

Exploring New Metaphors for a Networked World through the File Biography

Siân E. Lindley¹, Gavin Smyth¹, Robert Corish¹,

Anastasia Loukianov², Michael Golembewski³, Ewa A. Luger⁴ and Abigail Sellen¹

¹Microsoft Research, ²University of Surrey, ³Royal College of Art, ⁴Edinburgh University
 {sianl; gavin.smyth; rcorish; asellen}@microsoft.com;
 a.loukianov@surrey.ac.uk; mike.golembewski@gmail.com; ewa.luger@ed.ac.uk

ABSTRACT

We present a body of work undertaken in response to the challenge outlined by Harper *et al.* in their paper, ‘What is a File?’ [9]. Through a conceptual and design-led exploration of new file metaphors, we developed the ‘file biography’, a digital entity that encompasses the provenance of a file and allows the user to keep track of how it propagates. We explored this through prototyping and utilised it in two user studies. In the studies, we (i) asked people to sketch out file biographies for their own content, and (ii) deployed a tool enabling users to build their own simple file biographies across multiple versions of Word documents. We conclude that new file metaphors may need to play different roles for different types of digital content, with a distinction being drawn between content that is ‘in production’ and virtual possessions that are, in a sense, a ‘finished’ artefact.

Author Keywords

Grammar of action; virtual possession; patina; ownership; control; provenance; propagation; version control; fork.

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Increasingly the content we produce and use as part of work and leisure is digital. It is stored on computers, mobile devices and portable storage devices, posted to social media sites, put into shared repositories, and synced with cloud services. The ways in which content propagates and is distributed across these stores, as well as being copied and repurposed by others, all contribute to a sense that digital content is difficult to manage and keep track of. The aim of the research reported here is to explore ways of tackling these issues, with the overall goal of supporting a user experience of awareness and control.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CHI 2018, April 21–26, 2018, Montreal, QC, Canada

© 2018 Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-5620-6/18/04...\$15.00

<https://doi.org/10.1145/3173574.3173692>

The approach we take draws on the challenge outlined by Harper *et al.* in their paper, ‘What is a File?’ [9]. They argue that there is a need to rethink the file metaphor, which they position as a boundary object [23] that serves as a point of connection between computer scientists, for whom the file is one abstraction for persisting a bundle of data, and users, for whom the file presents digital content as a singular entity. Developed to serve as generic objects in the days of the Xerox Star [22], files could be acted upon in predictable ways through a set of generic commands (e.g., move, copy, delete) that also aid the user in forming a model of the system. However, in the current, highly networked, world, where content can be distributed across and accessed via a plurality of technologies, Harper *et al.* argue that the file metaphor is losing its salience, and the actions it supports are losing their clarity. This has consequences not only for how users understand their interactions with digital content, but also for the sense of ownership felt in relation to it.

Harper *et al.* propose that the work of reimagining the file is interleaved with the work of producing a new ‘grammar of action’ for a networked world; the challenge they articulate is the need to redesign the file metaphor alongside the means with which users can act upon it. Drawing on work by Odom *et al.* [16, 18], they propose a shift towards an abstraction that encompasses metadata, for example bundling a Facebook photo with comments, likes and tags, and that can be acted upon in ways that support a sense of ownership. They highlight the need to rethink actions such as ‘copy’ and ‘delete’; the former serving to dilute control over one’s content, the latter lacking power when content resides on multiple services. Extending this argument, we note that newer actions, such as ‘share’, also have ramifications for ownership. Share can imply any one of a diverse range of operations and permissions depending on the technology being used, including allowing someone to view, comment on, edit, download, or even delete the content in question. While differences such as these can lead users to choose tools expressly because of the sharing and collaboration model they permit [26], technologies are also frequently selected for pragmatic reasons, such as their being used by a critical mass of collaborators [25]. In these cases, the model of sharing adopted is that enabled by the technology, along with all that it implies for ownership and control.

The question of how we might create a new file metaphor to enable actions that underpin a sense of ownership is highly challenging but also deeply inspiring, and was the stimulus for the work reported here. In engaging with this question, we undertook design work, prototyping and user research, to explore what new file metaphors might be, how they might be presented, what actions they might support, and what user experiences they might facilitate. To contextualise this work, we give a brief account of related research.

RELATED WORK

Bolstering a sense of ownership by supporting awareness and control of digital content is the motivation for the work presented in this paper and, while we have taken a particular approach to this challenge, complementary efforts have been adopted in prior research.

One body of work has investigated new ways of organising digital content within a system, such as through using metadata to move beyond the constraints of a file hierarchy. This can be seen in the seminal Placeless Documents project [4], as well as in activity-based computing systems such as Tagtivity [19] and Giornata [27]. Further efforts have been developed to support organisation of content across devices and services, such as in Perspective [21], a distributed storage system for the home. Perspective presents ‘views’ over content, expressed much like search queries, which show a specific set of data and the device on which it is stored. The system uses semantic naming for both data access and management, and has rule-based data placement (e.g., ‘back up all the music files on this device’), so that the user does not need to know where files are in terms of actual (and often hidden) file hierarchies.

An alternative approach to supporting users in keeping track of digital content focuses on provenance. Jensen *et al.* [11] report findings from a longitudinal study of knowledge workers, tracking provenance events in their computer use. They define provenance as the history and genealogy of a file, and note its use in aiding understanding of the evolution and relationships between files, how and when different versions have been created, and how they build upon each other through copy-paste events. They use provenance as metadata that might assist keyword search, but also observe that surfacing of provenance has a role to play in tracking propagation of mistakes and evaluating the reliability and veracity of data. Provenance graphs were shown to be effective memory cues, revealed differences between similarly named files, and helped users recognise patterns in their workflow. Other attempts at making visible the histories associated with digital content include depictions of version transitions [2], computational wear [10] and traces [20], while the notion of a virtual patina comprising social metadata, such as likes and comments, has been considered as a means of contributing to the value and uniqueness of virtual possessions [18].

