216217 P2P-Next

D9.2d

Executive Summary of WP2-WP4 for the Project P2P-Next

Contractual Date of Delivery to the CEC: 30th April 2012

Actual Date of Delivery to the CEC: 30th April 2012

Author(s): Ulf Lindqvist, VTT

Participant(s): Thomas Look (MFG), Miriam Pelka (MFG), Njål Borch (NORUT), Johan Pouwelse (TUD)

Co-authors: Silvano Mignanti, Francesco Delli Priscoli, Gaetano Scarano, Alberto isidori, Andi Palo (UoR)

Workpackage: WP9
Est. person months: 2
Security: PU
Nature: R
Version: 1.0
Total number of pages: 27

Abstract:

The scope of the P2P-Next project is to develop an open source platform based on open standards and supporting future proof iterative integration for the next generation of interactive television. WP2-WP4 link the project to the market demands and the legal environment and ecosystem, map the user requirements, design the architecture of the project and, finally, develop the technology for content sharing merging content, communities, communication and commerce (Next-Share). This executive summary gives the results achieved in WP2-WP4 during the entire project in a nutshell.

Keywords:

Dissemination, Standardisation, Publications, Workshops

Document history

DATE	VERSION	STATUS	COMMENTS
08.08.2011	0.1	Draft	Disposition, UL
22.02.2012	0.2	Draft	WP2, MP
22.02.2012	0.3	Draft	WP3, EB
16.04.2012	0.6	Draft	WP4, JP
17.04.2012	0.61	Draft	WP4.4. L-EE
29.04.2012	0.9	Draft	Final draft, UL
30.04.2012	1.0	Final	Final version, JA

Executive Summary

The scope of the P2P-Next project is to develop an open source platform based on open standards and supporting future proof iterative integration for the next generation of interactive television. WP2-WP4 link the project to the market demands and the legal environment and ecosystem, map the user requirements, design the architecture of the project and, finally, develop the technology for content sharing merging content, communities, communication and commerce (**Next-Share**).

The main objectives of WP2 are the assessment of the legal situation, research on and development of business models, as well as observation of the market trends. Further work inlcuded service sketching of both pre-commercial services and applications, and targeted ad and editorial content devices. Market watch included study of general technological developments, a comparison of mobile payment application approaches as well as an examination of services and applications for P2P delivery.

The key outcome in the field of Public Interest and Industry Perspective were the identification of business requirements and issues found by analysing the market with regards to P2P systems. Trends in the area of digital media and IT related media services crucial for the project were identified and six generic business models were defined. For the tasks of Application and Service Development WP2 drew up several services and drafted associated business plans. Furthermore WP2 developed a number of pay-per-view and free view services as well as a first ad server version along with the build-up of a targeted editorial and promotional content device for the project.

WP3 provided the end user and system requirements specification for the Next-Share platform and the architecture specification and design. A set of fictive user personas were selected and their interactions with the system were described in scenarios. More technical and content provider requirements were identified based on input from other work packages and from the living lab trials. The architecture specifications were the basis for the Next-Share software, which was made extremely modular applicable to specific end-user interfaces. The work package produced three deliverables: the roadmap and design document in three iterations.

Work package 4 comprises the technical core of P2P-Next. The P2P-Next content sharing platform is called: **Next-Share**, which is a self-organizing system with complete decentralisation and hence lacks any central bottleneck or choke points that may hamper performance, induce setup cost, or require maintenance. The work was directed to integration with other work-packages, the IPvNext networking fabric design, implementation and testing, network awareness, locality and geography, development of a data exchange engine for data voice and video, and development of micro payments and resource credits.

Content

1 Executive Summary of WP2: “Business and Legal Ecosystems”

WP2 focused on regulation and legislative aspects, businesses models, market watch and P2P ecosystems. During the project lifetime WP2 has defined six business models to take the P2P-Next aspects into account and embedded them into specific cases likely to appear. Also emerging forms of monetisation like bundling and long-tail have been related to these models bringing further innovative aspects to the project and new options to refinance content. In addition current regulations have been considered, the most critical of which are the TV Directive, the Copyright Directive, and the so called Sanctions Directive. This has been completed by looking at relevant documents released by the European Commission on Europes legal future, attending events dealing with these topics as well as tracing current court cases on copyright infringements in online media.

WP2 linked the market developments *outside* of P2P-Next to the requirements and development process *inside* the project. As a result WP2 has produced a combination of:

a) comprehensive reports, studies and assessments,
b) a set of pieces of software in the form of applications and services combining the knowledge gained from developing business models in WP2 with RTD-work undertaken in work packages 3 – 6, and
c) actions when dealing with the user group, monitoring risk and defining contingency plans.

While the other work packages focused on developing, researching, validating, and showcasing the P2P own research results, technologies, and content, WP2 took the foreground of the project and linked it with third-party services and applications in a way that the complete offering would make P2P-Next superior to competing offers on the market.

The input of the WP2 sub-work packages was required at different points of the 4 years lifetime of the project and therefore the focus on the different tasks varied over time. The main focus of year 1 and 2 was to present a comprehensive view on the market environment of P2P-Next including the SotA in online media. This was provided by the sub-work packages 2.1 - 2.4 resulting in the 2.1.1 deliverables. From year 3 of the project onwards, efforts were shifted from theoretical research towards more practical work. This has resulted in a higher focus on the sub-work package 2.5 and the development of services and applications. A prototype device for targeted ads and editorial content (TAEPCD) has been developed and business plans for all P2P-Next services were developed.

WP2 activities had impact on all activities in the project ranging from WP3 - WP8. The work package has produced the following deliverables over the projects runtime, i.e. until M52:

M7 Milestone	• D2.1.1 Requirement, Business Model, Regulation, Legal Aspects Report (Stakeholder requirements, business model use cases, UGC, etc.)
M12 Milestone	• D2.1.1b Update on Report at Month 7 • D2.5.1 P2P Services v1.0 (Incl. Transaction Interaction, DRM, Billing, Micropayment • D2.5.2 Pay View and Free View Services • D2.5.3 Targeted ad and promotional content device
M18 Review	• D2.1.1 Amendment of the Initial Report at Month 7
M24 Milestone	• D2.1.1c Update on Previous Reports • D2.5.1b Update on M12 Report • D2.5.2b Update on M12 Report • D2.5.3b Update on M12 Report
M27 Review	• No amendments of the deliverables required
M36 Milestone	• D2.1.1d Update on Previous Reports • D2.5.1c Update on Previous Reports • D2.5.2c Update on Previous Reports • D2.5.3c Update on Previous Reports
M52 Milestone	• D2.1.1e Update on Previous Reports • D2.5.1d Update on Previous Reports • D2.5.2d Update on Previous Reports • D2.5.3d Update on Previous Reports

Table 1: Deliverables of WP2 until M52

The main year 1 research goal was to gather a wide array of State of the Art (SotA) information on the sub-work packages of WP2, mainly “Public Interest and Industry Perspective”, “Business Models” as well as “Regulatory and Legislative Aspects”. The results were guidelines as well as recommendations important for the project from its beginning. This stock of information was constantly updated and enhanced with further aspects. Regarding the targeted ad and editorial and promotional prototype content device (TAEPCD) the main result of year 1 was a proof of concept, which was provided along with the 2.5.3a deliverable. At first a comprehensive market watch was conducted presenting current stakeholders at the ad server market as well as existing standards in this area. Based on this a general description of the TAEPCD and possible modules of the device were defined.

A focus in year 2 of the project was to further assess aspects of the swiftly developing P2P market that gave important insight of the chances and challenges for P2P-Next. This research led towards explicit recommendations for the project to ensure P2P-Next keeping pace with market developments - and being able to move beyond, to provide QoS (Quality of Service) and QoE (Quality of Experience) as expected by

users, and to define those business opportunities which allow for a sustainable exploitation of the project results.

In terms of the TAEPCD a more detailed conception of the device in year 2 resulted in a first pre-commercial prototype implementation able to deliver standard linear and non-interruptive ad forms, like pre-rolls or overlay ads. The outcome of this task was the basis for the creation of the media planning engine of WP5 (D5.1.2). For the area of application and service development (WP2.5) a set of services - adjusted to the NextShare platform - was defined.

In year 3 WP2 focussed on issues such as impact and uptake of the NextShare technology. Regarding the market environment the focus was changed towards B2B-related aspects. In this course a business related evaluation on market leading media transport protocols was prepared and used for a SWOT analysis of generic P2P-systems and the NextShare platform as well as for a comparison with the P2P-Next Libswift protocol. With respect to the SWOT analysis objectives and approaches towards the P2P-Next project were defined. In the area of legal and regulatory environment, WP2 partners DACC and KTH attended several events dealing with future developments of EU media legislation and met decision makers of this area. An invited seminar on politicians responsibilities regarding creativity and remuneration in a rather confused legal/regulatory environment was held for EU Members of Parliament. Two meetings were held with DG INFSO and DG MARKET presenting WP2 findings and views on the regulatory and legal situation. Contributions have been made to the EU Digital Agenda Stakeholder Day.

Within WP2.5 payment and ad services were further developed. The main focus lay on the refinement of the AdServer technology and related integration work. First steps for AdServer Integration into WP4 and WP5 Limo were taken. Furthermore efforts towards an implementation of IAB VAST- and VPAID-standards into HTML5 were made.

In the final year of the project and during the extension until April 2012 focus on impact and uptake of NextShare technology continued. Market watch included study of general technological developments, a comparison of mobile payment application approaches as well as an examination of services and applications for P2P delivery. An update on NextShare stakeholder’s requirements towards Internet distribution systems, and P2P-delivery in particular, was achieved by a broadcaster survey accomplished by WP2. The survey for example confirmed that IP-based online services are a growing market. Also the approaches of P2P-Next regarding a dedicated payment system and the support of SotA DRM systems seemed to fit the demands of European audio-visual online services. Similar applied for the European approach of the project, since all participants of the survey indicated to have their main audience in Europe. Concerning bandwidth costs there existed big differences that allowed the conclusion that this is still an issue that for some service providers even today counts for more than 50% of their budget for IP-based distribution. On the other hand P2P seemed not attractive enough as an alternative distribution method, led by the requirement for a plugin and varying data rates linked to P2P technology.

In regard to cost aspects WP2 carried out OPEX/CAPEX analyses to find out about the efficiency of P2P and alternative delivery mechanisms. On the basis of three use cases - drafted as potential service using the NextShare platform - WP2 analysed capital and operational expenditures for models using CDN, purchase or p2p approaches. The main result was that p2p can drastically reduce costs, especially for traffic costs that can be reduced about 40% for offers that require high bandwidths.

Further efforts in the last period of the project went into the the refinement of requirements for P2P-Next in the light of latest media platform developments and the validation of NextShare features through test beds set up by IRT. The test results showed that the NextShare P2P algorithm is working well when there are no network restrictions. The use of NextShare clients can save around 50-75% on the server output data rate. The start and switching times were slightly higher than those of adaptive streaming protocols and the average offset of 20 seconds to the live signal was seen too high for a “live” TV service.

