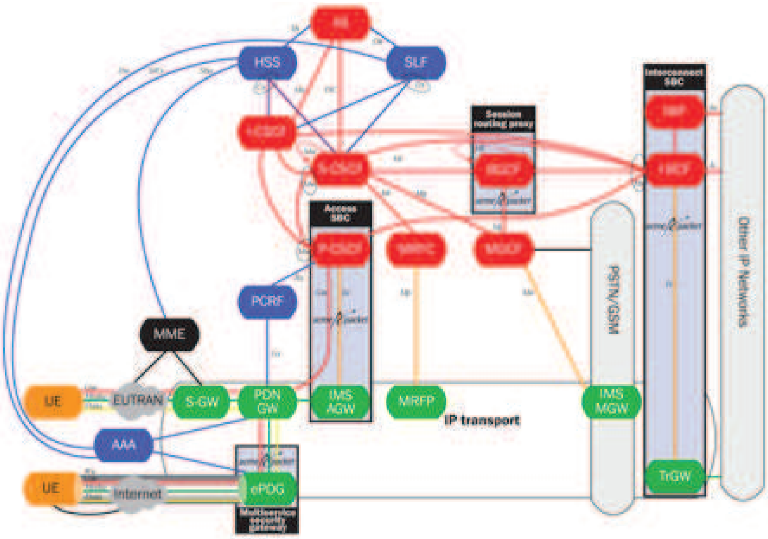

Figure 7.2: IMS Design as defined in 3GPP Release 8

7.4 The Evolution of the Telco Control Plane: A Roadmap

In this section, we use elements of the Telco control infrastructure presented in the previous section in order to suggest enhacements of the Telco control plane so that the three added-value services introduced in Section 7.2 are supported. Added-value services are related to user authorization, direct request routing and optimization of end to end QoS on the content path. The enhancements that we propose are well-adapted to the mobile context, and more specifically, to the EPS domain. This is mainly due to the fact that PCC is a mobile control architecture (PCEF is a mobile gateway like GGSN or P-GW) and ANDSF is an EPS entity. Nevertheless, similar control principles can be applied in a fixed context.

The main components of the new Telco control plane are two:

- Network APIs allowing OTTs subscription to added-value services proposed by the Telco
- A central brokering entity, similar to the one introduced in OMA, orchestrating the control plane operation and enabling the proposed services

A federation of CDNs can subscribe to one or many of the added-value services through an external API. The subscription scheme is shown in Figure 7.3.

Upon subscription, the controller provides to the Telco a list of premium users and contents and a list of services to which the CDN federation wants to subscribe. This

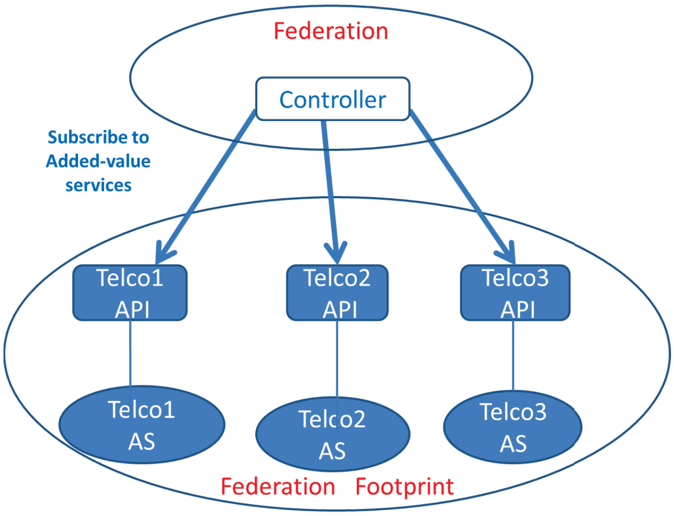

Figure 7.3: Subscription to added-value services

allows the Telco to identify, among all user sessions, the ones that will be subject to one or many of the added-value services. In a generic scenario, the list of users includes all the mobile subscribers of the Telco and the list of contents includes the content catalogs of the federation clients. For simplicity purposes, the controller provides to the Telco a list of root URLs referring to the federation clients. For instance, if YouTube is one client of the federation, one of the root URLs is *www.youtube.com*. Dependently of the agreement between the CDN federation and the Telco, any request targeting a URL that falls under this root is intercepted by the Telco as subject to one or many added-value services. The Telco communicates the registered root URLs to its DNS server. The goal is to identify the requests that are subject to added-value services and to redirect these requests to the broker.

The broker can be seen as an orchestrator of the Telco control plane. We suggest an inner design of the broker in order to enable the added-value services that we propose. The broker design that we suggest in shown in Figure 7.4.

The broker intervenes at the control and application layers. In fact, the broker embeds a proxy towards which all HTTP requests that are subject to one or many added-value services are directed. Once the broker intercepts a request that is subject to one or many added-value services, it interacts with different control entities, with the information system (SI) and, when requited, with third parties in order to fulfill and monetize this/these service(s).
At the reception of a HTTP request, the broker interacts with the SI in order to identify the service(s) to which the party that is authorative over the request has subscribed. The authorative party can be a CP, a pure-play CDN or a CDN federation. If the authorative party has subscribed to the *User Authorization* service, the broker interacts with user directories including the HSS and the SPR in order to authenticate the requesting user and authorize his access to the requested content. If the party is a CDN federation that has subscribed to *Direct Request Routing*, the broker uses the rules provided by the