Much of the above begins to suggest what it might mean to rethink the file metaphor, especially in terms of incorporating provenance. Karlson *et al.* [12] have been more explicit

in this, introducing the ‘versionset’ – the set of files that together represent a user’s concept of a document. They position this in the context of a copy-aware computing ecosystem, in which computers track and surface copy relationships between files. They report three main categories of copy creation based on in-depth interviews with information workers: copies for content preservation (through backups, emails to self, syncing); copies to share data across devices; and copies to share data with others. A prototype deployed to end-users demonstrated that the versionset can be inferred, and underpinned further findings on the nature of different copy events. For instance, ‘save as’ was often used to mark significant content changes, whereas ‘copy’ could be used to bridge file hierarchies. It was also found that not all copies are equal; for example, a PDF may represent the final version of a sequence of Word files. Of course, nuanced versioning is also enabled through source control systems such as Git [7], and more general efforts have been made to create unifying abstractions that encapsulate different representations of the same data. An example is the quFile [24], which returns the appropriate representation according to the context in which the file is read. In one implementation, the quFile underpins a copy-on-write file system, which enables previous versions to be retained and reverted to.

Complementary efforts to identify abstractions of the relationships between digital content come from the perspective of information studies. Feinberg [5] draws on the notion of ‘the intellectual work’ as a means of understanding how objects may exist as potentially vast sets of copies and almost-copies and, in doing so, identifies three levels of abstraction: document, text and work. A ‘document’ indicates a specific instance, e.g. a particular eBook, purchased by an individual. A ‘text’ indicates a set of symbols but not their physical embodiment, e.g. a particular edition of a book. ‘The work’ is the concept that links together all the documents and texts, e.g. ‘Hamlet’, including its different editions and translations. Feinberg argues that the work is best defined as a relationship between particulars, and is a relevant means of understanding document ecologies, especially for digital content, where ease of copying makes for extensive sets of similar files.

As noted in the Introduction, understanding the relationships between particular instances of digital content, and creating abstractions to represent them, can enable new possibilities for action. Examples include tools that permit the deletion of content regardless of how it has been synced and duplicated [6], entities that articulate a user’s data sharing preferences across a fragmented ecosystem of technologies [28], and systems that allow users to specify how services can access their data [1]. As noted in the Introduction, research into virtual possessions has also proposed the ‘bundling’ of content to form new abstractions [9], and the development of new actions as ways of underpinning ownership, including by reasserting ownership (by, e.g., withdrawing content from an online service) as well as by being able to relinquish it [16, 9]. In this paper, we draw inspiration from and extend this

work, through an exploration of what a file abstraction that enables a user experience of ownership might be, and what it may imply for an associated grammar of action.

CONCEPTUAL WORK

The overarching aim of this work was to explore and research new file metaphors and a new grammar of action for a networked world, with the goal of supporting a user experience of awareness and control. Our starting point was to engage in a process of brainstorming and scenario-development, to think about what an extended file abstraction could be and what actions we might develop alongside it. In focusing on file abstractions, we centred on single entities as well as bundles (such as Harper *et al.*'s [9] example of a Facebook photo with its social metadata, in turn inspired by Odom *et al.*'s virtual patina [18]), and included content that is user-generated, collected, bought, rented, streamed, or that accumulates through interactions with a system. Our thinking about action possibilities was organised around themes including actions for acquiring, keeping, ridding, organising, showing, sharing, finding, editing and collaborating on content. Following these initial attempts at scoping our thinking, we selected three examples of digital content to focus on in more depth. These were a Word document, as an exemplar of content that can be co-authored and published; a photo, as an exemplar of user-generated content that can be stored and posted in multiple sites and that acquires layers of metadata; and a playlist, as an exemplar of a bundle of content that takes its meaning from the pieces that comprise it.

We created typical scenarios for each of these, beginning at the point of their creation and encompassing editing, sharing or posting online, and removal from these spaces. As an example, in the scenario for the Word document, the document was initiated, had content pasted into it from elsewhere, was edited, and was then passed to a collaborator (a process that was repeated), before receiving comments from a supervisor and being published to the author's website and to an online library. As part of this scenario, we tried to capture some of the ad hoc ways in which users try to maintain control of their content (e.g., sharing a copy via email to keep one's own version independent, saving a Word file as a PDF before making it available online), as well as exploring opportunities for innovation (e.g., keeping track of what happens to a published version as it is downloaded, cited, and commented on).

Having developed these scenarios, we moved on to a process of sketching them (see Figure 1). This was primarily motivated by the need to move away from language in trying to unpack the atomic actions that made up our scenarios, and in doing so remove preconceptions associated with the current suite of actions we are so familiar with. Our aim was to produce a set of representations that combines current practice and the emergence of new actions (see Figure 2 for an example). One of the notable outcomes of this was that, although we had selected three types of content that are quite different, the resulting representations were visually very similar.

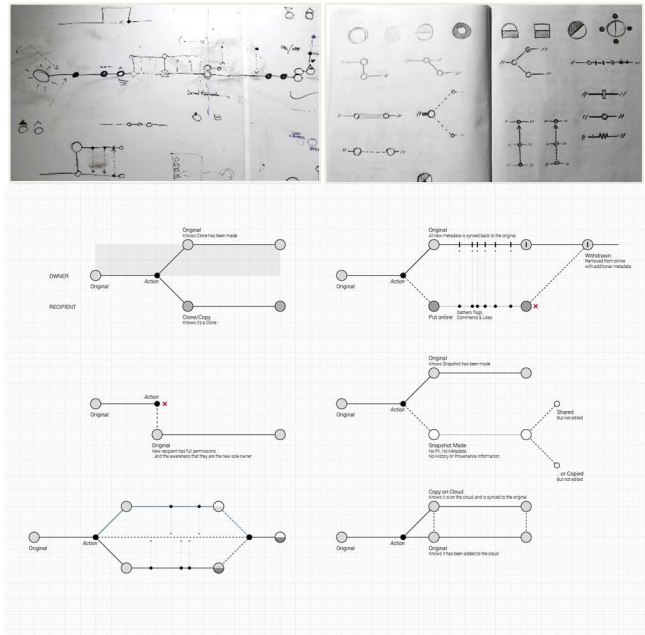


Figure 1. Initial sketches of the scenarios, including the development of visual representations of key actions.

The resulting representation, which we call the 'file biography', is an analytic tool conceptualised as an extended representation of the file, which can reveal information about its different versions (over time), its different instances (identical files stored in different locations), and what has happened to it (who has worked on it, commented on it, cited it, and how it has been transformed from one format to another). The file biography might be thought of as an abstraction linking multiple versions and instances such that, in a sense, it enables digital content to tell stories about itself.