Another key aspect in year 4 of the project and an important issue to turn the outcomes of the project into marketable products was the development of a business plan for the P2P-Next services which were drafted during the projects runtime. In this course a possible role model was developed, in which the P2P-Next operating service – realized in form of DACCs FairShare MediaWorld (FSMW) approach – builds the key service that is accompanied by enhancement services that can be optionally applied:

- Interactive TV service
- Payment services
- Targeted Ad and Editorial an Promotional Content Device (TAEPCD)

Details on the model and functionalities of the services can be found in the deliverables d2.5.1 and d2.5.2.

Finally, the current status of the legal and regulatory environment for P2P-Next was gained by updating the legal situation, a look at selected country-by-country cases (Germany, Slovenia and the Netherlands) as well as analysing the outcomes of the Digital Agenda Assmembly, June 2011. After all, the legal situation in Euope remains difficult: The Digital Agenda still is more or less in limbo and regulation for new media and net neutrality in the individual countries is not clearly defined or simply ignored.

The results of WP2 will be exploited mainly in the form of business plans and an operational service FairShare MediaWorld – an open marketplace for distribution and consumption of digital media and services.

2 Executive Summary of WP3: “User Requirements and Architecture”

WP3 is responsible for the user requirements and architecture of the P2P-Next project. Whereas within WP2 the focus is on the stakeholders, WP3 focuses on the user requirements, ensuring that we are putting the user first. The Living Lab (WP8) is providing an organized test ground for our design and software and the tests performed in the Lab in turn provide valuable insights into how well the user requirements are met.

The work package is divided into two sub packages; WP3.1 provides the requirements specification and WP3.2 provides the architecture specification and design. As P2P-Next is a 4 year long project, it was expected that users would change their media habits during the project, also as a result of having new technology available. This has also turned out to be the case. Being able to predict how these are going to be used is difficult, therefore lab trials with actual users have been of great importance. This activity has specified the implementation goals, expected experiments to be conducted, subsequent outcomes and identified impact at each iteration. This has provided a framework for the subsequent architecture design and development as well as metrics by which to measure the results. Both the requirements specification and the architecture specification have been revisited yearly through the project based on these trials.

WP 3.2 has produced four deliverables, (three iteration of the requirement and architecture document)

D3.1.1 Requirement and architecture specification, three iterations
D3.1.2 Living Lab Specification and Objectives

After the last revision, WP3 has collaborated with WP4 (development) and W8 (living lab) to ensure that all high level requirements are met, and that lab results are reflected in the requirements. Requirement specification

2.1 Requirement specifications

The objectives of WP3.1 were to establish a complete set of end-user and system requirements based on real user needs and to understand both the functional and non-functional requirements associated with the development of the Next-Share platform and associated user experiences (Next-ShareTV and Next-SharePC). The deliverables for this work package were specifications which form the basis of the development and deployment of a number of iterative proof-of-concept implementations. This task has exploited iterative processes including user involvement, ethnographic study, study of best-practice and state of the art, design and requirements analysis in order to formulate the requirements specifications.

In order to identify new and innovative technologies and business models, WP3 has selected a set of personas. The personas are fictive users of the system based on analysis both of current state-of-the-art, surveys, Living Lab feedback and partner experience. Interaction between the personas and the Next-Share system is described in scenarios, which has been used to identify and specify end user requirements for the system. Both personas and scenarios have been updated to reflect changes in technology, user behaviour and availability of content during the project. More technical requirements as well as content-provider requirements are identified based on input both from WP2 (Business and legal Ecosystem), WP5 (Content), WP6 (Next-SharePC prototype), WP7 (Next-ShareTV prototype) and WP8 (Living Lab trials). Also some requirements have been deleted or have changed the priority based on the same input. The resulting requirement specification has been used in WP3.2 to create and update an architecture and system design that is well aligned with the requirements and the envisioned scenarios.

The scenarios are based on the personas, top level use cases for P2P-next and business models from the project. The scenarios give concrete examples of user interaction with NextShare, and are easy to communicate. They point to central requirements and reflect changes in technology, user behaviour and availability of content. The scenarios are also easy to play out in trials.

Both personas and scenarios were revised after the first year and then again one year after based on changes in the technology and services, on results from Living Lab trials and on emerging user patterns.

More technical requirements as well as content-provider requirements and use cases are identified based on input both from other work packages. The work package contributions are in short:

- WP2 (Business and legal Ecosystem) – access control
- WP 3 (Requirements) - end user requirements
- WP4 (Tehchnology) - Scalability, fault tolerance, payments
- WP5 (Content) - Multi channel content
- WP6 (NextSharePC) - Browser integration
- WP7 (NextShareTV) - Resource usage
- WP8 (Living Lab trials) - User experience

Since the requirements come from several sources and some are overlapping, we needed a way to sort them so that it was easy to both find possible overlaps and identify them. The requirements are divided into the following phases / groups:

1. Preparation – done by content provider before material is injected
2. Injection – injection of contents into NextShare
3. Discovery of material in NextShare
4. Distribution / playing / data delivery

5. Commercial exploitation
6. Monitor / Evaluate
7. External contact
8. Platform and equipment
9. Interaction
10. Framework conditions

Each of the requirements can be attached to one or more primary actor:

- Provider
- Consumer
- Authoritative entity

The resulting requirement specification is used in WP3.2 in order to create an architecture and design that is well aligned with the requirements and the envisioned scenarios.

As a subtask, WP3.1 also provides a living lab specification based on the initial requirements specification and ethnographic studies obtained in task WP3.1. This includes the implementation goals, expected experiments to be conducted, subsequent outcomes and identified impacts. It provides a framework for the subsequent architecture development (WP3-6) and also metrics by which to measure the results. The living lab specification was completed in month 15 of the project as deliverable d3.1.2.

In the last project phase the main input to the requirements have come from living lab trials.

2.2 Architecture and design specification

Work package 3.2 has created a design of P2P-Next based on the Architecture Specification document produced by WP 3.1. The design has been kept modular, allowing partners to develop their solutions in parallel when possible. A road map that mapped the design out in time was created, allowing for several iterations in the development phase.

The architecture and design specification has been the basis for WP4 (Implementation) work performed within the project. Work package 3.2 created a design of P2P-Next based on the Architecture Specification document produced by WP 3.1.

The first year saw a top level architecture that has been modified and refined in the second and years. The main changes were that TCP has been exchanged with libswift, PC App with XPI Installer and Open access with access control. No major architectural changes were performed in the last year of the project.

The high level interaction between the core modules of the Next-Share software is shown in the following figure:

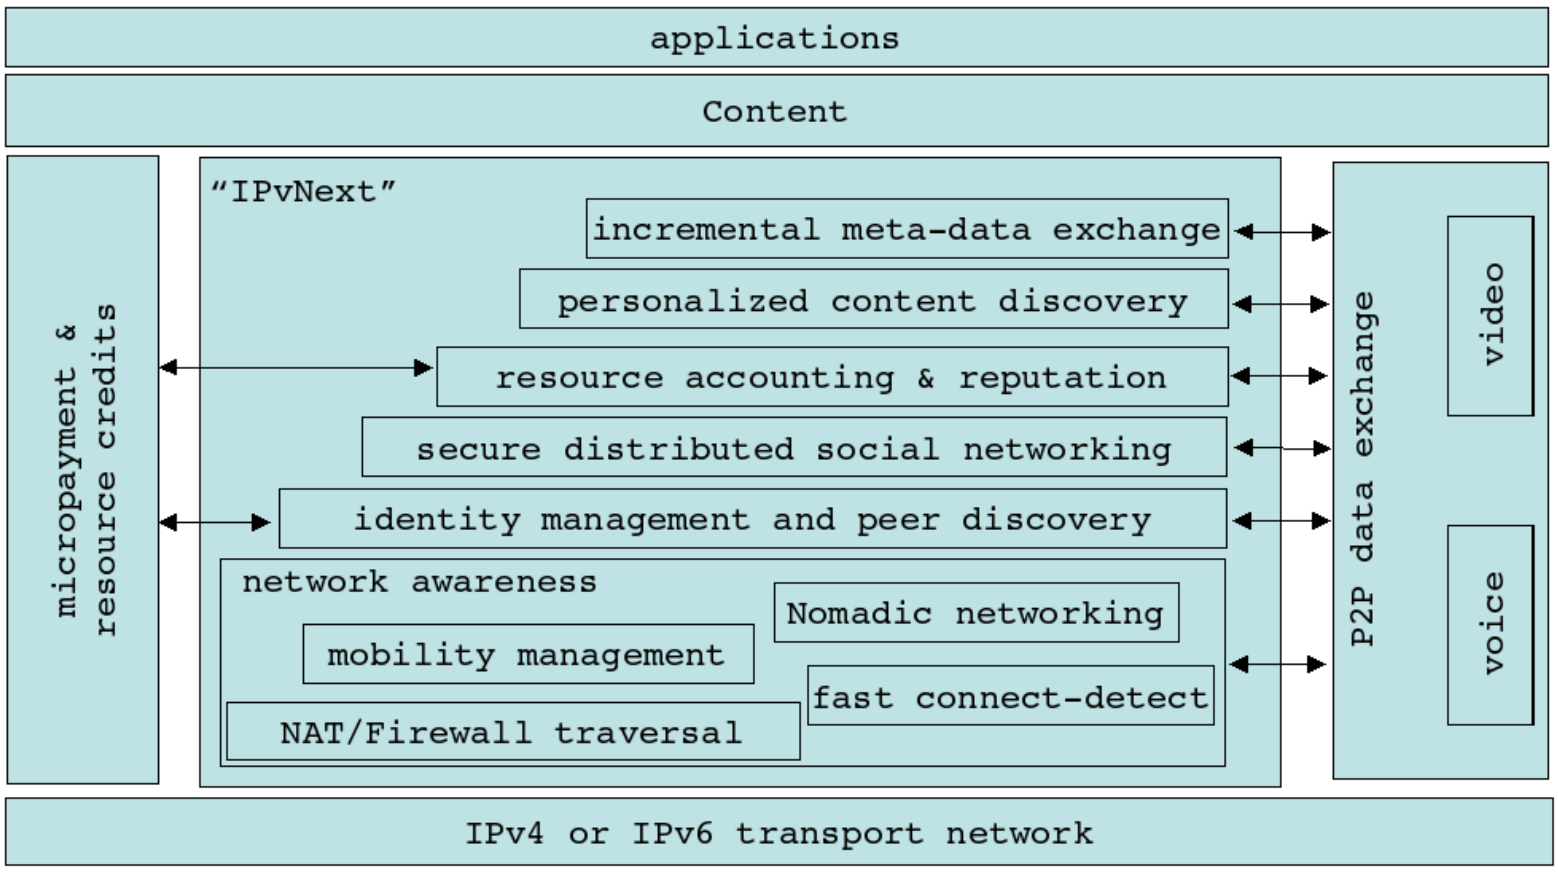

Figure 1: Core modules of Next-Share

The Next-share software can be divided into separate modules and functionalities depending on the specific end-user interface. While on one side we have the typical PC interface, provided by the WP6 (the Next-Share PC), on the other side we have the interface designed for CE devices, representing the future delivering tool for multimedia content.