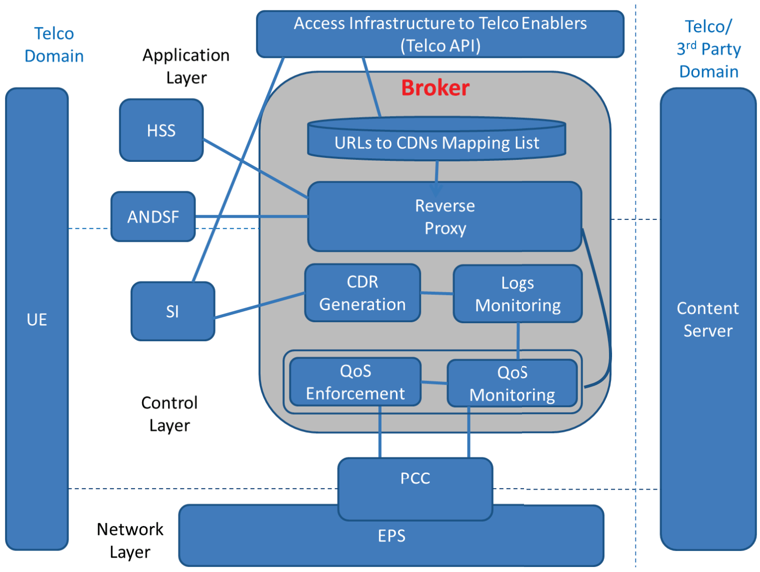

Figure 7.4: Design of the Telco Broker

federation controller in order to identify the CDN towards which the request should be routed. In case the target CDN is not the Telco-CDN, the broker uses DNS resolution for identifying the target CDN surrogate and directs the request towards this surrogate. Once the target CDN of a given request is identified, the broker can optimize, upon subscription, the end to end QoS on the path between the CDN surrogate and the requesting user. In this context, the broker interacts with the target CDN in order to negotiate video resolution at the beginning and during a video sesison. The broker also interacts with the ANDSF and PCC entities in order to enforce path selection and bandwidth allocation decisions. In particular, the broker can trigger, through the ANDSF entity, user's attachment to a new access and access reselection at the beginning and during an ongoing video session. The broker can also interact with the PCC in order to enforce bandwidth allocation between the user and the mobile gateway (P-GW in an EPS context) and to be notified of events, mainly referring to bandwidth fluctuations, occuring in its network.

In the next section, we instantiate a use case that illustrates the previously-listed roles of the broker.

7.5 Enabling Telco-based added-value Services: a Use Case

We assume that Orange offers, through a unified API, the three added-value services that we propose to third parties. We assume that a federation of CDNs that handles YouTube demand, subscribes, through its controller, to the three added-value services proposed by Orange. In particular, the federation delegates to Orange the authorization of its mobile users' access to different CDNs of the federation. The federation also delegates to Orange the direct routing of requests targeting any of the federation clients

towards the CDNs of the federation. Finally, the federation asks Orange to optimize end to end QoS for all sessions issued by its mobile users and targeting any of the federation clients.
When a mobile subscriber of Orange requests a YouTube video, the following operations will be performed (corresponding steps or call flows are shown in Figures 7.5 and 7.6):

- A DNS query is sent by the user to the local Telco DNS (Step 2). Through using the content URL, the Telco DNS identifies the request as subject to at least one added-value service. The Telco DNS redirects the request to the broker (Step 3).

- When intercepting the user's HTTP request, the broker contacts the SI in order to identify the service(s) to which the provider of the root URL has subscribed. In this example, the CDN federation that provided the root URL (*www.youtube.com*) has subscribed to the three added-value services.

- In order to authorize user access to YouTube content, The broker fetches the user authentication information in the HSS (Steps 5 and 6) . If the user is already authenticated, the broker procedes with the next steps. Otherwise, it sends an "Authentication Required" HTTP response to the user.

- The broker uses the requests routing data provided by the controller in order to identify the target CDN and to compute the IP address of the surrogate server. In this context, the broker plays the role of a DNS client with regards to the DNS name server of the target CDN (Steps 7 and 8). If the target CDN is the Telco-CDN, the surrogate server is directly selected by the broker based on geographic proximity criteria.

- In order to optimize end to end QoS on the content path, the broker fetches a *Resolution Description File* from the selected CDN surrogate (Steps 9 and 10). This file lists available resolutions for the requested video. If HTTP adaptive streaming is supported by the surrogate, This file correponds to the manifest file of the requested video.

- The broker triggers the ANDSF in order to gather information about the access networks that are in the user's context (Steps 12 and 13). These correspond to access networks to which the UE is already attached and to ones to which it can eventually connect. Provided information includes access types (cellular, wifi...), identifiers and conditions, in terms of available bandwidth.

- The broker positions access networks information with regard to the requirements, in terms of bandwidth, of different video resolutions. This positioning allows the broker to select the initial quality/resolution of the video, that is the highest resolution that can be supported by any access in the user's context (Step 14).

- If needed, the broker triggers the ANDSF in order to initiate user's attachment to a new access (Steps 15 and 16). The broker also provides to the PCRF information about the session (UE IP address, surrogate IP address) subject to QoS control as well as about the content path (access network identifier) and QoS rules to be enforced at this path level (Steps 17).

- The PCRF translates the QoS rules into parameters to be enforced at the traffic level (Guaranteed bitrate, Traffic class...). The P-GW uses these parameters in order to initiate the establishment of a dedicated bearer where video packets sent by the surrogate to the UE will be routed (Steps 18 and 19).

- The broker responds to the UE HTTP request with a 'HTTP Redirect' message that contains the IP address of the CDN surrogate (Step 20).

- The broker communicates the initial resolution to the CDN surrogate and requests the delivery of this resolution to the UE (Step 21).

- The PCC notifies the broker of an event occurring at the established bearer level. An event corresponds to an increase or decrease of the bearer end to end bandwidth (Step 35).

- The ANDSF informs the broker of an event corresponding to a change of the user's context (Step 24). Typically, an event refers to the UE attachment to a new 3GPP or non-3GPP access with better conditions than the one currently used for video delivery.

- Due to lack of streaming capacity at its level, The CDN surrogate updates the list of video resolutions that it can provide to the user (Step 42).

- When the broker is notified of any of the previously-listed events, it re-computes the content path and/or the QoS rules enforced at this path level (Steps 25, 36 and 44). The broker also selects a new video resolution which is inline with the new path and/or QoS rules. Typically, the broker can trigger the ANDSF in order to initiate user's attachment to a new access network (Steps 26, 27 and 28). In this case, the broker triggers the PCRF in order to initiate the establishment a new dedicated bearer and the removal the existing one (Steps 29, 30 and 31). If no change of the access network takes place, the broker triggers the PCRF in order to modify the QoS rules of the existing bearer Steps 37 and 45). In parallel, the broker notifies the surrogate of the newly-expected video resolution (Steps 40 and 48). In consequence, the surrogate adapts the resolution of the packets delivered to the UE (Steps 41 and 49).

- The PCC informs the broker of the end of the video session (Step 51). The broker sends signaling messages to ANDSF (Step 55), the PCC (Step 52) and the surrogate (Step 56) in order to terminate the signaling process.