Figure 2. A file biography of a shared Word document. Copies stored in different locations are represented vertically and time is represented horizontally. Different colours represent actions by different people on the shared document. Small circles represent edits; larger circles represent key versions (milestones) and important actions, such as sending copies to other people, merging edits, posting online and withdrawal from the web. Additional detail is provided in the video figure associated with this paper.

The brainstorming, scenario development and sketching that led to the file biography was also a means for us to explore a grammar of action for a networked world. We did this by documenting the basic actions associated with biographies for different content types, and by complementing these with a set of actions to be used in place of ad hoc ways of managing digital content. Our aim of supporting a user experience of ownership underpinned by awareness and control resulted in a focus on actions related to sharing and copying in particular. The grammar of action we produced is presented in Figure 3. Legacy actions on the left extend existing actions (‘transform’ encompasses actions such as ‘save as’), while the columns in the centre and on the right attempt to tease apart the circumstance where a user intends to create a distinct copy of a file, and the situation where that copy is the implicit result of some other act of sharing.

Legacy actions	Clarifying ‘share’	Relationships between versions and instances
New	Show	Milestone
Read	Entrust	Clone
Edit	License	
Delete	Give	Cite
Move		
Transform	Withdraw	

Figure 3. Initial grammar of action for a networked world.

The central column refers to ‘sharing’ as the management of access to a single instance of a file. We make a distinction between ‘showing’ and ‘entrusting’, the former encompassing mechanisms for sharing as read-only; the latter encompassing mechanisms for sharing that gives the recipient rights to edit and delete that content (such as is permitted in Dropbox and SharePoint). These actions are complemented by two new ones. ‘Licensing’ conveys an act of sharing whereby permissions are negotiated between the user and some other party, be that a person or a service (this is somewhat akin to the possibilities offered by Databox [1]). ‘Giving’ refers to the transfer of permissions from one person to another, and could underpin social practices such as gifting, loaning and inheritance. Notably, in a system where digital content remains under the user’s control, but is either shown, entrusted, or licensed to others, ‘withdrawal’ is also possible.

The column on the right refers to events where the action of ‘copy’ would typically be used to mark different versions of digital content, or to create different instances of it. The first action here is to ‘milestone’, which refers to marking versions in time. Versioning is already very common in some domains, such as software development, yet elsewhere users rely on ‘save as’ to manually create different versions. Milestoning encompasses this behaviour, and the resulting milestones support navigation via the file biography. The second action is ‘clone’, which we propose as an alternative to ‘copy’. While milestones deal with versions in a linear trajectory, clone entails a fork in the file biography. We propose that links between the resulting instances are maintained,

which implies a further iteration of possible actions, such as pushing and pulling changes. The final action here is ‘cite’. This is inspired by emerging practices that are made possible in a networked work, such as memeing and remixing [13], and allows the user to indicate a link between two pieces of content that would not otherwise be created.

DESIGN AND DEVELOPMENT

Following this, we designed three concepts to explore how the file biography could be used to portray content stored and copied to different places, and to support interactions with content that is versioned over time.

Atlas

Our first concept was initially formulated as a video of an Image Tracker (see Figure 4), which presented a unified view of a user’s photo collection (stored locally and online) and allowed the user to inspect the file biography of any image. This biography includes details of edits to the photo, the different sites it has been stored and copied to, and the comments and likes it has accrued. Based on this, we prototyped a system called Atlas, which imports images from local drives, OneDrive, Facebook, Flickr, Instagram and email accounts, identifies duplicates, and presents them as components of a single file biography (see Figure 5). The biography indicates when content has been stored in different locations, and documents the addition of likes, tags and comments. Additionally, Atlas enables the user to filter by storage location and photo album, and to pivot by the following: owner (the person who stored the photo in the given location), date the photo was taken, tags, camera model, people (those tagged in photos and the senders of emails), number of people, commenters, number of comments (on Facebook and Flickr), ‘likers’ and number of likes (on Facebook).




Figure 4. Still from the Image Tracker concept. The video is included in the video figure associated with this paper. Photo: Yulia Mayorava/Shutterstock.com.

Milestoner

While Atlas is primarily a way of keeping track of content stored and copied across multiple sites through a single representation, Milestoner deals with different versions of the same document over time. Milestoner supports the action of milestoning: marking a point in time to indicate that something noteworthy has happened. This is a simple action that

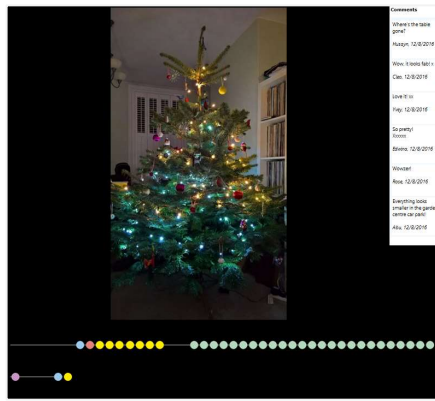


Figure 5. Simple file biography generated via Atlas, showing an image that was taken and stored on a mobile phone (lower line), and tagged, commented and liked on Facebook (upper line). Comments are also listed in the pane on the right.

is well understood by users, but one that can, when performed using ‘save as’, result in many redundant versions of files that are difficult to navigate retrospectively.

Milestoner was prototyped as an add-in for Word, which is accessed via the ‘Review’ ribbon and allows the user to create and name milestones (see Figure 6) and view the resulting file biography (shown in Figure 7). Milestones are presented on the biography as small circles labelled with a title, name and date. The file biography is colour-coded to indicate who created each milestone. Clicking on a milestone enables the user to open a previous milestone, compare the differences between it and the current document, and delete it.




Figure 6. Buttons added to the ‘Review’ ribbon in Word, and dialogue box that opens on clicking ‘Create Milestone’.

Activity Timeline

Our third concept, Activity Timeline, developed further some of the ideas we began to explore with Milestoner. Activity Timeline is designed for a collaborative writing context. Each document is associated with a simple file biography (shown above the text) as well as a chat stream (shown to the right), which can be used to create milestones (see Fig-




Figure 7. File biography shown at the bottom of a Word document showing colour-coded milestones created by three co-authors.

ure 8). When users navigate to earlier milestones via the biography, the chat moves back to provide context. Likewise, when clicking on an earlier milestone in the chat stream, the document reverts to a previous state and the biography also indicates that this is the case.

