

InterTwinkles: Online Tools for Non-Hierarchical, Consensus-Oriented Decision Making

by
Charles DeTar

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in
Media Arts and Sciences at the
Massachusetts Institute of Technology

September, 2013

©Massachusetts Institute of Technology 2013. All rights reserved.

Author
Program in Media Arts and Sciences
August 23, 2013

Certified by
Christopher Schmandt
Principal Research Scientist
MIT Media Laboratory

Accepted by
Pattie Maes
Associate Academic Head
Program in Media Arts and Sciences

InterTwinkles: Online tools for non-hierarchical, consensus-oriented decision making

by

Charles DeTar

Submitted to the Department of Media Arts and Sciences, School of Architecture and Planning,
on August 23, 2013, in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Media Arts and Sciences at the Massachusetts Institute of Technology

Abstract

Non-hierarchical, participatory, consensus-based decision making has seen an explosion in popularity in recent years. The traditional techniques of formal consensus, however, are limited to face-to-face meetings, which can limit organizations' capacity due to their time and cost. InterTwinkles is a set of integrated but composable online tools designed to assist small and medium-sized groups in engaging in formal group decision making processes online. In this thesis, I present a thorough investigation of the ethical and practical motivations for consensus decision making, and relate these to concerns of control and autonomy in the design of online systems. I describe the participatory and iterative design process for building an online platform for consensus, with particular attention to the practical constraints of real-world groups with mixed technical aptitude. I present the results of a three month field trial with six cooperative groups in the Boston area, and evaluate the results through the lens of adaptive structuration theory, with particular attention on the fit between the ethical motivations and performance outcomes.

Thesis Supervisor: Christopher Schmandt

Title: Principal Research Scientist, MIT Media Laboratory

Thesis Supervisor
Christopher Schmandt
Principal Research Scientist
MIT Media Laboratory

Thesis Reader
Sasha Costanza-Chock
Assistant Professor of Civic Media
Comparative Media Studies, MIT

Thesis Reader
Sepandar Kamvar
LG Career Development Professor of Media Arts and Sciences
MIT Media Laboratory

Contents

1	<i>Non-hierarchicalism</i>	17
2	<i>Consensus</i>	43
3	<i>Decision Making with Computers</i>	65
4	<i>Design of InterTwinkles</i>	89
5	<i>Field Trial and Evaluation</i>	115
6	<i>Conclusion</i>	139
7	<i>Acknowledgments</i>	143
A	<i>Consensus flowcharts</i>	145
B	<i>Meeting practices used by consensus-oriented groups</i>	153
	<i>Bibliography</i>	165

Introduction

Motivations

The Water Heater

For my first five years in Boston, I lived in a housing co-op in Dorchester. The house is a classic Boston triple-decker, which the 13 residents own collectively. That means that we didn't have a landlord, but were in charge of the mortgage, maintenance, utilities, and everything else. The house operated by consensus, and all decisions happened at weekly meetings.

About two years in, one of the water heaters that serviced the house died, which left our main kitchen without any hot water. One of the residents took on the task of fixing this, and called up a plumber for an emergency job. The plumber charged us the emergency rate; almost 2 thousand dollars to install a low-end water heater – a model which lists for around \$300.

In retrospect, we needn't have paid the emergency rate – we could have made do with only 2/3 of our water-heating capacity for a week or two, perhaps carrying dirty dishes to an upstairs sink to wash them. We could have gotten a much more energy efficient model which could have kept long term costs down and reduced our carbon footprint. We could have done all of this for less than half the cost. But the consensus process that could've allowed the house members to express those opinions was too slow, and we had no effective process in place to make decisions between meetings.

Could there be a way for a co-op to enjoy the benefits of creativity in problem solving conferred by its consensus oriented structure, but with the ability to respond to urgent needs?

Geographically dispersed board

For the last two years, I've served on the Board of Directors of a bi-national organization that supports housing cooperatives across the US and Canada. This professional non-profit draws its board



Figure 1: A water heater.

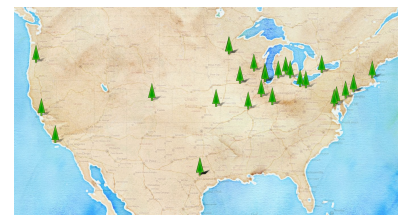


Figure 2: Geographically dispersed organization.

members from 4 time zones, and with a miniscule budget, runs a conference, a land trust, and a training and education network for tens of thousands of people.

The board operates by consensus, and holds only a few in-person meetings per year, due to the cost of travel. When the organization faces obstacles such as staff turnover or crises in its member co-ops, it can be difficult to sustain the creative energy and output necessary to sustain programming, leading to overworked staff and burnt out board members.

Could online tools allow a geographically dispersed board of directors to accomplish more, without having to meet more often?