It is important to note that, when authorizing user's access to the requested content, the broker generates a corresponding charging data record (CDR) and sends it to the SI. The generated CDR includes the IP address of the user and the requested URL.

Similarly, when selecting the target CDN for a given request, the broker generates a CDR that contains the IP addresses of the user and of the selected surrogate and sent it to the SI. Finally, at the end of a video session that was subject to a QoS optimization service, the broker generates a CDR that contains the session identifier referring to the user and surrogate IP addresses, the length of the session (in seconds) and the number of resolution adaptations witnessed during the session and sent it to the SI. Different CDRs are aggregated and later used by the SI in order to charge the parties that have subscribed to added-value services.

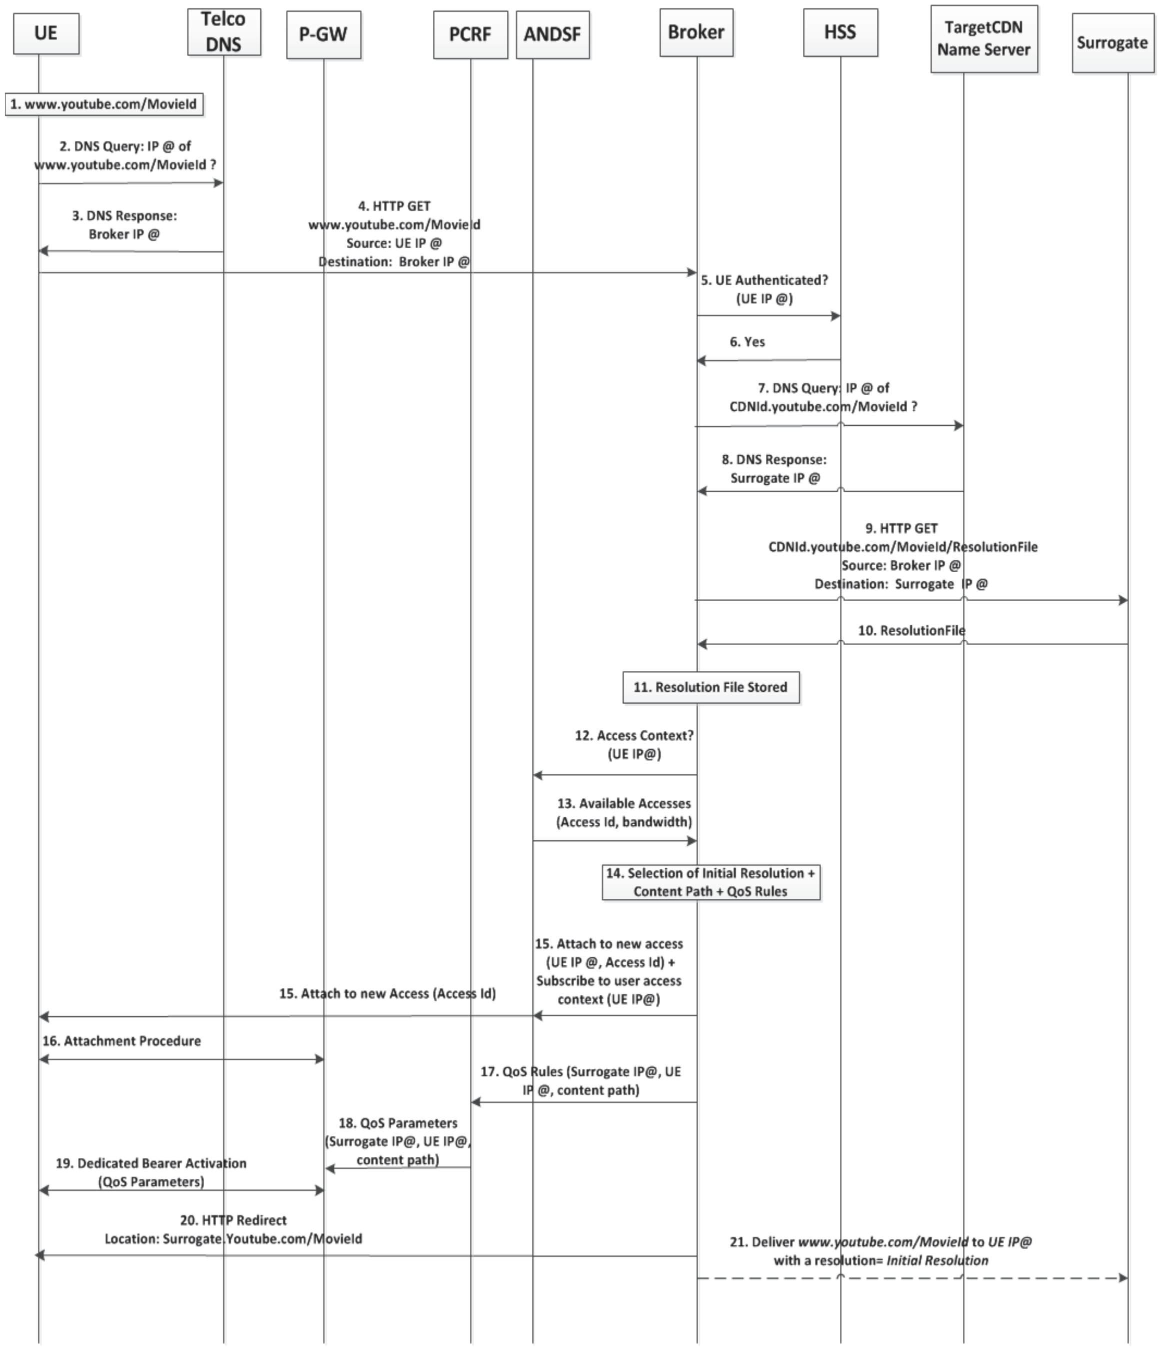

Figure 7.5: Call Flows (Part I)