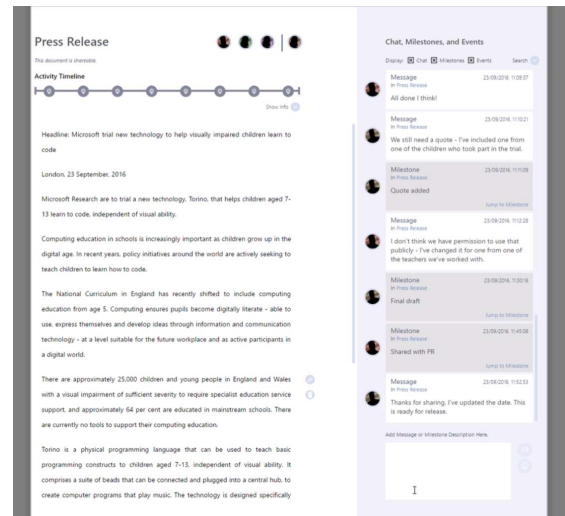


Figure 8. Activity Timeline showing milestones in a file biography (top) and chat stream (right). A video of the concept is included in the video figure associated with this paper.

USER RESEARCH

Alongside this strand of design and development work, we also undertook two user studies to inform and help us iterate the file biography concept. The first study was exploratory; we asked participants to sketch out file biographies about their own content. The second study was a deployment of Milestoner, which picked up on some of the issues raised in the first study, and allowed us to examine how file biographies are created over time and in a collaborative setting.

Study 1: Sketching the File Biography

The first study was an exploration of the file biography concept. We aimed to consider what should be shown and elided in the file biography, and whether this differs across file type; what might be visible in file biographies and how this differs by audience; and what user experiences a file biography might support, such as awareness and control and, relatedly, ownership. We aimed to understand how participants would sketch and talk about a variety of file types, including content with complex histories and with rich metadata, which is stored and synced across multiple locations, which may be created through the re-use or remixing of existing content, and which may involve collaboration or input from others.

Method

We asked participants to tell us about content they had produced (whether professionally, for education or as part of leisure) and to carry out two sketching activities. In the first, they sketched a biography-like representation relating to this content; in the second, they sketched a representation relating to a written scenario provided by us.

Participants

Ten people (three women and seven men) took part in the study, comprising three high school students (P8, P9, P10), three coders (P1, P3, P5), two makers designing content for laser cutting and 3D printing (P6, P7), and two gamers (P2, P4). Participants were aged from 17 to their mid-30s and were recruited through adverts placed on Gumtree, in a local maker space, and via an email sent to a local school. We recruited coders and makers because the content they produce is often modular, organised using version control systems, backed up and synced to various sites, and incorporates third party material. We recruited gamers to explore creative efforts such as the tailoring of characters and production of game mods. Finally, we recruited high school students because we expected digital natives to use cloud systems and online services, and to have a rich set of virtual possessions. Participants were given a £30 Amazon voucher.

Interviews

Interviews lasted from an hour to 90 minutes. We began by asking participants about an activity they had undertaken using digital technology; examples included producing a 3D model, creating a game as part of a small team, editing a photo, making a character for a videogame, and making a presentation with classmates. We then asked participants to sketch out a representation of what a file biography of their content would look like. They were shown the file biography in Figure 2 to give them a sense of how content history might be represented, but the meaning of the different elements was withheld. Participants were told they could draw something similar or something different, and include whatever they deemed relevant. Sketches produced by participants were used to ground further questions about what the representations might include and what values they might support. An example is given to the left of Figure 9.

Participants were then asked to sketch out the content history of a written scenario we had produced, which was about two writers collaboratively working on a piece of fan fiction. This was intended to introduce some consistency across the sketches, so we could look for points in common across participants. The scenarios were also useful for grounding questions about how a stranger’s digital content might be represented, and about why certain elements were omitted. An example is given to the right of Figure 9.

Analysis

Interviews were transcribed in full and analysed for emerging themes alongside the sketches, using methods described by Corbin and Strauss [3]. NVivo Pro was used for open coding, before the resulting ‘nodes’ were organised into axial

codes and then categories. Four categories were developed: process (example axial codes include phases, feedback, and sharing), collaboration (axial codes include awareness, communication and roles), organisation (axial codes include place, versions and file) and ownership (axial codes include control, credit and re-use).

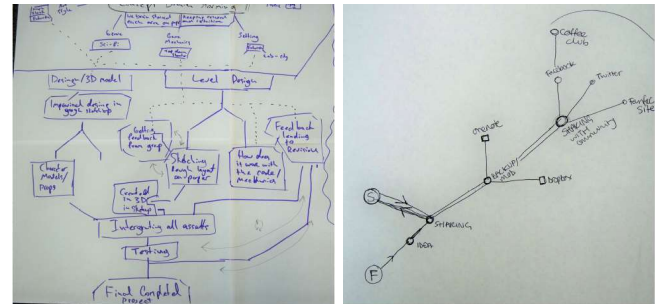


Figure 9. P3’s representation of the making of a videogame (left) and P4’s sketch of the fan fiction writing scenario (right).

Findings

Participants generated sketches about photo editing, generating shared presentations, coding projects, 3D model making, and their experiences with specific games. They were influenced in their sketches by their experience with content history tools (“*Hmm, yeah, it’s basically GitHub, isn’t it?*” – P1), but common points also emerged regarding what should be displayed and to whom, what would be useful to view in a representation of someone else’s content, and how the representations might be used.

In general, the sketches were conceptualised as timelines encompassing streams of activity:

“If there was a file format that would basically remember all the changes you’ve made, and so you could roll back and go up any branch, but it’s all inherently a single master file, and you could, you know you could explore the many variations you’ve looked at, I think that’d be really useful. So, you’d have one project file, and then it would have all these fantastic navigation tools that would allow you to navigate the history of that very same file” (P7).