Anti-war protests, 2002-2003

In 2002 and 2003, protests against the Iraq war erupted all around the world – some of the largest protests ever. Over a million people turned out in Rome, half a million each in London and New York, over 200,000 in San Francisco. For several months around the start of combat in Iraq, there were weekly mega-marches, with numerous smaller scale direct actions. The core of these protests were organized by clusters of small-scale affinity groups. Each affinity group, consisting of 10 to 20 people, decided what sort of role they were comfortable fulfilling, and then coordinated with other affinity groups to participate in larger actions.

These protests represented a tremendous mobilization of resources, and an inspiring moment of coordination. But as time dragged on, the energy required to maintain these affinity groups and clusters waned, and the protest movement faded. We saw history repeat itself in 2012, as the loss of centralized encampments in the Occupy movement impelled a loss of coordination energy.

Could a more robust, lighter weight coordination tool allow activists to build power over longer periods of time?

Designing tools for consensus

This dissertation represents the culmination of three years of work into designing online tools to support democratic decision making processes in non-hierarchical groups. Central to this work is the question: *How can we design effective online tools for democratic consultation?* Democracy as typically understood in the US consists of voting for a leader once every few years. Few people would accept a king or dictator as a satisfactory model for government, but nevertheless, many of us accept hierarchical control over our daily lives in the form of bosses, landlords, extortionate debt contracts, and other forms



Figure 3: San Francisco, Feb 16, 2003.
Photo: Chris Carlsson

of every day oppression. Instead of accepting this, many activists in the last 50 years have worked to develop effective organizational forms that bring democracy into every day life: worker's collectives, housing cooperatives, affinity groups, and more. Consensus decision making encompasses a set of techniques that these groups have innovated to make their daily work just, effective, and efficient. My goal is to extend these techniques online, opening new possibilities for liberatory organizing and extending organizations' capacities.

To contextualize this work, I will start with analyses of essential background that forms the basis of design. In chapter 1, I will discuss and problematize the notions of non-hierarchicalism and decentralization, and develop an analytical framework to identify mechanisms of control in both software and organizational systems. These mechanisms are key to the whether a system succeeds or fails as a tool for liberation. In chapter 2, I will develop a more refined analysis of consensus decision making practices, and distinguish the different ways the term and practices referenced by it are used by contemporary groups. In chapter 3, I will discuss the history of computer systems designed to support decision making processes, as well as some of the foundations for research involving groups.

In chapter 4, I will discuss the participatory and iterative design process that led to the current form of InterTwinkles. This will include a description of the design principles developed in partnership with target users, the functions of each of the tools, and the structure and capabilities of the platform. In chapter 5, I will discuss a three month field trial completed with six cooperative groups in the Boston area. I will analyze the results of this field trial in four ways: by comparing outcomes to the design principles developed during the participatory design process, by analyzing group outcomes using the lens of Adaptive Structuration Theory, by describing qualitative results based on solicited and unsolicited feedback from users, and by analyzing the system using the analytical framework for control developed in chapter 1.

Throughout this work, I have sought to maintain a perspective suitable for informing future designers of systems. The analytical work describing non-hierarchicalism and consensus are intended to map the territory of these large and fuzzy notions such that designers can more easily develop tools suitable for particular target groups and meanings; the historical analysis of decision support systems is intended to identify broad failures or contributions pertinent to the development of new systems; and the documentation of the design and field trial of InterTwinkles strives to inform future designers about what worked and what didn't with the approach taken here.

In my experience working with consensus-oriented groups, I've

observed a tendency of participants to throw their hands up in futility when faced with the prospects of doing group work online. My hope is that this work can inspire more development to bring the same creative enthusiasm that groups have for in-person work into online communication.

1 *Non-hierarchicalism*

So Yertle the Turtle King, lifted his hand
And Yertle, the Turtle King, gave a command.
He ordered nine turtles to swim to his stone
And, using these turtles, he built a new throne.
He made each turtle stand on another one's back
And he piled them all up in a nine-turtle stack.
And then Yertle climbed up. He sat down on the pile.
What a wonderful view! He could see 'most a mile!

"All mine!" Yertle cried. "Oh, the things I now rule!
I'm the king of a cow! And I'm the king of a mule!
I'm the king of a house! And, what's more, beyond that
I'm the king of a blueberry bush and a cat!
I'm Yertle the Turtle! Oh, marvelous me!
For I am the ruler of all that I see!"¹

Beginning in the 1960's, in an effort to reconstitute inter-personal politics in more egalitarian ways, activist groups in the US began experimenting with non-hierarchical, democratic forms of organization and decision making.² From the civil rights movement in the 1960's, to the feminists and anti-nuclear movements in the 1970's, the global justice movement in the 1990's and 2000's, and all the way to Occupy Wall Street in 2011, so-called "leaderless" and horizontal organizational forms have become a fixture of contemporary activism.