The two remaining modules represent the common functionalities for both interfaces. The NextShare Core, goal of the WP4, represents the common distribution tool for delivering and injecting content through the desired interface of the Next-share product. On the other edge, as an overlay to the interface, we have the common Services shares between the PC and TV interfaces.

The key point of this architecture is its extreme modularity.

The design has been kept modular, allowing partners to develop their solutions in parallel when possible. A road map maps the design out in time, allowing for several iterations in the development phase. This might again lead to an updated architecture specification, roadmap and design. It was expected that only more advanced functionality would be adjusted by these updates, in order to not disturb the development process of P2P-Next.

The illustration below shows a general overview of the architecture.

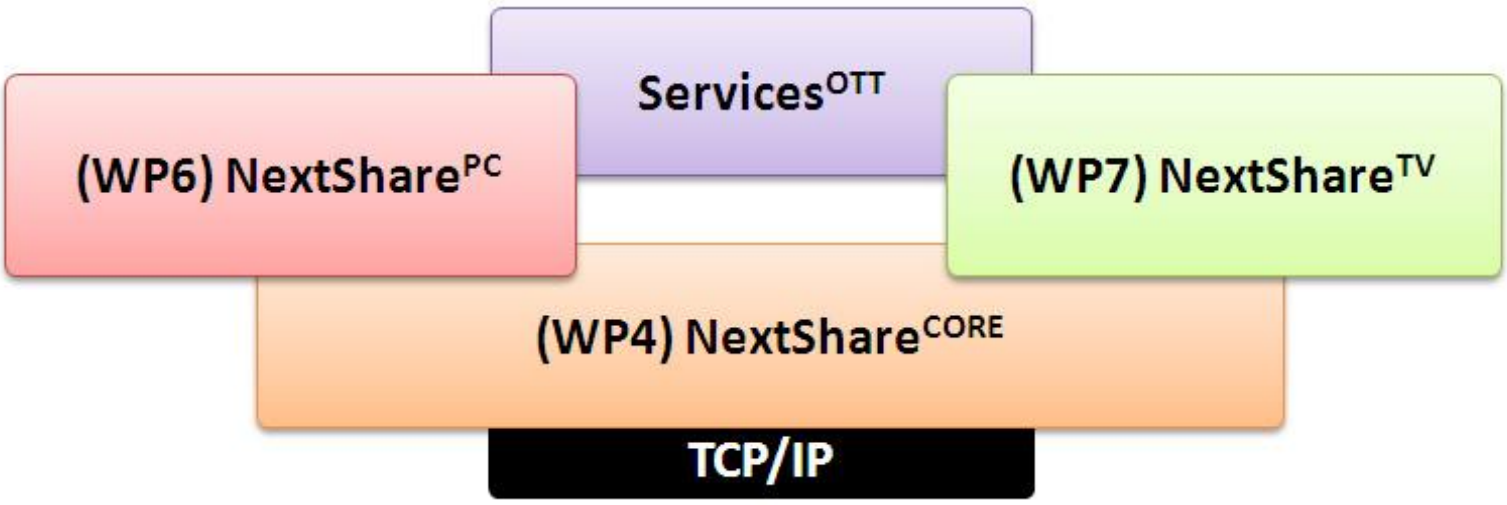

Figure 2: NextShare Architecture

The architecture and design specification has been the basis for WP4 (Implementation) work performed within the project.

Each of the elements is detailed in the 3.1.1b deliverable.

Work package 3.2 has produced three deliverables: the roadmap and design document (three iterations).

3 Executive Summary of WP4: "Peer-to-Peer and IPv Next Networking Fabrics"

Work package 4 comprises the technical core of P2P-Next.
The P2P-Next content sharing platform is called: **Next-Share**, which is a self-organizing system with complete decentralisation and hence lacks any central bottleneck or choke points that may hamper performance, induce setup cost, or require maintenance.

Three challenges drove the WP4 effort: 1) Living Lab code, 2) Next-Gen engine code (Libswift) and 3) forward looking ground breaking science. Resources within WP4 where stretched to the limit in 2010 due to the need of having to divide our limited resources over these challenges. We dubbed this the *three challenges problem*. During 2011 Libswift became mature and was promoted on 19 December 2011 to become an upcoming IETF Internet standard. In the final phase of the project the team focussed on near industry-grade streaming performance and real-world deployment.

A great deal of effort was spent on integration of WP4 with other workpackages, the first challenge. In the third year of the P2P-Next project we managed to deploy the complete end-to-end architecture in the WP8 LivingLab. During M8 of this project the Next-ShareCore for P2P live streaming was created. This M8 core is improved and refined in several WP8 Living Lab trials. Second challenge was creation of the Next-Generation Next-ShareCore, named libswift. Libswift is our new clean-slate P2P engine. These challenges are a cause of tension between WP4 and WP8 where support for the old M8 core needs to be balanced against the investment into the new clean-slate implementation. Furthermore, time spent on refining the libswift implementation into nearly production-level software and producing clear documentation implies less time can be spent on re-thinking the foundations, architecture and scientific output in general. The third challenge is conducting forward looking scientific research and obtaining the breakthroughs required for high-profile Journal publications. As indicated during the project reviews in Brussels the scientific output required a significant increase, both in volume and quality of publication venue.

Fortunately we succeeded in combining scientific breakthroughs with creation of an novel clean-slate P2P engine and supporting LivingLab with a reliable 24/7 operational platform. In the first year WP4 scientists operated mostly as individual groups at various locations. During the second and especially the third year both software and scientific publications became a team effort. In the final year, the team delivered beyond expectations. The following WP4 highlights after years of work could not have been achieved by these groups working in isolation:

- Realisation of a complete end-to-end platform. The successful deployment of WP4 next-generation technology in the LivingLab validates our assumptions and scientific findings.

- Creation of the first full-featured P2P engine capable of running on embedded devices. Libswift is the first P2P streaming technology to run on an iPad device, which is available on the Google app market for Android smartphones

plus the TV settop box created by WP7. It does not incrementally improve the state-of-the-art, but represents the next generation.

- Libswift is accepted as an upcoming IETF Internet standard. https://datatracker.ietf.org/doc/draft-ietf-ppsp-peer-protocol/

- Detailed power consumption experiments with smartphones shows the strength of our technology. Smartphone power consumption while streaming video using our P2P engine is equal or sometimes even superior (less power) to streaming from a CDN in a dedicated app.

- Wikipedia.org now runs our P2P browser extensions on their production servers. Several thousand of people installed WP4 technology to stream videos on wikipedia in their browser (http://arstechnica.com/tech-policy/news/2010/09/wikipedia-now-distributes-its-videos-using-p2p.ars). A new startup exclusively relies on P2P-Next technology: www.bombykol.nl Public trials with Dutch public broadcasters are conducted with several thousand people (http://torrentfreak.com/dutch-public-television-tries-bittorrent-downloads-100629).

- The WP4 team managed to get several algorithms and solutions published at the top conference in the field for several consecutive years: IEEE P2P. Over the years we created several high-impact publications, including the highest venue in the field: **SIGCOMM**.

- Significant advances are made in the area of WP4 streaming algorithms. We devised an improved piece picking algorithm, offline flashcrowd detector and improvements to peer selection algorithms for improved streaming support during flashcrowd conditions.

- Our self-learning reputation algorithm in a fully distributed setting is expected to be ready for experimental validation soon. Current progress levels allow for a publication in 2012 with the strong claim as “the worlds-first deployed distributed reputation system with real-world usage”.

3.1 Trials and workpackage integration

Numerous trials have been conducted to test WP4 technology in the Living Lab. Firstly, in M8 the novel WP4 unifying algorithm supporting Bittorrent downloads, live streaming and video-on-demand was tested. Numerous people participated in this first test. Over 50.000 unique IP numbers successfully installed and ran our code. This first trial was a big success and proved the validity of both our algorithm and implementation. In M14 we conducted a trial with NAT/Firewall testing involving 3500 people. We determined that 90% of those participants where behind a NAT or firewall and that 40% of those where of the “port-restricted cone” type and overall a mere 9% of participants had a NAT timeout setting of 1 minute or less. In M16 we conducted a first NAT puncturing trial and successfully established a connection between 1800 unique pairs of peers both behind NAT devices.

In M24 a large public video-on-demand field trial was conducted centred around WP4 technology in collaboration with other work packages. Content was provided by the BBC, along with encoding and packaging in WP5. WP6.5 was responsible for code release management, the software installer and browser integration. WP8 software was successfully used to recorded valuable behaviour of the over 1000 people which successfully installed our Next-Share software. Key outcome of this trial is that our video-on-demand algorithms operated as expected outside laboratory conditions. However, performance needs to be improved further. We recorded the average time between initiation of video-on-demand and the actual start of the playback, dubbed *playback delay*. This delay was measured to be 8.51 seconds on average.

Key conclusion is that for reduction of the playback delay we need to move to a permanent state of significant overcapacity or overseeding in the system. Playback delay should be in the order of 2-3 seconds for acceptability. To stimulate people to seed we need to move beyond the limits of tit-for-tat and direct reciprocity. A reputation system is needed for enabling indirect reciprocity by keeping track of upload and download ratios in a robust manner. Significant progress has been made in this direction, see additional detail in the WP4.3 section.

The first PhD. thesis which is financially supported by P2P-Next has been finalised at the end of 2009. This thesis by Dr. Mugurel Andreica includes experimental research conducted within the P2P-Next living lab on novel network congestion control algorithms.

For our final technical trial we demonstrated the maturity of our WP4 P2P engine on mobile devices. During M52 we deployed a Libswift streaming application on the Android market (play.google.com), as shown in the screenshot below. Furthermore, two key consortium members secured funding for 2012 from the European Institute of Technology (EIT) to make this near industrial-grade Open Source technology further ready for commercial usage.