While P7 uses the word ‘file’ here, a broader abstraction is indicated. Participants often included multiple concurrent tasks and their associated files in their sketches. These were related to the same goal but involved different activities and, in some cases, different individuals. In addition to person, place was a key element in these representations, encompassing local and online stores (the latter being personal or shared), as well as websites, networks, and devices. So, while file and folder structures played a role in content management, the underlying sites for storage were also an important component of organisation. As P10 noted:

“So, I think you’d say I have kind of like a local photo album on my laptop where I just have everything, then I have a few iCloud albums which are all albums that I’ve shared with like myself and my girlfriend or myself and my sister, where we are both putting photos into it, and that’s usually for a

particular thing like a holiday [...] Whereas Google [Drive] tends to be once I've made it all and put it all together and chosen everything, [...] that's kind of the final cut as it were."

As activities were sketched across multiple places, the resulting representations were understood as capturing history by serving as a record of which files had been stored where. Included in this were the origins of files, key versions, changes, and references to related content. Capturing this was understood as potentially useful for record keeping, by helping track content and navigate to different versions, especially when those are distributed across multiple places. P1 noted, *"It can get so messy if you don't keep track of it: 'Wait, what did I plan to do with this version? And what did I add to it?' You forget that you add changes in different places, that it's not the main version."*

Representations were also sketched as having multiple phases, including initial activities such as brainstorming, sketching and producing notes that would often go undocumented or captured on paper that is then discarded. Participants noted that it would need to be extremely easy to include paper-based content in a digital representation if this were to be realistically accomplished. Activities also often began with the repurposing of existing content, and again, this was not captured in any meaningful way in participants' current practices. These initial phases of ideating and remixing were often followed by productive work and making, including processes of refinement and iteration. Within this, content might be discarded or go unused in response to feedback or shifts in direction. Version control was seen as especially important in complex collaborations or activities, although not all participants kept content histories: *"I do tend to delete all the old versions once it has been submitted, or once it's finished then I'll just keep the finished version usually as a PDF, and delete all the old junk"* (P8).

In addition to keeping track of their content, participants also suggested additional roles that the representations they produced might play. One of these was in managing collaboration. For participants, collaborators ranged from skilled co-authors and contributors, to commentators, audience members and strangers, who may nevertheless be providers of feedback and other forms of input. Therefore, roles that might be expressed in a representation included close collaborators but also people with more distant relationships to the participants or activity. In terms of close collaboration, participants expressed that by making activity visible, the representations could facilitate awareness and so, support understanding and communication. This might partly be done by highlighting important moments. P1 noted that *"I kind of like the idea of being able to mark yourself which moments are important in the biography"*, and P8 described how he captured key communications through a specific action: *"If there are particularly important messages with important details on them or something then I'll screenshot it on my phone and it will save to the camera roll"*. The representa-

tions were also seen as potentially providing a means of documenting communication in both digital and non-digital channels, as P6 noted: *"I guess if it was... you know someone had a really great modification that they suggested, I probably would go and write it down, or sketch it out. Yeah, I probably would record it in some way."*

Other values the representations might support include indicating progress through an activity and enabling reflection on the approach taken, as noted by P5: *"So personally I think it would provide valuable retrospective input. So, for example, just drawing the timeline as you asked me to really made me sort of realise that there is a lot of time I wasted on silly things. I mean I guess it's quite specific, but I guess in terms of my developmental process, if every time I finish a project I can just say 'generate timeline' and then view the timeline that'd be super helpful"*. Participants spoke of zooming in for greater detail, but also felt that the level and quality of detail might differ by content type and the viewer's relationship to the content. Regarding the latter, it was expected that producers of content would have a different view to other people, including colleagues, recipients and wider audiences. However, participants did not see the need for representations to capture every detail. For example, P2 described how *"the stuff that I didn't write down is probably the names of the different random people"*, meaning the names of people posting via social media in the fan fiction scenario (although the expertise of that person was deemed worth noting). The finer details of how someone else has re-used one's work were also seen as surplus to requirements, as P5 explained:

"If it's someone you're collaborating with, you might want to have more information on what they are doing. You really want to know exactly what they've been working on so that you don't sit down for three hours and do the same thing they did the day before. I think it's also useful to have a good idea of where your collaborative partner has researched things. If, for example someone makes a pull request on your GitHub page, and they take something and modify it, I guess you don't really need to see all the process, maybe just some highlight of what they've done. [...] I mean you don't really care, it's not your project. I mean it is to an extent, but it's their modification. Yeah, I think it does make a difference."

Representations were seen as supporting a different set of possibilities when depicting content produced by others. Examples include support for learning by giving an overall view of how a task has been accomplished by others:

"In fact, seeing how other people create things as well would be interesting and useful because it would help me to be more efficient and productive to take influence from them. [...] Especially music, because the kind of music I make is mainly instrumental, and it's kind of long pieces that are repeating. I think it would be interesting to see how other people who create similar music are able to link everything in a non-linear fashion." (P3)

When reusing content produced by others, either as a template, by repurposing it for a new activity, or by directly using open source or third-party materials, participants felt that rich detail could aid trust in the information, and that this might include manually entered details for more complex activities. As P6 mentioned:

“If you had a bit more details, or history for a stranger, you might almost kind of trust their designs a bit more? You know, you usually see their finished projects as a picture or whatever, I’m not quite sure if it’s actually as easy as they said it is, or if it actually works quite as well as they said. If it just kind of looks right, but actually there are kind of issues with it. Yeah, maybe that would give you a bit more confidence that they’ve kind of tried it out properly... to see some of their previous mistakes and then seeing them being fixed.”

However, while participants could see the value in viewing other people’s mistakes, they were not always keen to share their own, as P6 acknowledged: *“I probably wouldn’t share with the world my initial slightly rubbish versions”*. Keeping content private was also seen as useful to avoid overwhelming any viewer of the work, and, in some cases, to avoid compromising ownership. Ownership over digital content was bound up with originating it, having a role in its production, and feeling that one should be credited for this role. It was additionally related to control, and this includes having access to content, knowing what has happened to it, being able to act upon it, and specifying how others can act in relation to it, for example, by expressing copyright. P10 commented:

“Then the other thing is just who has that piece of information, because obviously with a photo it’s not necessarily as important, but as far as I’m concerned with any data or any information it’s important to know who’s got it at any point in time, especially in kind of a modern developing world where there’s a lot of kind of cybersecurity issues. Being able to know exactly who has it and where it is at any one point in time is quite important.”