But activists aren't the only ones excited by the possibility of flat, democratic forms of organization. "Decentralization" has found massive success in peer production projects like Wikipedia, free and open source software, social aggregators like twitter and reddit, and even the fundamental architecture of the Internet itself. This has led technology prophesiers to unleash a torrent of books touting the benefits of non-hierarchicalism – including Clay Shirky ("Here Comes Everybody"), James Surowiecki ("The Wisdom of Crowds"), Steven Johnson ("Future Perfect: The Case for Progress in a Networked Age"), Jeff Howe ("Crowdsourcing: Why the Power of the Crowd is Driving the Future of Business"), Don Tapscott ("Wikinomics: How Mass Collaboration Changes Everything"), and Yochai Benkler ("The Wealth of Networks").

¹ Seuss. *Yertle the turtle, and other stories*,. Random House, New York, 1958. ISBN 0394800877 9780394800875

² Francesca Polletta. *Freedom Is an Endless Meeting: Democracy in American Social Movements*. University Of Chicago Press, 1 edition, May 2004. ISBN 0226674495; and Andrew Cornell. *Oppose and Propose: Lessons from Movement for a New Society*. AK Press, April 2011. ISBN 9781849350662

Given the wealth of attention around decentralization, we should expect a similarly growing expertise in understanding the operations of decentralized groups. But for all the rhetoric surrounding the notion of flatness, there is a surprising paucity of precision in their mechanisms. Instead, there is a tendency among pundits to conflate organizations and movements that operate at vastly different scales, with vastly different structures, and vastly different purposes, all under the various banners of “leaderless”, “non-hierarchical”, “decentralized”, “horizontal”, and “peer-to-peer”, including:

- Protest movements like Occupy Wall Street
- Ad-hoc organizations like Anonymous
- Peer content projects like Wikipedia
- Free and Open Source Software
- Standards processes like IETF, IEEE, or ICANN
- “Flat” companies like 37signals or Valve
- Peer-to-peer file sharing services like BitTorrent
- Peer-to-peer currencies like Bitcoin
- Federated social networks, like Diaspora or StatusNet
- Decision making processes in cooperatives and collectives
- The Internet itself

Even the harshest critics of the strain of utopianism running through these accounts (such as Evgeny Morozov, who damns them all as “internet-centrism”³) have tended more towards lumping-together than teasing-apart, all under a single fuzzy banner of horizontalism. This ideographic mixing of meanings draws broad axes of alignment between disparate and often incompatible political positions, connecting peer production to free software to workplace solidarity to global justice movements to startup culture. I believe it goes far enough to regard the terms of horizontalism as *condensation symbols*⁴ – emotionally-laden terms of imprecise meaning that enable people with vastly different goals and politics to cherry-pick their preferred referents.

Condensation symbols, however, are a terrible place to start for software or process design – where the underlying phenomena express such variation, we need more precision in our language. There are important distinctions in the decentralism that powers BitTorrent, the atomization of labor in Amazon’s Mechanical Turk,

³ Evgeny Morozov. Why social movements should ignore social media. *The New Republic*, February 2013b. ISSN 0028-6583. URL <http://www.newrepublic.com/article/112189/>; and Evgeny Morozov. The meme hustler. *The Baffler*, 22, March 2013a

⁴ Doris A Graber. *Verbal behavior and politics*. University of Illinois Press, Urbana, 1976. ISBN 0252002628 9780252002625

Wikipedia's governance, and the models of organizing and decision making favored by activist affinity groups – and these distinctions need to be teased out with far greater clarity if one is to design tools that support their different needs. This is important both on technical dimensions (for example, deciding between server architectures for tools designed to support politically motivated horizontalists) and on design dimensions (for example, designing the mechanisms groups use to do administrative maintenance of group membership in a decision making platform).

My goal in this chapter is to provide some of the clarity needed to design systems for decision making in non-hierarchical groups. I will undertake the task of semantic distinction (identify the different ways people use the words flat, horizontal, decentralized, non-hierarchical), and construct an analytical framework to describe control in these systems. As the terms are often value-laden, I will pay particular attention to the motivations behind projects' and groups' appropriation of these terms. To guide analysis of decentralized systems, I will identify three mechanisms for decentralization (replication, federation, and subsidiarity), three dimensions of control which can be expressed hierarchically or non-hierarchically (access, evolution, and vision), and four analytical lenses which must be considered for any dimension of control (technical capabilities, rules and protocols, exclusive meanings, and structural inequalities). I will then discuss how these concepts relate to non-hierarchical decision-making processes and leadership.

Motivations for flatness

Since terms like “hierarchy” and “decentralized” are often used in a value-laden, emotionally thick manner, it's instructive to start from the motivations for flatness in order to begin to tease apart the dimensions around which these concepts operate. We can broadly divide them into two groups: *normative* motivations, which involve moral or value-based claims about decentralization, and *instrumental* motivations, which involve claims about the practical effectiveness of decentralization for particular outcomes.