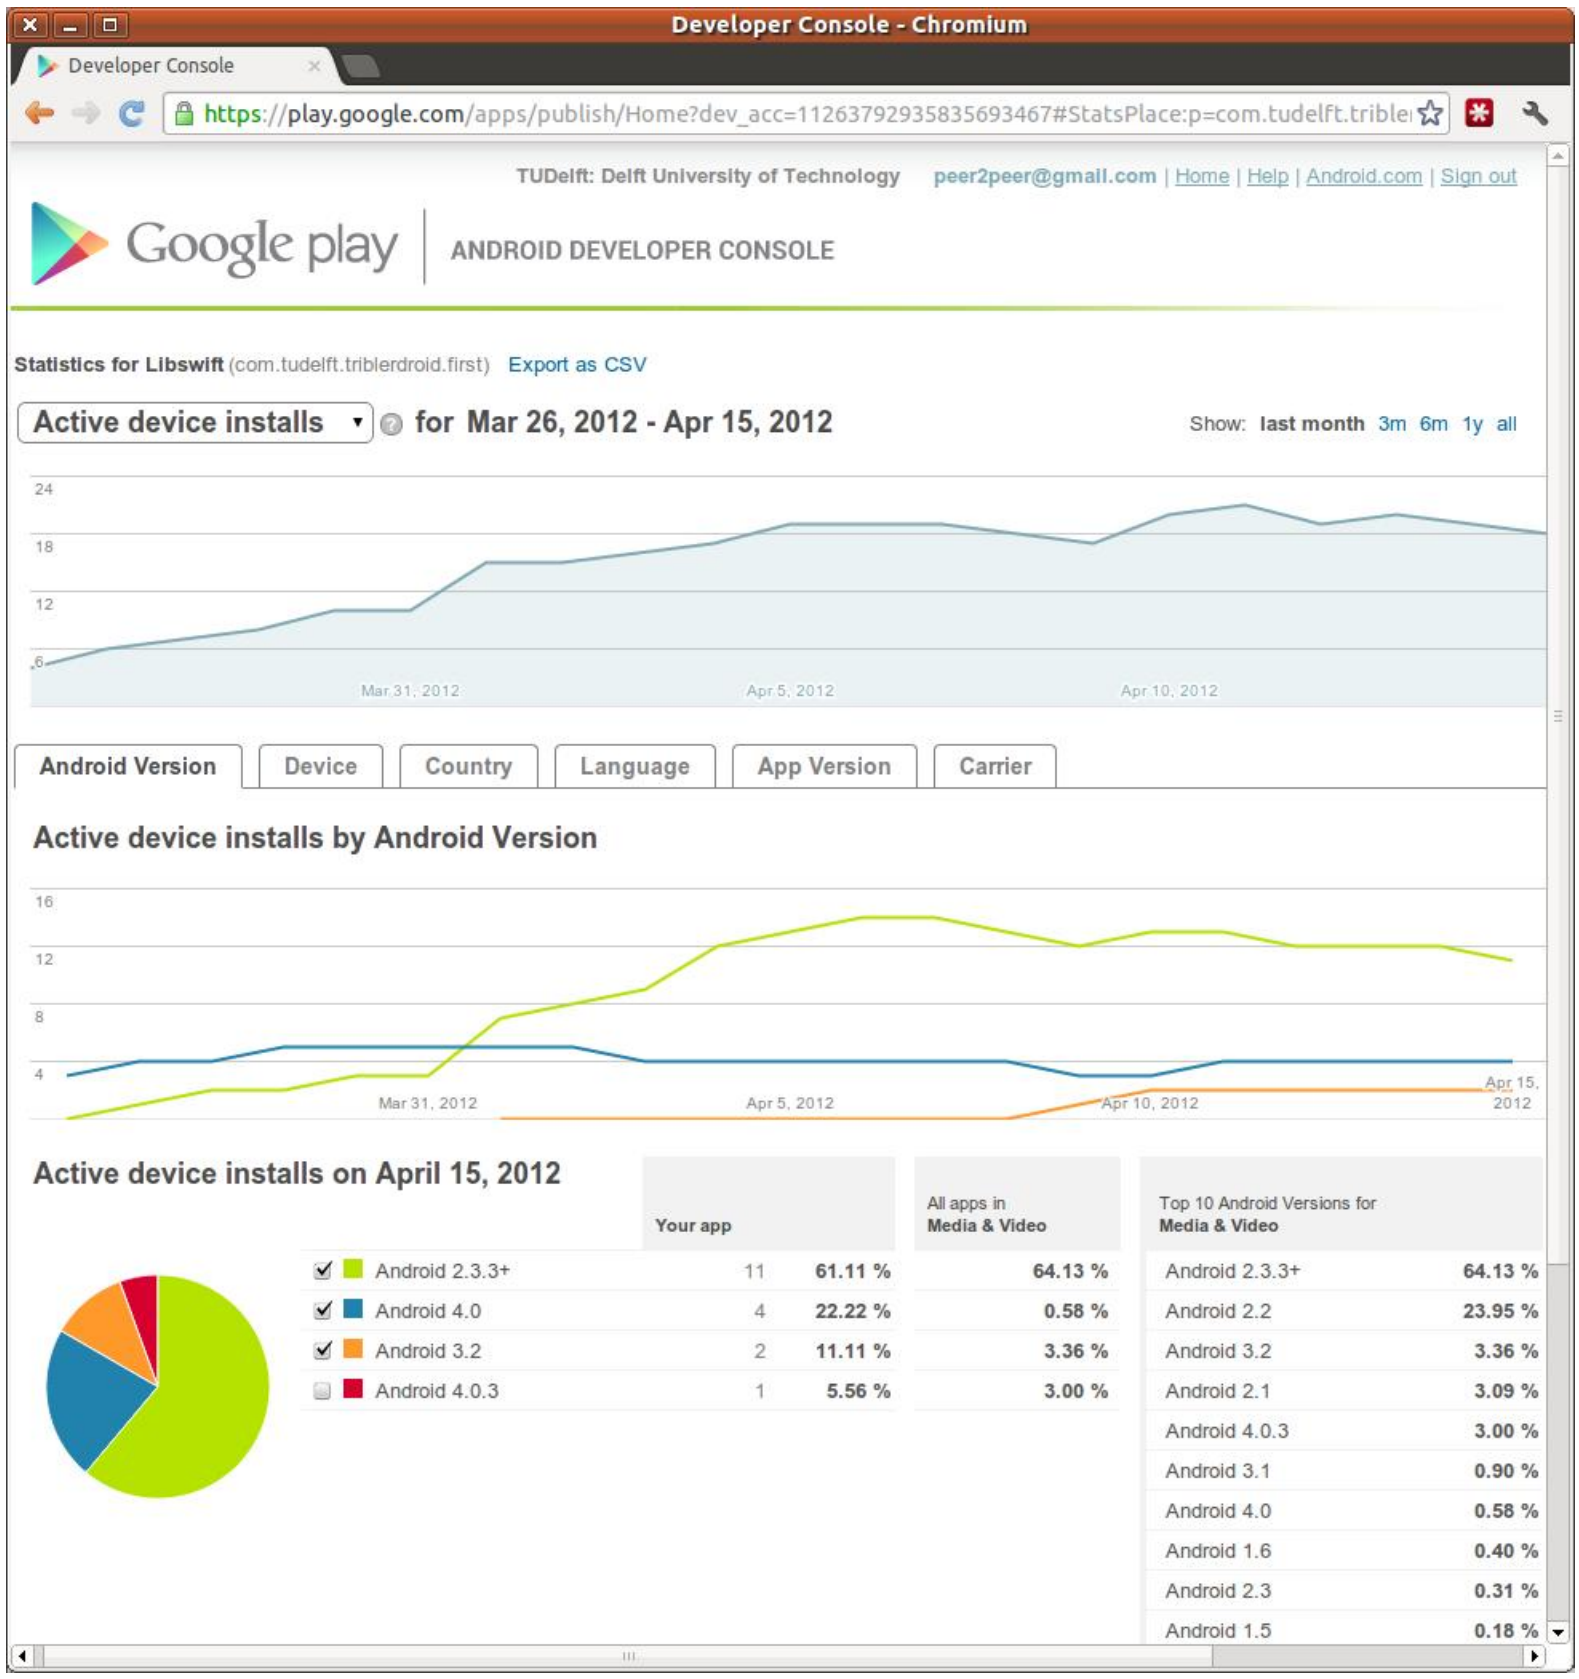

Figure 3: Open trial of P2P engine for Android devices.

As shown above our streaming demonstration application is downloaded several times. The first curve in this Figure shows the number of active devices running our Libswift app over time, plus a second curve providing a breakdown per Android OS version.

3.2 WP4.1 IPvNext networking fabric design, implementation and testing

Peer discovery has been proven to be a difficult P2P problem. Commonly, the peer discovery process relies on a centralized server called *tracker*. A tracker is a centralized web server which keeps track of every peer in a given swarm. Setting up a central tracker is technically complex and too difficult for non-technical users. There is already a distributed tracker system supported by most of the popular BitTorrent clients. This system is based on *distributed hash tables* (DHT). The problem with

BitTorrent's DHT is that it is very slow and unreliable. According to some studies it takes around ***one minute*** to get a list of peers from the DHT. This delay is, of course, unacceptable for the Next-Share platform.

WP4.1 has shown that it is feasible not only to transition the DHT used by BitTorrent from a role as a secondary fall-back mechanism, to the primary peer discovery mechanism, but to improve it even further to where it can play a role in on-demand streaming applications.

Our initial study of MDHT node implementations revealed that Utorrent is the most common implementation currently in use, with a measured "market" share of over 60 %, and also the best (fastest) of the implementations in common use according to our measurements. Clearly, this makes UTorrent the ideal candidate as the state-of-the-art benchmark for us to beat. Our most aggressive implementation is remarkably faster then UTorrent's. Not only are our median lookup time only 1/6 of UTorrent's, but our 99\% percentile is just below 500 ms when UTorrent's is over 4 seconds. While this comes at a somewhat higher lookup cost (257%), when we consider both lookup and maintenance traffic, our implementation actually uses significantly less traffic (1/3) than UTorrent.

Resource accounting and reputation

Finding peers with a sufficient reputation has proven to be a hard problem. Significant resources are put into developing a fully distributed reputation system in P2P-Next. Prime focus is on enabling video-on-demand by providing an incentive for seeding and preventing hit-and-run behavior. Therefore, discovery of honest and cooperative peers is key.

Many proposals of Internet reputation systems has been proposed, some of them made it to a prototype phase, but no system exists which actually works.

This research is split into 1) validating the existing Tribler public/private key system [project background knowledge] and signature mechanism for it's security strength, 2) understanding the limits and enhancing our self-reporting framework called Bartercast, and 3) building policies which rewards good seeders with a robust video-on-demand experience. An implementation has been made of this resource accounting framework. Depicted below are 11 peers along with the data transfers for video-on-demand and epidemic protocol for self-reporting using bartercast messages.