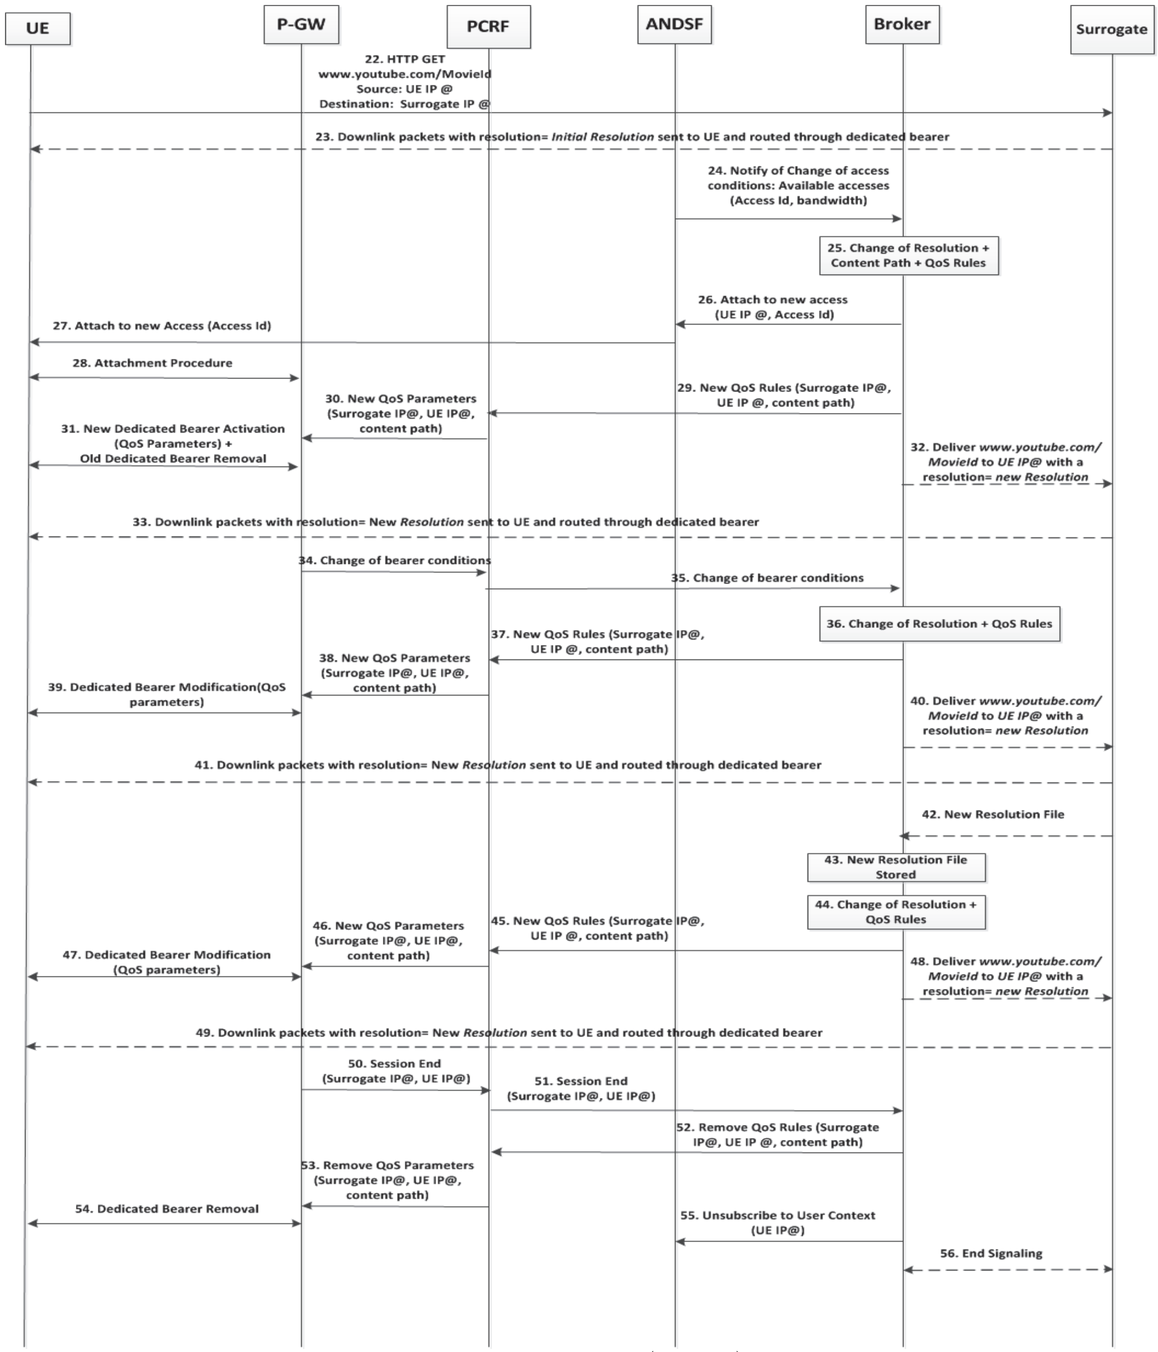

Figure 7.6: Call Flows (Part II)

Chapter 8

Conclusion and future work

8.1 Conclusion

The collaboration of CDN providers is a main aspect of evolution of content distribution services. Despite its attractiveness for both CPs and CDN providers and its rising interest among IT manufacturers, federation is one form of CDNs collaboration that has not been deeply investigated in the litterature.

In this thesis, we addressed the federation of autonomous and distinct CDN providers, a CDN provider being a player that owns vacant in-network or overlay resources at one or many geographic PoPs. We considered a system formed, on the one hand, of CDN providers with different capacities, footprints and cost models and, on the other hand, of CPs with different service requirements over a global or local footprint. Given this system, we introduced a CDN federation solution based on a centralized control architecture. This solution allows taking two types of decisions: static decisions of federation computation and provisioning and dynamic decisions of federation control. Federation computation consists of selecting the CDN providers that should federate together and the market jointly targeted by the so-formed federation. It also consists of computing an optimal policy of revenue sharing within the federation. Federation provisioning consists of deciding where, inside the federation, contents should be placed and how future users' requests should be directed. Dynamic decisions aim at ensuring a real-time routing of incoming requests towards different CDNs based on the provisioning phase output and on a dynamic vision of different CDNs state.

In order to enable static decision-making, we introduced an optimization model that allows computing and provisioning federations on a monthly basis. The model aims at maximizing the joint gain of the federation while taking into account CPs service requirements and CDN providers constraints in terms of capacity and economic fairness. Economic fairness, as defined in our model, means that each member of the federation is assigned a revenue that is proportional to its load. It also means that each member of the federation is given enough incentives to federate. A variant of the optimization

model is used for re-provisioning CDN federations on a daily basis.
We used the optimization model that we introduced in order to investigate three use cases of federation of interest for the CDN industry. The first use case addresses the federation of Telco-CDNs located in the same country. The second use case addresses the federation of pure-play CDNs with overlapping footprints. The third use case addresses the federation of Telco-CDNs and pure-play CDNs. For each of the use cases, we assessed the gains, in terms of extra revenue, achieved by CDN providers through moving from a separate scenario to a federation. We demonstrated that, when facing a high market demand, CDN providers have always an interest in federating. In particular, some CDN providers can double their economic gains through federation.