Summary

These findings indicate that the file biography might be expanded beyond the notion of a single entity, versioned over time or duplicated across different locations, to encompass multiple entities that have an activity in common. These entities might be associated with different places but also with different people. The level of detail included with regard to individuals especially is highly contextual; it may be more useful to include detail pertaining to collaborators, but having sufficient information about strangers also helps interpretation of their contribution, for example by being aware of their role or level of expertise.

The study also suggests an expanded set of user experiences that might be underpinned by the file biography. The representations that participants sketched were not just about awareness, control, or navigation, but also touched on reflection, exploration, learning, and support for re-use. This sur-

faceted tensions in what should be made visible in a file biography; while participants thought it useful to learn from the mistakes of others, they were not so keen to share their own.

Study 2: Co-authoring with a File Biography

Our second study was a deployment of Milestoner, and picks up on the themes of collaboration and making work visible that emerged in the first study. This study was an opportunity to focus on how participants would work with a file biography over time and in the context of co-authoring. We explored the different possibilities offered by the tool (marking and displaying progress as well as navigating to prior versions) and the user experiences that these underpinned.

Method

We deployed Milestoner within our research group over a period of six months. (Unfortunately, due to reasons of confidentiality, we were not able to deploy Milestoner outside our lab at the time of the study.) Six of our colleagues (three women and three men) used the add-in to produce five papers as part of their ongoing work. Three of the papers were written by multiple people using Milestoner (P1, P2, P3 and P5); two were sole-authored (P4 and P6). Co-authors stored their documents on OneDrive for Business so that they could edit the same document and see the same file biography.

We conducted 20 interviews, interviewing each participant at least twice and up to four times, approximately each month over periods ranging from two to six months (times varied with the length of the particular writing process). Retrospective interviews were used because Milestoner tends to be used in very short bursts but over a very long time period. The aim of the first interview was to understand the nature of the writing that was planned, what work had already occurred, whether the collaboration was well-established or new, and what practices were in place to support version control, backup, use of shared repositories, commenting and group relations. Follow-up interviews examined how milestones had been created, by whom and for what purpose, and how writing was progressing and being managed. As time passed, we asked whether the relevance of milestones had changed in hindsight, and whether the file biography had supported re-visitation of earlier versions.

Analysis

Interviews were conducted face to face and, in one case, via Skype. They were audio-recorded and transcribed. The data were analysed with a focus on why participants created milestones, why they revisited documents associated with milestones, and what value they gained from the file biography.

Findings

Milestones were principally created to mark progress and, for co-authors, manage collaboration. Milestones created to mark progress were added at certain time intervals, such as at the end of each day (P3), or to correspond with a completion point, such as the end of a section. This provided reassurance that changes could be easily undone, both in sole authorship, when *“making changes in the document that might be not easily recoverable”* (P4), and when co-authoring, in

case co-authors made unwanted changes. Additionally, milestones were labelled to indicate the presence of sections that the author expected to be reduced or even removed, but that might later be re-used in other documents. Adding milestones on completion of key sections meant that authors would easily be able to re-find and re-use this material:

“We are thinking of taking quite a lot of the material I’m cutting out and putting it in another paper and at first, I was sticking it in the OneNote saying, ‘Here are extra bits’, but actually now I’ve just said ‘Okay, I know that in Version 2 that old section is still there, and I can just take it’.” (P2)

For the management of collaboration, milestones were created to support turn-taking. For instance, in one collaboration, a process was agreed in which milestones were inserted whenever someone started working on the document, to *“save the last person’s work”* (P3). This made it obvious who was working on the document at any one time, and meant that lost work could easily be recovered: *“There were things like that when someone said, ‘Well I wrote this’, and I went ‘Did you?’ and I had to go back and find it”* (P2). Additionally, the descriptions associated with milestones were noted as providing a *“good way to get an overview of what the changes are”*. Milestones were highlighted as a *“way of talking for the document as a whole [...] to add a meta-layer of annotation”* (P1). This was also useful when working on a paper that had been started by a co-author:

“I think the milestone notes have helped with seeing, ‘Okay, this is the order in which she actually put these things down. These big chunks of text, this is how she wrote them.’ If I read the document as she has meant to portray them to the reader, from top to bottom, as with every first draft, there are some strange inconsistencies and skips in the narration, but they make much more sense when you see the order in which she actually wrote the paper and how her thoughts developed as she wrote the paper. I think that’s one thing that I’ve definitely gotten from Milestoner.” (P1)

Understanding the evolution of a document was also personally useful. For example, reviewing the process later could lead to insights such as *“how long it might take next time when you are doing similar work”* (P6).

Participants also highlighted aspects of Milestoner that could be extended. These include adding support for associating additional content with milestones, and making some milestones private. Participants highlighted numerous reasons to include additional (and sometimes separate) documents in a file biography, including adding documents to the beginning or end. For example, P6 wished to affix his MSc thesis to the paper he was writing about the work done towards it, whilst making it clear that the two documents are distinct (and therefore not part of the same biography). Other examples include adding literature reviews to the beginning of a biography and published documents to the end, as well as connecting paragraphs sent by email, to make activity outside of the document visible to co-authors. Finally, participants

wished to include cleaned-up versions of the document when showing it to others, suggesting a need for branching and potentially merging.

The desire to withhold some content was also expressed by participants who saw value in making some milestones private. *“Selectively exposing milestones”* (P1) could allow authors to hide milestones they didn’t expect to be relevant to their co-authors, or that they were not yet ready to share. Relatedly, milestones could be linked to access control, such that the work being done would lock the document.

Summary

These findings suggest implications for the file biography and grammar of action, as well as some next steps for Milestoner. As a first step, a representation closer to that shown in Figure 2 would address some of the issues raised in relation to branching and merging. However, iterating the representation further would allow it to better handle occasions where, for example, a cleaned-up instance of a document is created for sharing, but the feedback received is applied to another ‘working’ instance.

Our findings also imply that we might consider how to tailor file biographies for different audiences and different purposes. Fading out the finer details to support a general overview, enabling reflection on process, and supporting the creation of both visible and undisclosed milestones, could underpin richer and more flexible user experiences.