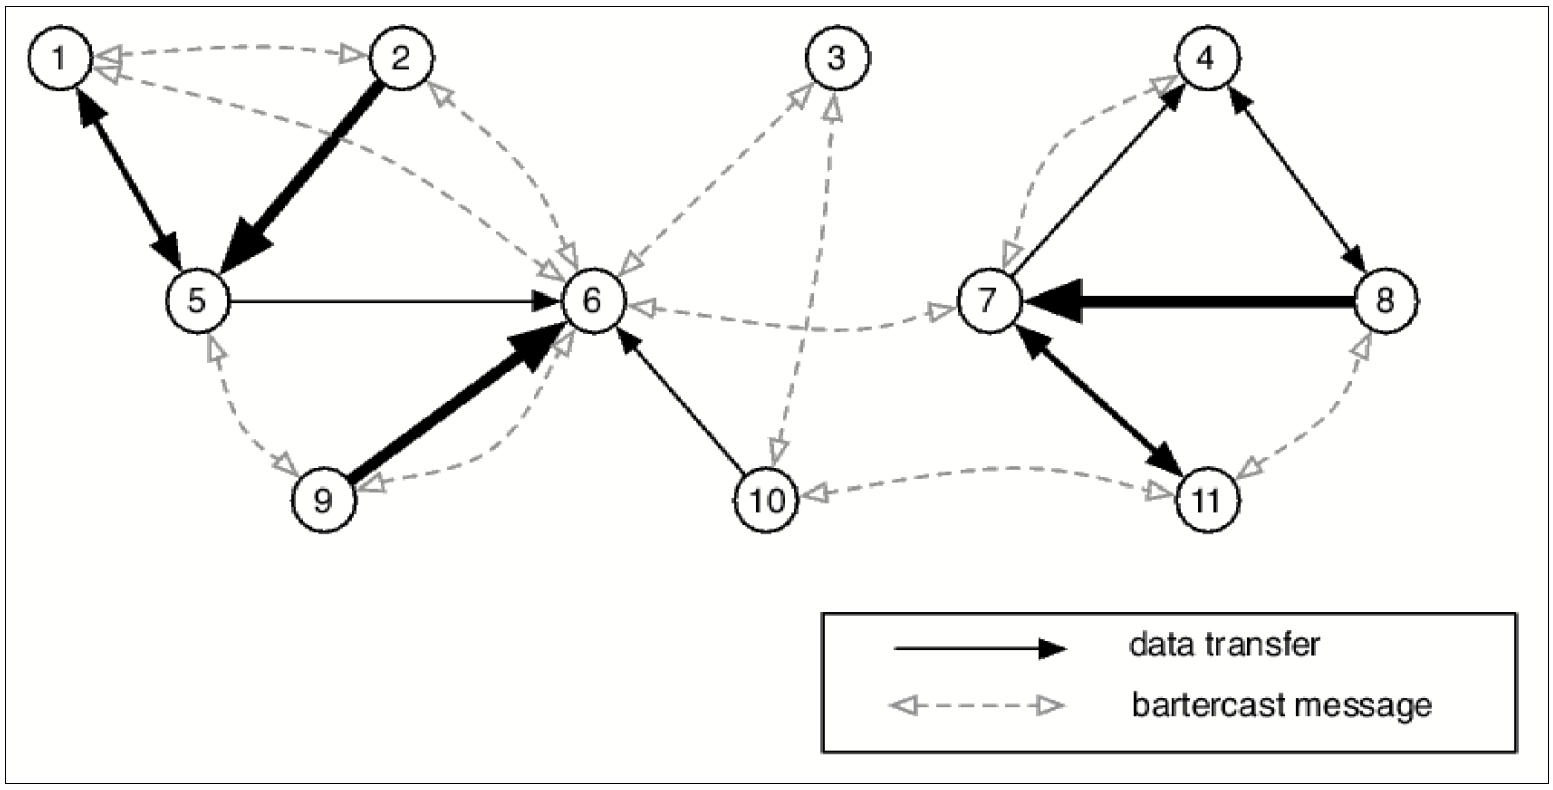

Figure 4: : illustration of data transfer between peers and exchange or BarterCast messages.

The target is to ensure that hit-and-run is limited by ensuring freeriders are caught. By combining many of the above BarterCast messages it is possible to detect freeriders from good contributors. Shown below are the results of years of incremental improvements to our scalable and distributed reputation system. We focus on the crucial problem of providing peers with all the required information for accurate reputation evaluation (e.g. BarterCast messages). Scalable solutions require that a single peer only needs to store a subset of all available reputation input information from all other peers.

Furthermore, as the number of peers increases this subset should be stable. The mechanism to select the subset is critical to the accuracy of a scalable system. We pioneered a similarity-based approach for targeted dissemination of reputation information. In our approach each peer maintains a list of similar peers and gives higher priority to those when exchanging messages and thus building a graph.

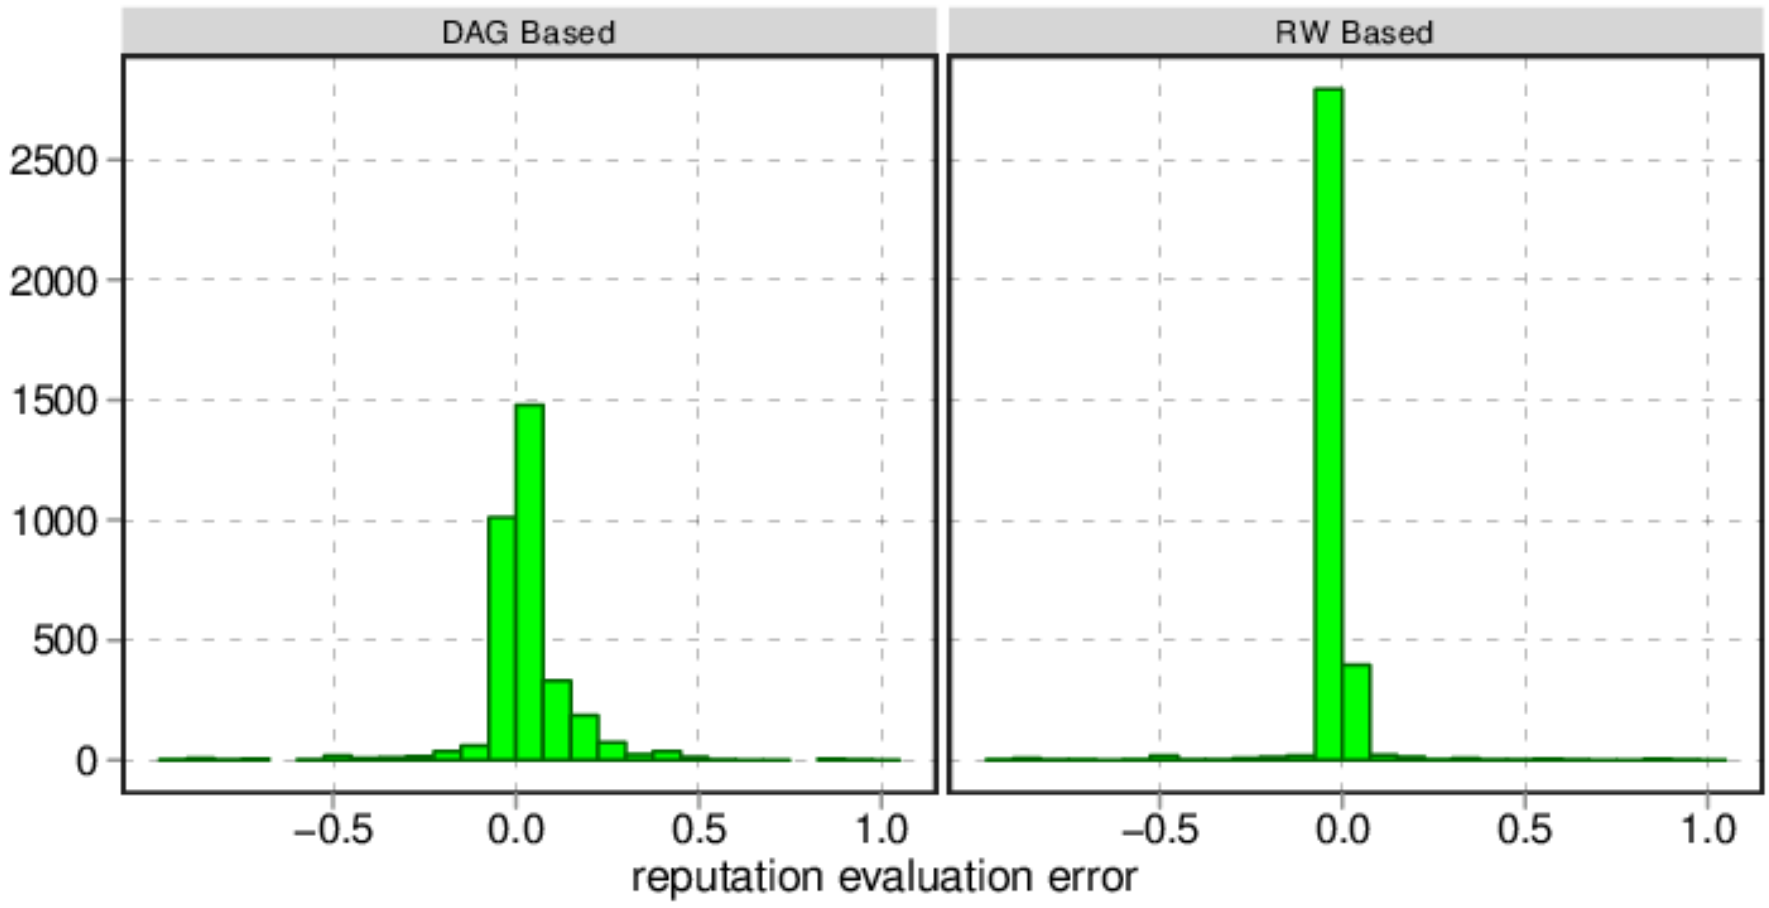

Figure 5: Accuracy of our distributed reputation system

Shown above are the results using two methods to derive peer similarity in the partial graph of a peer. The first method is based on creating and incrementally maintaining a directed acyclic graph (DAG), and the second method is based on using multiple non-uniform random walks (RW) in this graph. We evaluate the accuracy and the overhead of the new dissemination using trace-driven simulations based on traces from the Tribler network that uses Bartercast. As the results shows, our method has only a marginal error. Due to the usage of subsets it also has very low overhead and performs like the case of providing complete knowledge to all peers.

3.3 WP4.2 Network awareness, locality and geography

Understanding the network is the key to efficiency and high performance when using P2P technology. An extensive test infrastructure was developed to understand the effects of the network.