In the context of dynamic federation control, we focused on the control of peak events within a federation of CDNs. We introduced three frameworks that allow dealing with peak events at the federation level. One framework advocates an intra-CDN control of peak events referring to a dynamic adaptation of video resolution at individua CDNs level. The other frameworks privilege a joint control approach involving different federation members. We conduct trace-driven simulations in order to assess the performance of different frameworks. Two performance indicators are used in this context: the number of sessions rejected by the federation and the average video resolution witnessed by end users. We demonstrated that, when a joint approach for events control is adopted within a federation of CDNs, the proposed federation is better resilient to peak events. This translates into a higher hit ratio of the federation and a better video resolution witnessed by end users.

Given the important role of the Telco in the content provisioning value chain, we focused on the added-value services that it can provide to federations of CDNs and to individual OTTs including pure-play CDNs and CPs. Major added-value service have been identified in this context. These include: user authorization on behalf of OTTs, direct routing of requests to CDNs in a CDN federation context and optimization of end to end QoS between CDNs and mobile users. In order to enable these services, we suggest enhancements of the existing control infrastructure of the Telco. These enhancements take the form of network APIs, new Telco control entities and adaptations of existing control entities. We introduced a use case that illustrates the operation of the new control plane of the Telco.

In summary, we have addressed a new scheme of CDNs collaboration, that is the federation, which, in the light of the evolution of content distribution services, is attractive to both CDN providers and CPs. The federation solution that we propose presents the main advantages of being economically beneficial to different CDN providers and of ensuring a high resilience to unexpected traffic peaks which are continously occuring in the Internet. Given its many assets, the Telco could perfectly be part of this solution either through directly participating to a CDN federation or through providing added-value services to an existing federation.

8.2 Future Work

Even though we did the best we could in order to cover our research topic, some interesting research issues remain unsolved. We would like to list some of these issues as part of the future work.

Horizontal federation of content players In the context of this thesis, we have investigated the federation of CDN providers, a CDN provider being any player than owns vacant storage and streaming capacity. According to this new definition, a cloud provider can be a CDN provider. A carrier that owns in-network caching resources can also be a CDN provider. Nevertheless, in the use cases that we addressed (Chapter 5), we assumed that all the members of a federation have homogeneous charging models. These models were inspired by cost models that exist in the CDN market.
As a direct continuity of our research, it would be interesting to consider use cases of federation involving players do not belong to the same category and that do not adopt homogeneous business models. A federation of a cloud provider and a CDN provider (in the traditional sense of the term) falls in the category of horizontal federation that we describe. The same reasoning applies for a federation of a CDN provider and a carrier that does not own an overlay CDN. Issues like the joint business model of the so-formed federation and the policy of revenue sharing should be tackled. The positioning with regards to the *Separate Scenario* should be tackled as well given the fact that the meaning of *Separate Scenario* is not clear for all categories of players that can be involved in an horizontal federation.

The competition of CDNs In the context of this thesis, we have assumed a high market demand referring to a high number of CPs that want to delegate their content delivery to third-party CDNs. Given the capacity and footprint limitations of CDNs, we concludes that CDNs have technical and economic incentives to federate.
Let us consider a scenario where the market demand is low or where different CDNs extend their own infrastructure in order to overcome their limitations in terms of capacity and footprint. In this case, CDNs no longer have incentives to federate. Instead, we face a competition scenario where all CDNs have a strong position in the market and where one or a few of these CDNs is/are able to monopolize the market demand.
The competition of 'strong' CDNs is a research topic that can be addressed using mathematical tools like non-cooperative game theory. In this context, it would be interesting to investigate the strategies adopted by competing CDNs given the fact that they all have a strong and (almost) similar position in the market. It would also be interesting to translate these strategies into static and dynamic decision-making policies. An eventual control scheme, that can be centralized or distributed, can also be proposed.

Enabling in-network CDNs through SDN In the context of this thesis, we have focused on the Telco positioning in the CDN ecosystem and more particularly in a context of CDN federation. Given that the SDN concept is gaining popularity among ISPs, it would be interesting to investigate whether SDN can allow new positionings of the Telco in the content provisioning value chain.

The impact of peak events on CDNs performance has been widely discussed in this dissertation (mainly in Chapter 6). We showed that this kind of events can lead to the overload of a CDN or even of a federation of CDNs. When traffic offload is no longer possible, dealing with these peaks requires degradating the resolution of contents delivered to end users.
We suggest a pro-active, network based approach for dealing with peak events. This approach requires the support of SDN on the content path (path between CDNs and end users). It consists of the following: When an intermediate ISP (typically a Telco) detects an unusual surge of a given CP's traffic in its network, it activates, through a SDN like architecture, in-network caching for this CP contents. CP contents will be progressively stored in virtualised nodes in the ISP network. Future requests for these contents will be delivered, in an ICN-like fashion, by the network nodes. This allows alleviating the CDNs load and preventing overload and denial of service scenarios at CDNs level.
The main advantage of this approach with regards to transparent caching is related to the fact that it is controlled. Indeed, SDN allows an on-demand activation and deactivation of in-network caching. Furthermore, in-network caching can be performed for selected contents. Finally, in-network caching upon peak events can be proposed by the ISP as an added-value service to overlay CDNs.
Given the cited advantages, it would be interesting for a Telecom operator to further investigate this approach in terms of enabling architecture, decision-making logic (which contents should be stored, when to active/deactivate in-network storage etc), impacts on SDN controller and on network nodes etc.

Bibliography

[1] http://blog.streamingmedia.com/2011/02/att-building-out-their-content-delivery-network-using-edgecasts-software.html.

[2] http://blog.streamingmedia.com/2012/08/updated-list-of-vendors-in-the-content-delivery-and-transparent-caching-markets.html.

[3] http://drpeering.net/core/ch2-transit.html.

[4] http://highscalability.com/youtube-architecture.

[5] http://www-eio.upc.edu/lceio/manuals/cplex-11/html/.