While we had envisaged the file biography as relating to a single file, our findings again highlight the value of linking to related content. The file biography might automatically indicate circumstances in which text is copied and pasted from one document to another; however, emails that relate to the writing process might need to be specifically connected by users. This suggests the need for further research into the action of ‘citing’ (Figure 3) as a means of linking content.

DISCUSSION

The aim of this research is to explore new file metaphors and an associated grammar of action, with the aim of underpinning a user experience of awareness and control. As outlined in the Introduction, we have been inspired by the challenge put forward by Harper *et al.* [9] to consider these as intertwined: a new metaphor should enable certain types of action. In beginning to explore this space, we have focused first on developing a metaphor and grammar of action that is user-centred. We have developed a metaphor, the file biography, through conceptual, design and development work, and have explored it through interviews with users and through a deployment of Milestoner. In this Discussion, we draw on this body of work to consider what we have learnt about new file metaphors and what this implies for associated actions.

The first issue we address here is what a new abstraction for the file should comprise. Harper *et al.* posit that a new abstraction should bundle digital content in ways that are more expansive than that represented in the existing file metaphor.

As noted, they give the example of a photo posted to Facebook, which encompasses comments, tags and likes when downloaded. We focused on place and time as key aspects of our initial metaphor. Place was explored through Atlas, by mapping the different locations photos are stored in and posted to, and time was explored through Milestoner, by enabling writers to build a file biography over versions, and through Activity Timeline, by allowing users to recreate the context around a milestone by surfacing related chat.

However, when participants sketched their own representations and used Milestoner as part of their writing, we found that they wished to extend the scope of what might be included. In their wish to incorporate communications, documents that feed into the writing process, and even checklists, we see an abstraction that goes beyond the file, the versionset [12], the quFile [24] and ‘the work’ [5], all of which relate to documents and texts that are understood as related and being, in some sense, ‘the same’. Instead, by wishing to encompass related but conceptually different content (such as the thesis that feeds into a paper but is understood as a distinct piece of work), participants expressed ideas that resonate with activity-based computing [c.f. 19, 27]. Furthermore, by presenting a visualisation of activity, the file biography becomes a tool not just for navigation or keeping track of content, but for reflection, learning, and support for collaboration, a set of opportunities that emerged in both studies.

This emphasis on activity suggests a reframing of the file biography, from a means of allowing files to tell stories about themselves, to a tool that sits between individuals, and is an enabler for the activities that span them. One challenge raised by this perspective relates to how new file abstractions should be presented to different audiences. What should be included and omitted from content histories? And how would this tie into what actions are permitted? Possibilities that arose from the Milestoner deployment include management of turn-taking in co-writing by ‘locking’ milestones or otherwise hiding work that the author is not yet ready to share. These suggestions resonate with more complex version control systems, such as Git [7], in which work can be checked into and out of branches that are shared with others, as well as with Cimetric, a file synchronisation platform for scholarly collaboration [14]. Cimetric enables user control of change integration, while provenance information enables files to be checked in and out of different repositories.

An emphasis on activity also highlights how the grammar of action we initially proposed might be rethought. Our aim of supporting a user experience of ownership and the resulting emphasis on the need to clarify sharing and copying, which could then better support actions such as giving and withdrawal, led us to focus on permissions. In teasing apart different modes of sharing, we considered the ways in which content is made available to others through a user-centred lens. While it would be possible to apply such actions in productive work (and indeed, our Word document scenario focused on precisely this), the cycle of entrusting content to a

collaborator, and then withdrawing it from them when one wishes to work privately, would be laborious. Drawing inspiration from Git and Cimetric, we might consider alternative ways of checking in content in a revised grammar of action. This could be approached by breaking down ‘sync’, which, like ‘share’, is an action that carries ambiguities and is poorly understood by users [15].

These reflections highlight the challenges of creating new file metaphors for a networked world. On the one hand these must cater for artefacts such as photos, music and books, which we might consider as largely complete or finished, and which have been exemplars in much of the research on virtual possessions [e.g. 16, 17, 18, 8]. On the other hand, new metaphors must cater for artefacts such as papers and source code; works that are still in production. So, while the model of sharing that was brought to the fore in our initial grammar of action might have especial relevance to finished artefacts, enabling actions such as giving (transferring ownership), or entrusting (and so enabling an element of joint ownership), a richer and more accessible model of distribution is essential to developing a grammar of action for the production of content in a networked world. A grammar of action that works across ‘finished’ and ‘in production’ content seems essential, if we are to cater for instances where the former becomes material for the latter, in cases of remixing and re-use.

Our next steps in this project are to explore these questions. We have interviewed people who work with others to produce content, be that source code, documents or slide decks, to deepen our understanding of the actions that are interwoven with shared work. We are developing diagrams of the workflows that are described, and using these to push further our thinking of what new file abstractions might be.

CONCLUSION

The question of how to design new file metaphors for a networked world is challenging to address. In this paper, we have outlined our approach, which combines research through design, development, and with users, to explore this problem space. Our work centres on the file biography, used as an analytic tool to ground prototyping and interviews with users. Our research has led us to reflect on how new file metaphors can work across components that are stored in different places or that represent different points in time. It has also pushed us to consider how metaphors might play different roles in different contexts, with a distinction being drawn between content that is ‘in production’ and virtual possessions that are, in some sense, a finished form. We argue that these different lenses require different emphases when developing a new grammar of action; while the former might be bound up with ways of syncing content that enables work to be done in private, the latter might be bound up with more nuanced models of sharing and ownership.

ACKNOWLEDGMENTS

With thanks for thoughtful input to Richard Harper, Richard Banks, Phil Gosset, Tim Regan and Jade Alglave.