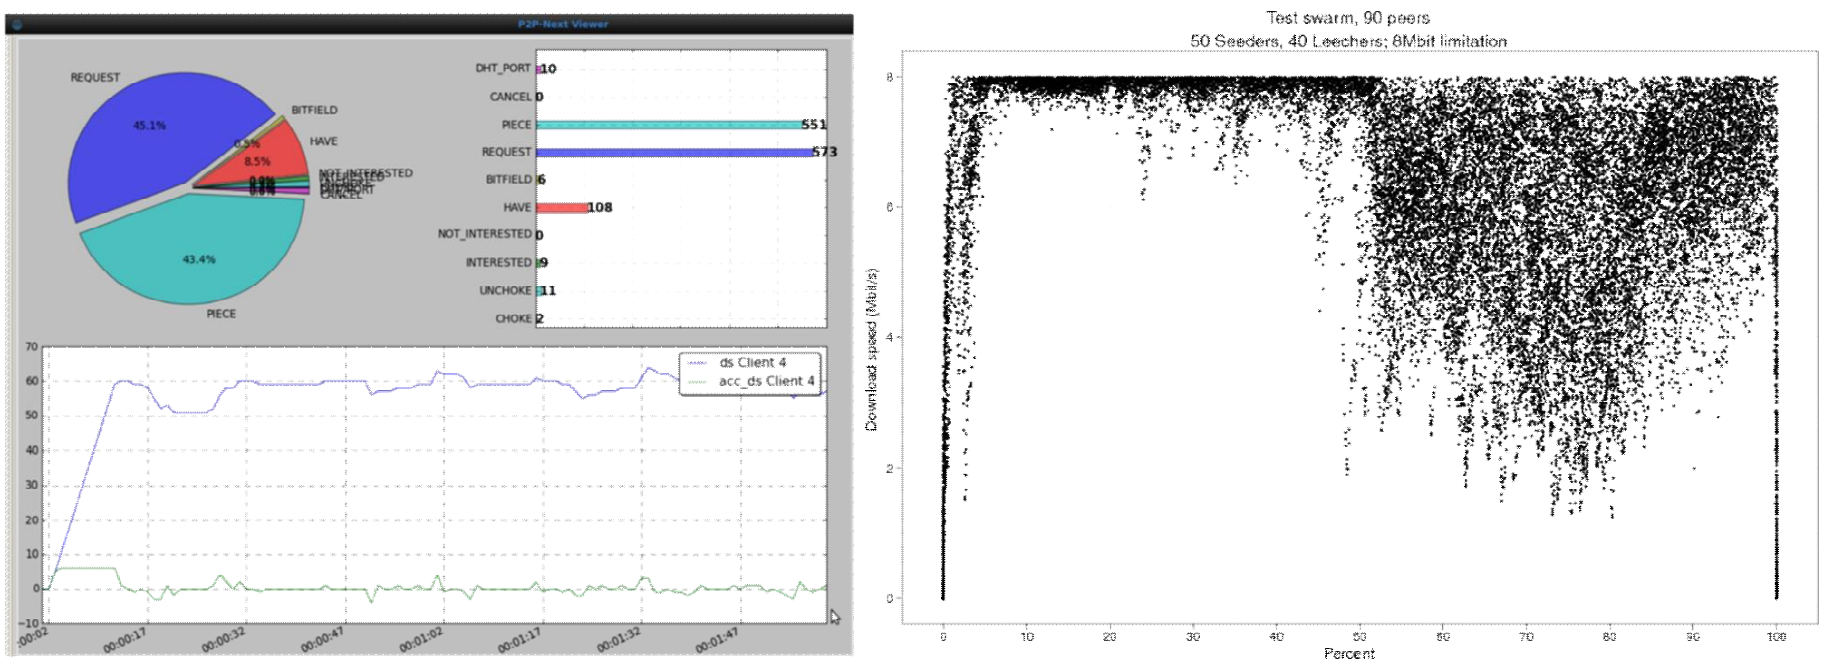

Figure 6: Network utilization in the P2P testing.

Our P2P Testing Infrastructure for starting and managing the swarms has been deployed on a hardware infrastructure by consortium partner UPB and used for various experiments regarding P2P swarms and their characteristics. The above right image present the download speed evolution with respect to percentage of a 90 peer swarm (50 seeders, 40 leechers). The image on the left presents the GUI component of the message processing engine (allowing parsing, analysis and graphical rendering of protocol messages), displaying the types of BitTorrent messages and their number and the evolution of the download speed and download acceleration (variation of download speed).

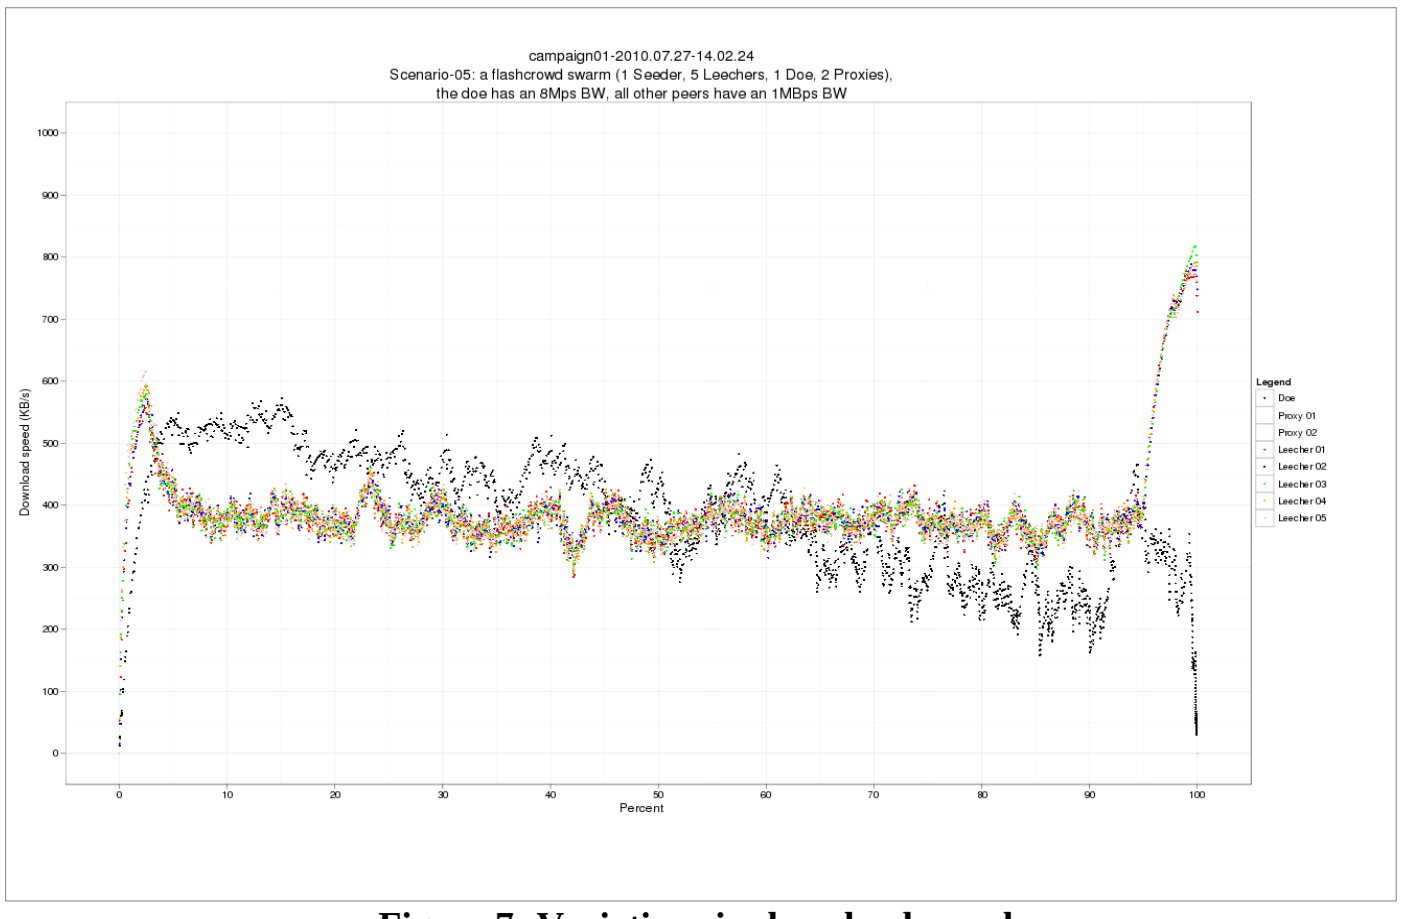

Figure 7: Variations in download speed.

The ProxyTechnology was improved and integrated in the NextShare core. This technology is capable of overcoming strict limits due to limited upload capacity of Internet peers. This unique technology enables the boosting of download performance by introduction of multiple relays or proxies.

A command-line version of the NextShare was used for proxy performance evaluation in various scenarios. The complex nature of the testing environment needed required a dedicated testing infrastructure to be designed and implemented. The P2P Testing Infrastructure presented above was used to run a number of experiments that offered a feedback on the behavior of the Proxy core (above image).

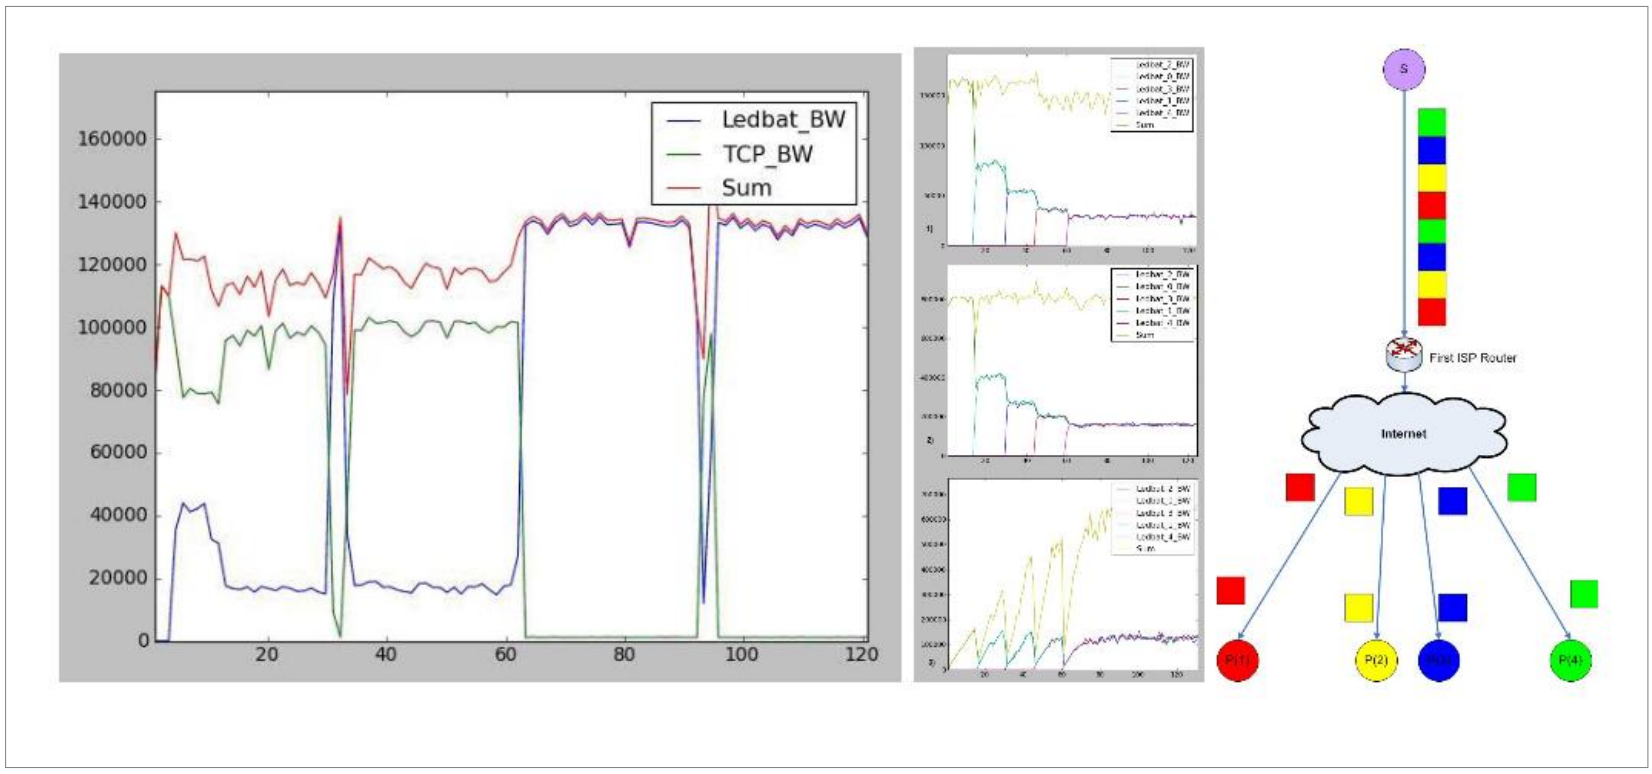

Figure 8: Testing of the LEDBAT network.