[6] http://www.cnet.com/news/netflix-reaches-streaming-traffic-agreement-with-comcast/.

[7] http://www.marketwatch.com/story/onapp-powered-cloud-providers-join-forces-to-accelerate-web-performance-across-109-global-locations-2012-10-24.

[8] http://www.streamingmedia.com/articles/editorial/featured-articles/netflixs-new-cdn-is-good-news-for-everyone-84306.aspx.

[9] http://www.streamingmedia.com/articles/editorial/what-is-.../what-is-adaptive-streaming-75195.aspx.

[10] http://www.tcpipguide.com/free/.

[11] http://www.v-net.tv/ericsson-establishes-partnership-programme-for-global-cdns/.

[12] http://www.velocix.com/vx-portfolio/solutions/cdn-federation-interconnect.

[13] *Optimizationâ€”theory and applications.* Springer, 1983.

[14] *Multiple criteria optimization: Theory, computation, and application.* Krieger, 1989.

[15] *Internet Economics.* Massachussets Institute of Technology, 1997.

[16] Service accessibility. Technical Specification 22.011, 3GPP, 2008.

[17] *The IMS: IP Multimedia Concepts and Services*. 2009.

[18] Policy and charging control (pcc) over s9 reference point. Technical Specification 29.215, Third Generation Partnership Project (PCC), 2011.

[19] Cisco visual networking index: Forecast and methodology, 2012-2017. Technical report, Cisco, 2012.

[20] Cisco visual networking index: Global mobile data traffic forecast update, 2013-2018. Technical report, Cisco, 2013.

[21] General packet radio service. Technical Specification 23.060, Third Generation Partnership Project (PCC), 2013.

[22] Policy and charging control architecture. Technical Specification 23.203, Third Generation Partnership Project (PCC), 2013.

[23] The state of the internet. Technical report, Akamai, 2013.

[24] Telecommunication management, charging management, charging architecture and principles. Technical Specification 32.240, Third Generation Partnership Project (PCC), 2013.

[25] Telecommunication management, charging management, online charging system (ocs): Applications and interfaces. Technical Specification 32.296, Third Generation Partnership Project (PCC), 2013.

[26] Access network discovery and selection function (andsf). Technical Specification 24.312, 3GPP, 2014.

[27] Ip multimedia subsystem (ims). Technical Specification 23.228, 3GPP, 2014.

[28] 3GPP. http://www.3gpp.org/technologies/keywords-acronyms/100-the-evolved-packet-core.

[29] 3GPP. http://www.3gpp.org/technologies/keywords-acronyms/97-lte-advanced.

[30] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, and B. Ohlman. A survey of information centric networking. *IEEE Communications Magazine*, 2012.

[31] M. Ain. Architecture definition, component description and requirements. Technical report, Deliverable, PSIRP 7th FP EU-funded project, February 2009.

[32] Akamai. http://www.akamai.com/.

[33] V.K. Akhidari, Y. Guo, F. Hao, M.Varvello, V.Hilt, M. Steiner, and Z.L.Zhang. Unreeling netflix: Understanding and improving multi-cdn content delivery. In *IEEE INFOCOM'12*, Helsinki(Finland), August 2012.

[34] Alcatel-Lucent. www.alcatel-lucent.com.

[35] H.A. Alzoubi, S. Lee, M. Rabinovich, and J.Van der Merwe. A practical architecture for an anycast cdn. *ACM Trans. Web*, October 2011.

[36] Amazon. http://aws.amazon.com/cloudfront/pricing/.

[37] Amazon. www.amazon.com.

[38] A Data-Oriented (and Beyond) Network Architecture. T. koponen and m. chawla and b.g. chun and a. ermolinsky. In *SIGCOMM'07*.

[39] Apple. https://developer.apple.com/streaming/.

[40] I. Ari, B. Hong, E.L.Miller, S.A.Brandt, and D.E.Long. Managing flash crowds on the internet. In *IEEE MASCOTS 2003*.

[41] L. Atzori, A. Lera, and G. Morabito. The internet of things: A survey. *Computer Networks*, 2010.

[42] F.L.Faucheur B. Niven-Jenkins and N.Bitar. Content distribution network interconnection problem statement. Technical report, IETF, January 2012.

[43] BBC. http://www.bbc.com/.

[44] T. Berners-Lee, L. Masinter, and M. Macahill. Uniform resource locators (url). RFC 1738, IETF, 1994.

[45] G. Bertrand, E. Stephan, T. Brudbridge, and P. Eardly. Use cases for content delivery network interconnection. RFC 6770, IETF, 2012.

[46] BitGravity. www.bitgravity.com.

[47] BT, CenturyLink, China Mobile, Colt, Deutsche Telekom, KDDI, NTT, Orange, Telecom Italia, Telefonica, Telstra, and Verizon. Network functions virtualisation: An introduction, benefits, enablers, challenges and call for actions. In *SDN and OpenFlow World Congress*, 2012.

[48] P. Calhoun, J. Lougney, E. Guttman, and J. Zorn. Diameter base protocol. RFC 3588, IETF, 2003.

[49] C. Camaron and D. De Miguel. Mobile virtual network operators (mvno) basics. Technical report, Valoris Viewpoint, 2008.

[50] CDNetworks. fr.cdnetworks.com.

[51] CDNPlanet. http://www.cdnplanet.com/.

[52] Cedexis. http://www.cedexis.com/.

[53] Cisco. Wholesale content delivery networks: Unlocking new revenue streams and content relationships. Technical report, 2012.

[54] Conviva. http://www.conviva.com/.

[55] DailyMotion. http://www.dailymotion.com/.

[56] J.K. Deane, T.R. Rakes, and A. Argawal. Designing content distribution networks for optimal cost and performance. *Information Technology and Management*, 2012.

[57] F. Dobrian, V. Sekar, A. Awan, I. Stoica, and D. Joseph. Understanding the impact of video quality on user engagement. In *ACM SIGCOMM'11*.

[58] C. Drake, T. Brown, D. Mavrakis, and G. Cottle. Cdn future scenarios: A growing market delivers new business opportunities. Technical report, Informa Telecoms and Medias, July 2012.

[59] Dyn. http://aiscaler.com/dcc/dyn-cdn-manager.

[60] D. Eastlake. Domain name system (dns): Iana considerations. RFC 5395, IETF, 2008.

[61] Ericsson. www.ericsson.com.