REFERENCES

1. Amir Chaudhry, Jon Crowcroft, Heidi Howard, Anil Madhavapeddysnip, Richard Mortier, Hamed Haddadi, and Derek McAuley. 2015. Personal data: thinking inside the box. In *Proceedings of The Fifth Decennial Aarhus Conference on Critical Alternatives (AA '15)*, 29-32. <http://dx.doi.org/10.7146/aahcc.v1i1.21312>
2. Fanny Chevalier, Pierre Dragicevic, Anastasia Bezerianos, and Jean-Daniel Fekete. 2010. Using text animated transitions to support navigation in document histories. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*, 683-692. <https://doi.org/10.1145/1753326.1753427>
3. Juliet Corbin and Anselm Strauss. 2008. *Basics of Qualitative Research* (3rd ed.). Sage Publications.
4. Paul Dourish, W. Keith Edwards, Anthony LaMarca, and Michael Salisbury. 1999. Presto: an experimental architecture for fluid interactive document spaces. *ACM Trans. Comput.-Hum. Interact.* 6, 2 (June 1999), 133-161. <http://dx.doi.org/10.1145/319091.319099>
5. Melanie Feinberg. 2013. Beyond digital and physical objects: the intellectual work as a concept of interest for HCI. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, 3317-3326. <https://doi.org/10.1145/2470654.2466453>
6. Roxana Geambasu, Tadayoshi Kohno, Amit A. Levy, and Henry M. Levy. 2009. Vanish: increasing data privacy with self-destructing data. In *Proceedings of the 18th conference on USENIX Security Symposium (SSYM'09)*, 299-316.
7. Git. Retrieved September 18, 2017, from <https://git-scm.com/>
8. Jane Gruning and Siân Lindley. 2016. Things we own together: sharing possessions at home. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*, 1176-1186. <https://doi.org/10.1145/2858036.2858154>
9. Richard Harper, Siân Lindley, Eno Thereska, Richard Banks, Philip Gosset, Gavin Smyth, William Odom, and Eryn Whitworth. 2013. What is a file? In *Proceedings of the 2013 conference on Computer Supported Cooperative Work (CSCW '13)*, 1125-1136. <https://doi.org/10.1145/2441776.2441903>
10. William C. Hill, James D. Hollan, Dave Wroblewski, and Tim McCandless. 1992. Edit wear and read wear. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '92)*, 3-9. <http://dx.doi.org/10.1145/142750.142751>
11. Carlos Jensen, Heather Lonsdale, Eleanor Wynn, Jill Cao, Michael Slater, and Thomas G. Dietterich. 2010. The life and times of files and information: a study of desktop provenance. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*, 767-776. <https://doi.org/10.1145/1753326.1753439>
12. Amy K. Karlson, Greg Smith, and Bongshin Lee. 2011. Which version is this?: improving the desktop experience within a copy-aware computing ecosystem. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*, 2669-2678. <https://doi.org/10.1145/1978942.1979334>
13. Lawrence Lessig. 2008. *Remix*. Penguin Books.
14. Catherine C. Marshall, Ted Wobber, Venugopalan Ramasubramanian, and Douglas B. Terry. 2012. Supporting research collaboration through bi-level file synchronization. In *Proceedings of the 17th ACM international conference on Supporting Group Work (GROUP '12)*, 165-174. <http://dx.doi.org/10.1145/2389176.2389202>
15. Cathy Marshall and John C. Tang. 2012. That syncing feeling: early user experiences with the cloud. In *Proceedings of the Designing Interactive Systems Conference (DIS '12)*, 544-553. <https://doi.org/10.1145/2317956.2318038>
16. William Odom, Abi Sellen, Richard Harper, and Eno Thereska. 2012. Lost in translation: understanding the possession of digital things in the cloud. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*, 781-790. <http://dx.doi.org/10.1145/2207676.2207789>
17. William Odom, John Zimmerman, and Jodi Forlizzi. 2014. Placelessness, spacelessness, and formlessness: experiential qualities of virtual possessions. In *Proceedings of the 2014 conference on Designing Interactive Systems (DIS '14)*, 985-994. <https://doi.org/10.1145/2598510.2598577>
18. William Odom, John Zimmerman, and Jodi Forlizzi. 2011. Teenagers and their virtual possessions: design opportunities and issues. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*, 1491-1500. <https://doi.org/10.1145/1978942.1979161>
19. Gerard Oleksik, Max L. Wilson, Craig Tashman, Eduarda Mendes Rodrigues, Gabriella Kazai, Gavin Smyth, Natasa Milic-Frayling, and Rachel Jones. 2009. Lightweight tagging expands information and activity management practices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*, 279-288. <https://doi.org/10.1145/1518701.1518746>
20. Daniela K. Rosner, Miwa Ikemiya, Diana Kim, and Kristin Koch. 2013. Designing with traces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, 1649-1658. <https://doi.org/10.1145/2470654.2466218>

21. Brandon Salmon, Steven W. Schlosser, Lorrie Faith Cranor, and Gregory R. Ganger. 2009. Perspective: semantic data management for the home. In *Proceedings of the 7th conference on File and Storage Technologies (FAST '09)*, 167-182.
22. David Canfield Smith, Charles Irby, Ralph Kimball, and Eric Harslem. 1982. The Star user interface: an overview. In *Proceedings of the National Computer Conference (AFIPS '82)*, 515-528.
<http://dx.doi.org/10.1145/1500774.1500840>
23. Susan Leigh Star. 1989. The structure of ill-structured solutions: boundary objects and heterogeneous distributed problem solving. In *Distributed Artificial Intelligence (Vol. 2)*, Les Gasser and Michael N. Huhns (eds.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 37-54.
24. Kaushik Veeraraghavan, Jason Flinn, Edmund B. Nightingale, and Brian Noble. 2010. quFiles: the right file at the right time. *Trans. Storage* 6, 3 (September 2010), Article 12.
<http://dx.doi.org/10.1145/1837915.1837920>
25. Amy Volda, Judith S. Olson, and Gary M. Olson. 2013. Turbulence in the clouds: challenges of cloud-based information work. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*, 2273-2282.
<https://doi.org/10.1145/2470654.2481313>
26. Stephen Volda, W. Keith Edwards, Mark W. Newman, Rebecca E. Grinter, and Nicolas Ducheneaut. 2006. Share and share alike: exploring the user interface affordances of file sharing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*, 221-230.
<http://dx.doi.org/10.1145/1124772.1124806>
27. Stephen Volda and Elizabeth D. Mynatt. 2009. It feels better than filing: everyday work experiences in an activity-based computing system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*, 259-268.
<https://doi.org/10.1145/1518701.1518744>
28. Jason Wiese, Sauvik Das, Jason I. Hong, and John Zimmerman. 2017. Evolving the ecosystem of personal behavioral data. In *Hum-Comput. Interact.* 32, 5-6, 447-510.
<https://doi.org/10.1080/07370024.2017.1295857>