A Python implementation of the LEDBAT network congestion control algorithm was developped. It was evaluated with real users and in emulated netwok environment, based on the following metrics: behavior towards TCP, performance and fairness. A generic data transport protocol development framework was also developped (in Python, over UDP). The implementation was used to test multiple congestion control algorithms. A collaborative technique for estimating the upload capacity of a machine was also developped.

3.4 WP4.3 P2P Data exchange engine for data, voice and video

The key effort within WP4.3 is focussed on facilitating seamless and robust Video-on-Demand. After our experiences of the first year we were able in the second year to start work on a new clean-slate P2P engine based on a novel protocol named Libswift. The Libswift architecture is shown below.

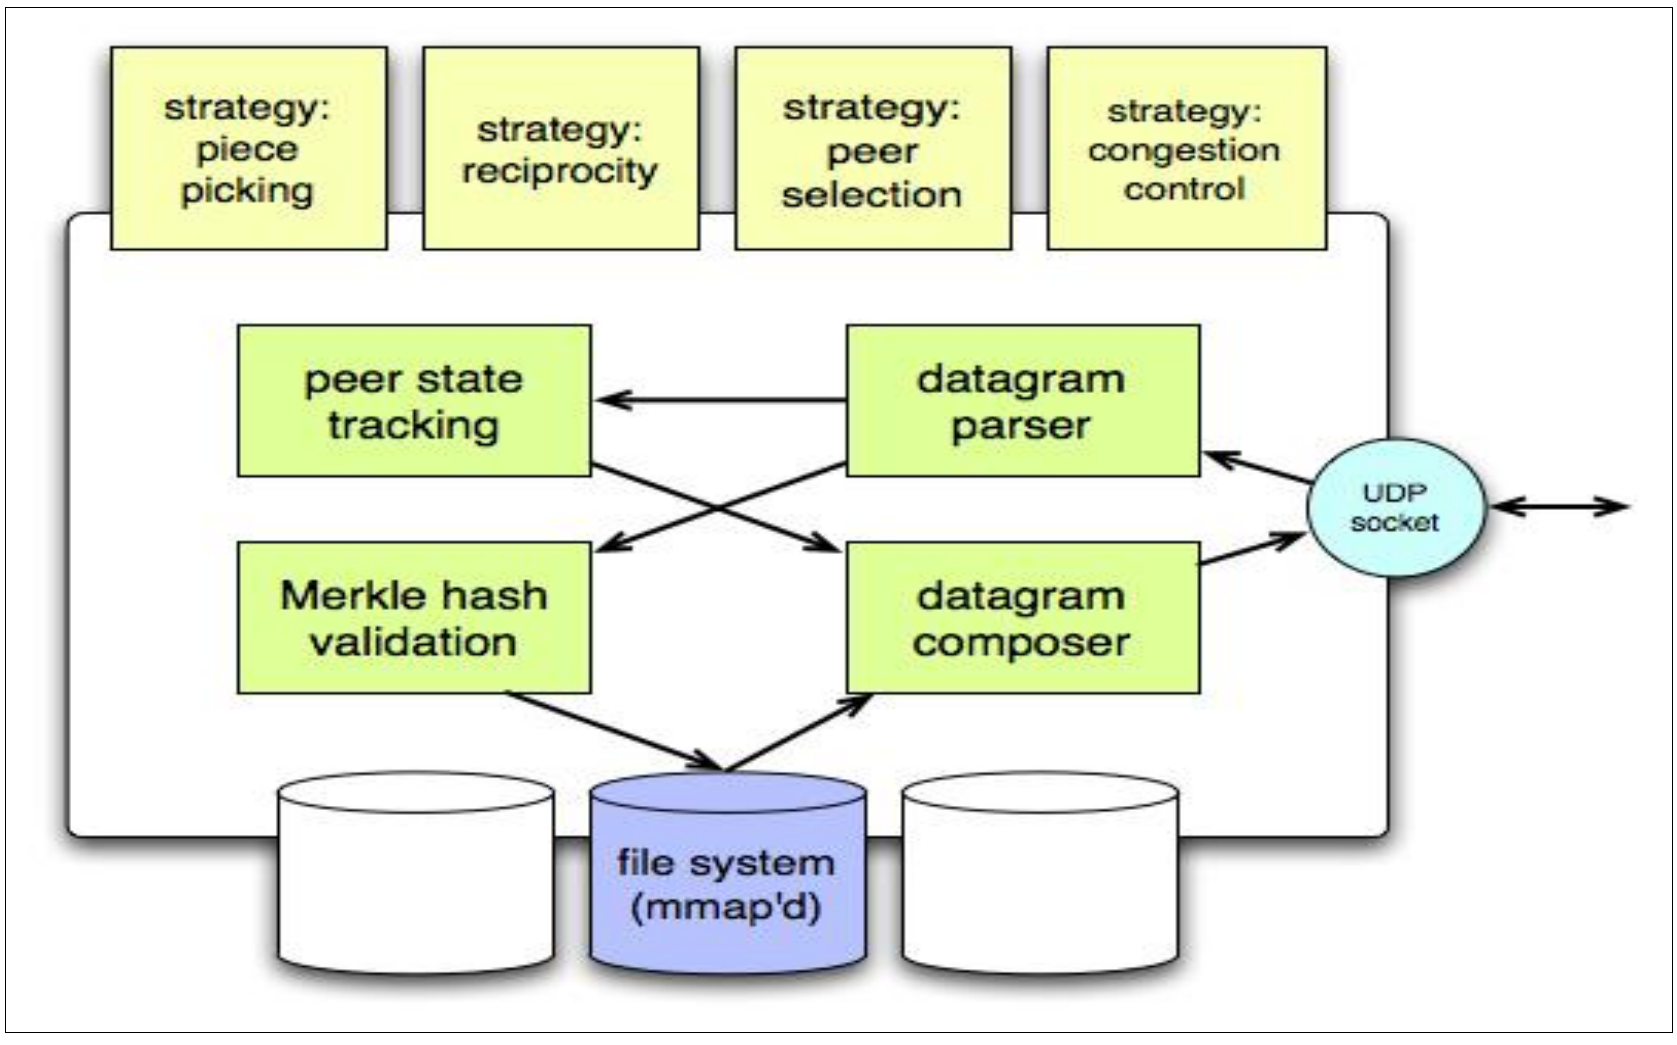

Figure 9: The Libswift architecture.

The objective of Libswift is to create a content-centric multiparty transport protocol to allow seamless, effortless data dissemination on the Internet. Primary usecase for the protocol is video-on-demand (VoD). The protocol is architected for ubiquitous, embedded use in internet browsers, operating systems and set-top boxes and thus compliant with the P2P-Next ambition level and requirements. The protocol is architected to be technically generic, not bound to any single scenario or usecase. It might be employed in the settings of content distribution, peer-to-peer or peer-assisted network. Considering that embedded use is the priority, the protocol is optimized for low footprint in multiple aspects, including non-intrusive congestion control. Because of the current network realities, NAT hole punching is considered a part of the protocol's core functionality. The technical objectives led to a decision to abandon TCP as an underlying transport protocol.

While TCP supports the abstraction of pairwise conversations, we need a content-centric protocol built around the abstraction of a cloud of participants disseminating the same data in any way that is convenient to them. Thus, the choice was made to use UDP, the low-level datagram transport, which provides better possibilities for custom

congestion control, NAT penetration etc. That in turn led to the concept of atomic unreliable datagrams (as opposed to the data pipe). An atomic datagram is either immediately accepted, stored and relayed or dropped. No data context spans multiple datagrams. The objective of non-intrusiveness assumes that the protocol instance might work round the clock, still it must not cause degradation of shared resources, primarily the Internet bandwidth.

To meet the objective we employed the LEDBAT congestion control algorithm (low extra delay background transport) according to the internet draft. To fit the protocol into the constraints of the transport layer it was made free of technical metadata. It was achieved through the use of Merkle hashes, exclusively single-file transfers and additional techniques. As a result, a transfer is identified and bootstrapped by its root hash only. As the startup delay is a key quality metric for VoD applications, special steps were taken to minimize the protocol's warm-up time. The entire protocol's state machine was implemented in March 2009, first successful transmission being done on April 2009. Libswift has been successfully ported to the Apple iPad in 2010. This proves both the viability of our concept and the light-weight nature of our implementation. Near the end of 2010 we also ported Libswift to Android-based mobile phones. This is significantly more challenging as the iPad has a 1 GHz processor and abundant memory, simple threading model and superior support for C++ and standard development libraries when compared to smartphones. Finally, on April 2011 Libswift was launched live to the Android app market.