[62] ETSI. http://www.etsi.org/technologies-clusters/technologies/nfv.

[63] Facebook. http://www.facebook.com/.

[64] P. Faratin, D. Clark, S. Bauer, W. Lehr, P. Gillmore, and A. Berger. The growing complexity of internet interconnexion. *Communications and Strategies*, 2008.

[65] R. Fielding, J. Gettis, J. Mogul, and H. Friestik. Hypertext transfer protocol - http/1.1. RFC 2616, IETF, 1999.

[66] Open Networking Foundation. https://www.opennetworking.org/.

[67] N. Fujita. Coarse grain replica management strategies for dynamic replication of web contents. *Computer Networks*, 2004.

[68] S. Gosselin, T. Menouni, p. Bertin, J. Torrijos, and D. Breuer. Converged fixed and mobile broadband networks based on next generation points of presence. In *Future Network and Mobile Summit*, 2013.

[69] M. Handley, V. Jackobson, and C. Perkins. Sdp: Session description protocol. RFC 4566, IETF, 2006.

[70] K. Hosanagar, R. Krishnan, M. Smith, and J. Chuang. Optimal pricing for content delivery network (cdn) services. In *37th Hawaii International Conference on System Sciences*.

[71] G. Ibrahim, D. Kofman, Y. Chadli, and A. Ansiaux. Toward a new telco role in future content distribution services. In *ICIN'12*, 2012.

[72] IETF. www.ietf.org.

[73] NFV ETSI ISG. Network function virtualisation: Use cases. Technical report, ETSI, 2012.

[74] V. Jacobson, D.K. Smetters, J.D. Thornton, and M.F. Plass. Networking named content. In *CoNEXT'09*.

[75] J.W. Jiang, S. Ioannidis, L. Massoulie, and F. Picconi. Orchestrating massively distributed cdns. In *CoNEXT'12*, Nice(France), December 2012.

[76] J. Jung, B. Krishnamurthy, and M. Rabinovish. Flash crowds and denial of service attacks: characterization and implications for cdns and web sites. In *WWW '02*.

[77] J. Kangasharju. Object replication strategies in content distribution networks. *Computer Communications*, 2002.

[78] M. Latouche, J. Defour, T. Renger, T. Verspeht, and F. LeFaucheur. The cdn federation: Solution for sps and content providers to scale a great customer experience. Technical report, Cisco, 2012.

[79] K. Lee. Making a technological catch-up: Barriers and opportunities. *Asian Journal of Technology Innovation*, 2011.

[80] Level3. http://www.level3.com/.

[81] S. Levy. Going with the flow: Google's secret switch to the next level of networking. *WIRED ENTREPRISE*, 2012.

[82] Limelight. http://www.limelight.com/.

[83] H. H. Liu, Y. Wang, Y. R. Yang, H. Wang, and C. Tian. Optimzing cost and performance for content multihoming. In *SIGCOMM'12*, Helsinki(Finland), August 2012.

[84] Y. Le Louedec, N. Amman, C. Hawinke, and S. Nopert. Technical report, OCEAN European Project, 2012.

[85] A. Manzalani, R. Minerva, E. Deckel, Y. Tock, and N. Crespi. Manifesto of edge ict fabric. In *17th International Conference on Intelligence in Next Generation Networks*, 2013.

[86] OnApp Marketplace. http://www.onapp.com/cdn/.

[87] MaxCDN. www.maxcdn.com.

[88] U. Michel. Sdn and nfv: How things fit together. Technical report, Deutsche Telekom, September 2013.

[89] Microsoft. http://www.microsoft.com/silverlight/smoothstreaming/.

[90] Mirror-Image. www.mirror-image.com.

[91] P. Mockapetris. Domain names: Concepts and facilities. RFC 1034, IETF, 1987.

[92] P. Mottishaw. Policy control and charging for lte network. Technical report, Openet, 2009.

[93] Netflix. http://www.netflix.com/.

[94] Limelight Networks. http://www.limelight.com/services/orchestrate-content-delivery.html.

[95] R. Nibarger and M. Teubner. An introduction to long term evolution (lte). Technical report, Community Oriented Policing Services, 2012.

[96] OMA. www.openmobilealliance.org.

[97] EdgeCast OpenCDN. http://www.edgecast.com/solutions/opencdn/.

[98] Orange. www.orange.com.

[99] Ovum. http://ovum.com/2012/11/20/akamai-solidifies-its-position-in-carrier-cdn-with-two-moves/.

[100] G. Pallis and A. Vakali. Insight and perspectives for content delivery networks. *Communications of the ACM*, 2006.

[101] A. Mukaddim Khan Pathan and R. Buyya. A taxonomy and survey of content delivery networks. Technical report, 2006.

[102] C. Perkins. Ip mobility support. REC 2002, IETF, 1996.

[103] R.S. Peterson and E.G. Sirer. A content propagation metric for efficient content distribution. In *SIGCOMM'11*.

[104] PlanetLab. www.planet-lab.org.

[105] S. Podlipnig and L. Boszremenyi. A survey of web replacement strategies. *ACM Computing Surveys*, 2003.

[106] I. Poese, B. Frank, B. Ager, G. Smaragdakis, and A. Feldmann. Improving content delivery using provider-aided distance information. In *IMC'10*, 2010.

[107] L. Qiu, V.N. Padmanabhan, and G.M. Voekler. On the placement of web server replicas. In *20th IEEE INFOCOM Conference*, April 2001.

[108] Telco 2.0 Research. Broadband 2.0: Mobile cdns and video distribution. Technical report, August 2011.

[109] R. Rodriguez and P. Druschel. Peer to peer systems. *Communications of the ACM*, 2010.

[110] J. Rosenberg, H. Schulzrinne, G. Camarillo, A. Johnston, and J. Peterson. Sip: Session initiation protocol. RFC 3261, IETF, 2002.

[111] Y. Sanchez, C. Hellge, T. Schierl, and W.V. Leekwijck. Scalable video coding based dash for efficient usage of network resources. In *Third W3C Web and TV Workshop*.

[112] H. Schulzrinne, S. Cazner, R. Frederik, and V. Jacobson. Rtp: A transport protocol for real-time applications. RFC 1889, IETF, 1996.

[113] H. Schulzrinne, A. Rao, and R. Lanphier. Real time streaming protocol (rtsp). RFC 2326, IETF, 1998.