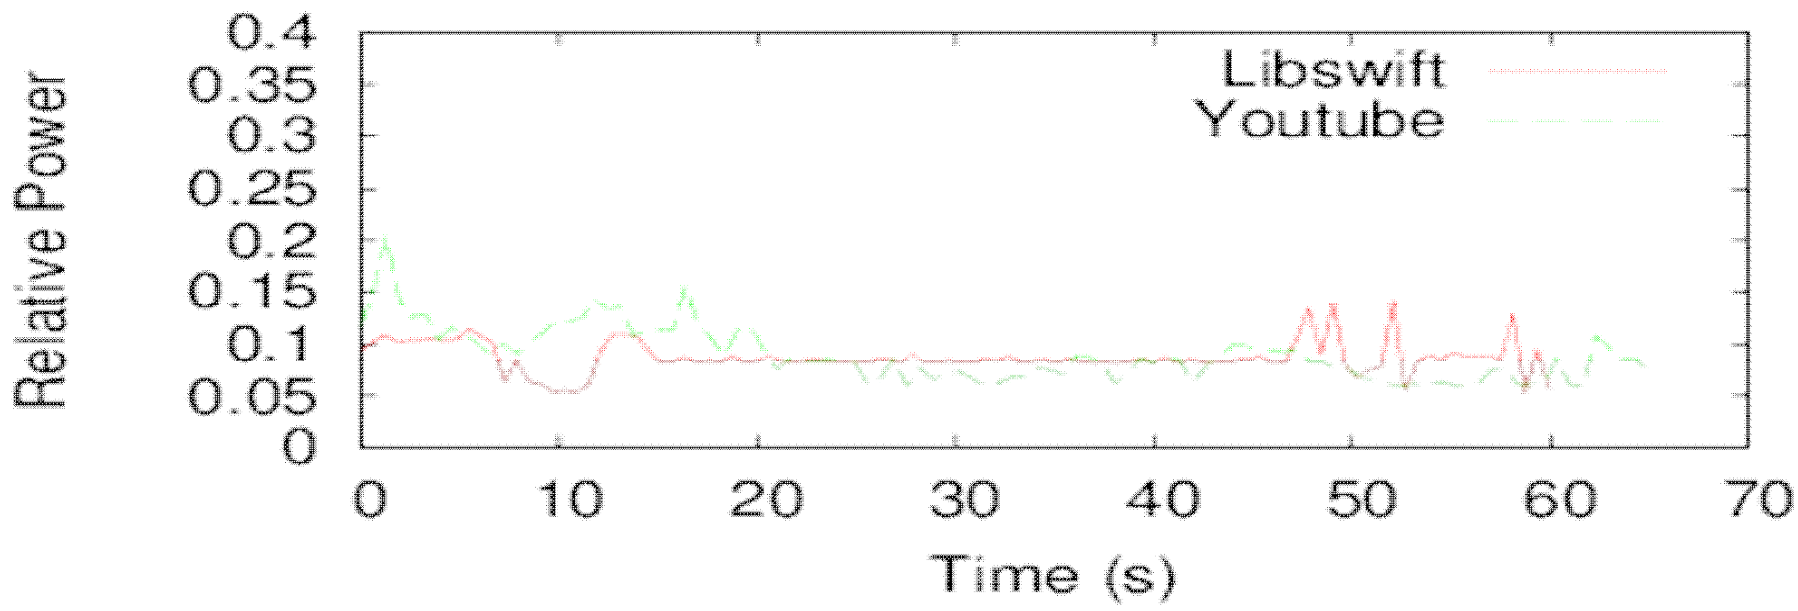

Figure 10: Relative power consumption of Libswift and Youtube on a Samsung Nexus Android device.

Shown above is evidence that peer-to-peer video streaming is comparable or at some timepoints even superior in terms of power consumption to CDN-based video streaming. For this experiment we measured the power consumption of the flagship Android 4.0 device. We recorded the power consumption (without costly accurate absolute power calibration) while streaming using both Libswift and Youtube.

Especially upon startup Youtube is more costly then the NextShare technology.

Libswift has been ported succesfully on not only iPad and settop boxes, but also the caching hardware of Oversi. This fully Open Sourced P2P Cache is delivered within P2P-Next as a modified subset of Oversi's flagship OverCache P2P caching and delivery platform, currently deployed at numerous Tier-1 service provider sites around the world. P2P-Next Cache is targeted at an ISP requiring an integrated solution for P2PNext traffic optimization and caching within its network core.

P2P-Next Cache supports not only the ubiquitous BitTorrent based VoD and the P2P-Next live streaming BitTorrent flavor, but also the new NextShareCore Libswift protocol. This allows the ISP to use a single unified solution for P2P-Next traffic optimization and caching within its network core. P2P-Next Cache performance is tested over a relatively low-end platform (2.2Ghz Core 2 Duo, 8GB memory, 4x500GB disks). The Cache performance is close to 1Gbps while providing data from memory (where there are less than 500 peers) and around a little more than 400Mbps when disks need to be read.

As a final breakthrough, WP4 managed to demostrate real-time live streaming with Libswift at an IETF event. In M52 we have succesfully implemented both live streaming and on-demand streaming in a single code base.

3.5 WP4.4 Micro-payments and resource credits

The objectives for WP4.4 are to foster sustainable business opportunities and additionally support payments between any two users in NextShare. The NextShare micropayment system has been up and running for quite some time and incremental improvements have been made. The payment system architecture is presented in the figure below.

The payment system architecture is presented in the figure 11 below.

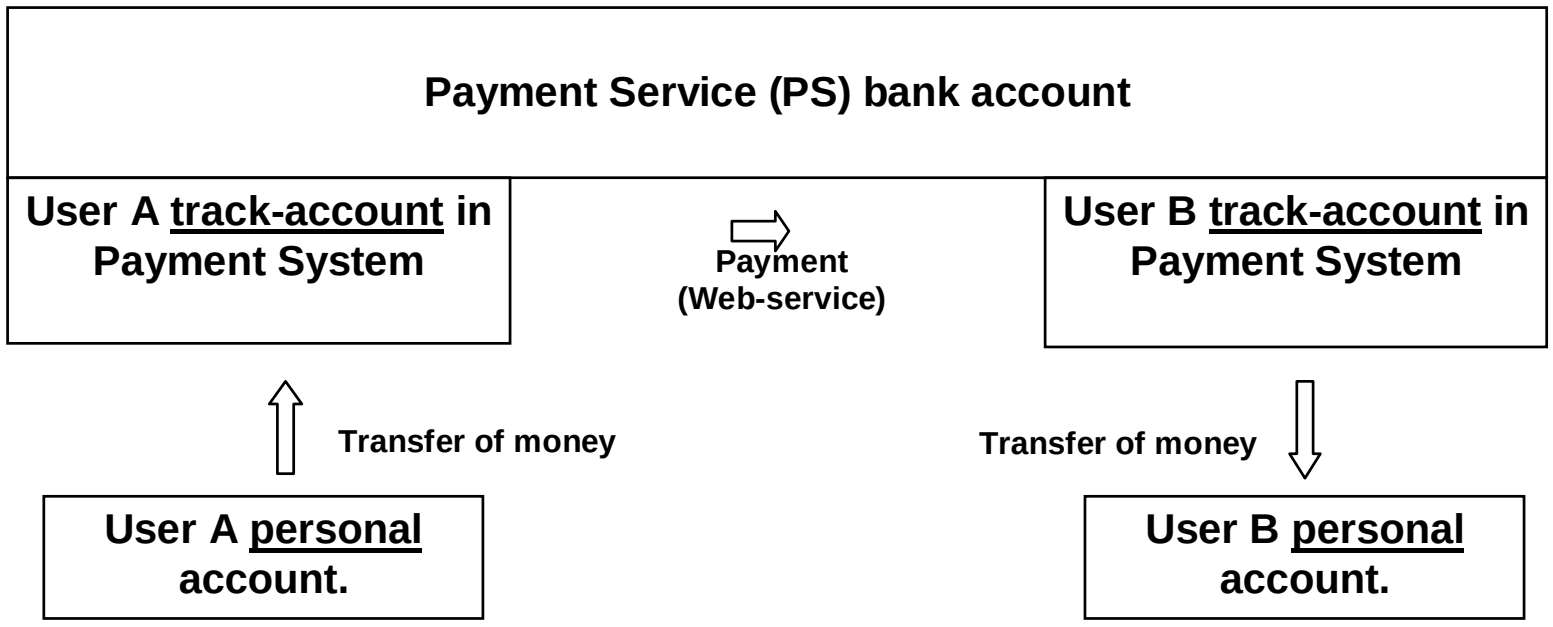

Figure 11: The payment system architecture.

In the payment system there is a “buyer’s” side (A) and a “seller’s” side (B). Every buyer has a User A track-account and similarly every seller has a User B track-account. In the following no difference has been made between content providers, service providers, providers of experiences and ordinary users (consumers). They are all called users. The rational is that ordinary users are expected to contribute in NextShare as prosumers in a similar way as we see users participate today in occurrences like Youtube, Facebook etc. In the following content will be used as an umbrella term for experience, service and content.

Within our paymnet system we implemented pre-paid accounts, ease-of-use registration functions, mobile phone payment support, integration with already existing mobile micro payments such as SMS payments, and a prototype Android mobile payment application based on credit cards.

Our mobile micropayment technology does not require a CVV/CVC code, similar to payments in parking meters, vending machines, etc. The (Android) payment application is installed from a web site and/or a “Market”. At the first run the user data gets registered. Registration takes place on the phone. Credit card details are divided and one part sent to the sales server and stored there and the remaining part in the phone. The credit card details are complete only in combination. A push notification starts the application when the user chooses to pay with the application on a web page or elsewhere. All purchases are stored in a payment queue specific for each application. The user requests the payment queue on the phone and can choose to approve or cancel a payment and this action is immediately recognized on the sales system, typically a web page. Approval means that the remaining part of the credit card details stored in the phone is transmitted to the sales server. The sales server can then charge the credit card and complete the purchase. The application can be protected by a code (PIN or alpha numeric) if the business approaches so requires.

We conducted a series of payment trials. Our largest trial revolves around the concept of “Closed Swarms”, where peers receive preferential treatment and fast downloads if they are “authorised”.

The ability to distinguish between authorized and non-authorized nodes in a swarm is the most fundamental challenge, as this allows content providers control of their resources within Next-Share.
As a part of the commercial requirements, Closed Swarms was defined and implemented with the Next Share Core in 2009. This work has continued, with improvements to the implementation in the core, embedding into the NextSharePC branch of the project using HTML integration. In early 2010, a payment demonstration was performed as well, where a mobile phone payment scheme was demonstrated in collaboration with a graded Closed Swarm service.

In the autumn 2009 a student project was set up to integrate the KTH Access Control Server (ACS) and DACC Micro-payment System. The goal of the project was to demonstrate a number of business models – advertising supported free video content, pay-per-view and subscription while ensuring content integrity. During early 2010 the knowledge of the project was transferred to DACC.

In spring 2010 a follow up project with NORUT was set up where the Micro-payment System was integrated with Closed Swarms and used in a living lab test.

A pre-commecial test with NextShare PC and the payment system from DACC was planned to be performed during the second half of 2011 at the Swedish town Sundvall together with Mid Sweden University, Acreo (A Swedish research organization), and the local network operator Servanet. A first technical test (Alpha) was done in the beginning of 2011 together with Acreo within their closed test network in Hudiksvall in Sweden with about 50 users but no payment due to lack of valuable VoD material.

4 Conclusions about WP2-WP4

The P2P-Next project was carried out in 52 months starting in January 2008. Each WP represents an integrated part of the project. The three first technical work packages (WP2-WP4) have delivered the following main inputs to the final goal:

- Assessment of the legal structure of P2P applications
- Business requirements and models for the exploitation of P2P systems
- Observations of market trends in digital media and IT related media
- Market watch to evaluate the take-up of the P2P-Next applications
- Specifications of user requirements for P2P applications
- Specification for the architecture of the Next-Share platform design and the living lab trials
- Design of the content sharing platform Next-Share and its decentralised self-organising system
- Integration of network awareness, locality and geography in the system
- A P2P data exchange engine for data, voice and video
- A system for micro payment and resource credits
- Exploitation plan for the results including presence at industry events, websites, Tweeter and Facebook, version controlled repository, packaging for multiple operating systems, adaptation by Depian Linux Operating System, and developer mailing list

The work packages have fulfilled all their goals specified in the original plan and delivered the corresponding reports, hence contributing to the final goals of the P2P-Next project. The final conclusions of the project are given in the Final report (D9.4).