[114] H. Schwarz, D. Marpe, and T. Wiegand. Overview of scalable video coding: Extension of the h264/avc standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 2007.

[115] K. Sripanidkulchai, B. Maggs, and H. Zhang. An analysis on live streaming workloads on the internet. In *IMC'O4*, 2004.

[116] T. Stockhammer. Dynamic adaptive streaming over http. In *Second annual ACM conference on Multimedia systems*.

[117] A. Su, D.R Choffnes, A. Kuzmanovic, and F.E Bustamante. Drafting behind akamai (travelocity-based derouting). In *SIGCOMM 2006*, 2006.

[118] British Telecom. www.bt.com.

[119] Deutsche Telekom. www.telekom.com.

[120] T.R.Puzak. *Analysis of Cache Replacement Algorithms*. PhD thesis, University of Massachusetts - Amhrest, 1985.

[121] Twitter. http://www.twitter.com/.

[122] Video on demand, 1994. Patent.

[123] VentePrivee. http://www.venteprivee.com/.

[124] WebRTC. http://www.webrtc.org/.

[125] M.B. Weiss and S.J. Shin. Internet interconnection economic model and its analysis: Peering and settlment. *Netnomics*, 2003.

[126] XDN. http://www.xdn.com/.

[127] YouTube. http://www.youtube.com/.

[128] Y. Zhang, R. Yu, S. Xie, and V. Yao. Home m2m networks: Architectures, standards and qos improvements. *IEEE Communications Magazine*, 2011.

Appendix A

In order to linearize the cost function, we introduce an intermediary binary variable that we denote by ζ_{jpt}^{iml}. ζ_{jpt}^{iml} is the product of the y_{jp}^{iml} and λ_{jp}^{it} variables.
Given the new variable, the monthly cost $Cost_2^{ji}$ paid by a given j $\mathcal{S}$ for routing its content to a given zone i ∈ $\mathcal{F}$ is expressed as follows:

$$Cost_2^{ji} = \sum_{p\in\ \mathcal{P}_j} \sum_{t\in\ [1-24]} \sum_{m\in\ \mathcal{G}} \sum_{l\in\ \mathcal{C}_m} C_{jpt}^{iml} \times \zeta jpt^{iml}$$

Where C_{jpt}^{iml} is a constant that depends of the model parameters. C_{jpt}^{iml} can be expressed as follows (A_{jp}^{i} can be calculated through equation (4.5) and C_{band} is a constant introduced in section 4.2):

$$C_{jpt}^{iml} = Peak_m^t \times D_m^i \times V_l^i \times A_{jp}^i \times C_{band}$$

The following three constraints are added to the model constraints listed in sections 4.3, 4.4 and 4.6. These constraints aim to stress a linear relationship between the intermediary variable ζ_{jpt}^{iml} and the initial variables y_{jp}^{iml} and λ_{jp}^{it}.

$$\zeta_{jpt}^{iml} \leqslant y_{jp}^{iml}$$

$$\zeta_{jpt}^{iml} \leqslant \lambda_{jp}^{it}$$

$$\zeta_{jpt}^{iml} \geqslant \lambda_{jp}^{it} + y_{jp}^{iml} - 1$$

List of Publications

Published

- A Centralized Architecture for Establishing Federations of Content Players, Ghida Ibrahim and Daniel Kofman, Patent, France, FR 13 53215
- A Telco-based Solution for end to end QoS Optimization, Ghida Ibrahim, Patent, France, FR 14 55483
- Improving Content Delivery Through Coalitions, Ghida Ibrahim and Daniel Kofman, Paper, IEEE/IFIP Network Operations and Management Symposium 2014 (NOMS 2014), May 2014, Krakow, Poland
- Federating Heterogeneous, Distributed Content Players, Ghida Ibrahim and Daniel Kofman, Presentation, Third ETSI Workshop on Future Networks, April 2013, Sophia Antipolis, France
- Toward a New Telco Role in Future Content Distribution Services, Ghida Ibrahim and Daniel Kofman and Alexandra Ansiaux and Youssef Chadli, Paper, 16th International Conference on Intelligence in Next Generation Networks (ICIN'12), October 2012, Berlin, Germany
- Enhancing Resources Placement in a Multi-CDNs Context, Ghida Ibrahim, Poster, 24th International Teletraffic Conference (ITC'12), September 2012, Krakow, Poland

Ongoing

- Managing Unexpected Events in Multi-CDNs Systems, Paper to be submitted to the 5th International Conference on Network of the Future (NOF 2014)
- Optimizing Content Delivery in a Mobile Context: a Telco-based API, Paper to be submitted to the 18th International Conference on Intelligence in Next Generation Networks (ICIN'15)
- Federating Distinct, Distributed CDN Providers, Paper to be submitted to IEEE Transactions on Network and Service Management

Evolution of the Control Plane for Future Content Distribution Services

Ghida IBRAHIM

RESUME : L'écosystème de distribution de contenus évolue rapidement. Un axe majeur d'évolution concerne la fédération de fournisseurs de réseaux CDNs avec des capacités et couvertures différentes. Dans le contexte de cette thése, nous introduisons une solution technique qui permet la gestion des aspects statiques de mise en place et d'approvisionnement de fédérations de fournisseurs de réseaux CDNs ainsi que des aspects de contrôle d'évenements dynamiques au sein de fédérations établies. Au delà de la solution proposée, les différentes contributions de cette thèse permettent de quantifier les gains, en termes économique et en termes de performance, d'une approche fédérée par rapport a une approche non fédérée

MOTS-CLEFS : Fédération, CDN, OTT, Telco, Controle

ABSTRACT : Content distribution services are rapidly evolving in various directions. A major evolution trend concerns the federation of Content Delivery Network (CDN) providers with different capacities and footprints. In the context of this thesis, we introduce a technical solution that allows dealing with static aspects of CDN federation establishment and provisioning as well as dynamic aspects of control of unexpected events within a federation of CDNs. Beyond the proposed solution, different contributions of this thesis allow quantifying the gains, in economic and performance terms, of a federated approach with respect to the status quo thus enabling the translation of the federation concept into a reality in the CDN market.

KEY-WORDS : Federation, CDN, OTT, Telco, Control

TELECOM ParisTech

ParisTech
INSTITUT DES SCIENCES ET TECHNOLOGIES
PARIS INSTITUTE OF TECHNOLOGY