Normative motivations for decentralization (usually expressed using terms like “non-hierarchicalism” or “horizontalism”) are based around the idea that it is desirable to wrest control from a privileged group or individual, and cede that control to a wider group. Proponents of normative motivations might speak of *democracy*, *sociality*, or *liberation* as goals of their systems, in contrast to *authoritarianism*, *individual profit*, or *proprietaryness*. As an example, the Diaspora social network advertises itself largely in terms of the desire for individual

control – and even expresses this in moralistic terms: “Connecting socially is human nature. You shouldn’t have to trade away your personal information to participate.”⁵ Similarly, the Free Software Foundation expresses the motivation for free software in terms of liberation: “As our society grows more dependent on computers, the software we run is of critical importance to securing the future of a free society.”⁶ Advocates of normative motivations for decentralization emphasize the social harm caused by allocating power to the few, and the liberating potential of distributing that power more equitably.

Instrumental motivations for decentralization (usually expressed using terms like “peer-to-peer” or “distributed”) focus on the potential for decentralized structures to function as better solutions to particular technical challenges. These include notions such as *robustness*, *creativity*, *efficiency*, and *quality*. The packet switching networking algorithms that power the Internet were first developed as a means of creating computer networks which would be robust against a nuclear strike that could destroy any central network hub.⁷ This was not based on any notion of liberation or freedom from government authority; on the contrary, it was to preserve government authority in the event of major infrastructural damage. Other instrumental motivations for decentralization include the potential for creativity, efficiency, and quality to be improved by broadening the base of people contributing ideas, reducing the friction of communication between layers of management, and reducing the transaction costs for micro-work. The “Open Source Initiative” and its rhetoric around “open source software” embodies an instrumental emphasis on the advantages of liberally licensed software, in contrast to the normative emphasis on freedom promoted by Free Software Foundation.

While proponents of a particular decentralized process or technology often move easily between normative and instrumental descriptions, a key difference is in that instrumental motivations are contingent on successful outcomes, while normative motivations hold regardless of outcome. A democratic worker’s collective might hold steadfastly to their horizontal meeting practice, even if efficiency and profits would increase by hiring a manager instead, because the group believes that it is morally better to retain equitable decision making. By contrast, groups motivated toward decentralism primarily for instrumental reasons would drop that strategy if evidence showed it to be less effective than a centralized alternative.⁸ The historical swing between paradigms of mainframe computing to desktop computing to cloud computing demonstrate mainstream commercial agnosticism to ideological decentralization, and the danger (from the perspective of a normatively motivated horizontalist) of reliance on purely instrumental motivations.

⁵ The diaspora project. URL <http://diasporaproject.org/>

⁶ Free software is a matter of liberty, not price - free software foundation - working together for free software. URL <https://www.fsf.org/about/>

⁷ Alexander R. Galloway. *Protocol: How Control Exists After Decentralization*. MIT Press, 2004. ISBN 9780262072472

⁸ Benjamin Mako Hill provides elucidation of this distinction in the context of free/libre open source software, arguing that claims by proponents of “Open Source” that liberal software licenses result in higher quality code break down when, as is very often the case, free software projects aren’t in fact higher quality than proprietary alternatives. Hill [2009]

An example of a conventionally instrumental focus on decentralization is how the startup 37signals describes their management structure:

Besides being small, 37signals has always been a flat organization. In fact, flatness is one of our core values. We have eight programmers, but we don't have a chief technical officer. We have five designers, but no creative director. We have five people on our customer support team, and no customer support manager... We've experimented with promoting a few people to manager-level roles. In some cases, this has worked out; in others, it hasn't. But one thing we've found is that groups that manage themselves are often better off than groups that are managed by a single person. So when groups do require structure, we get them to manage themselves.⁹

Not only does 37signals describe instrumental reasons for flatness, they also discuss experiments in which they hired management to see if it would work better. By contrast, the language C.T. Butler (a prominent writer and trainer in formal consensus decision making) uses to describe non-hierarchical group process in his article "A Revolutionary Decision-Making Process" speaks about "equal access to power", accountability, and ensuring that "resolutions were in the best interest of everyone involved."¹⁰ For Butler, considerations of short term pragmatic efficiency are of secondary importance to the long-term liberating potential of horizontal processes.

⁹ Jason Fried. Why i run a flat company, April 2011. URL <http://www.inc.com/magazine/20110401/jason-fried-why-i-run-a-flat-company.html>

¹⁰ C.T. Butler. A guide to formal consensus. URL <http://consensus.net/revolutionary.html>

Ways of being flat

When we say that an organization or project is decentralized or non-hierarchical, what do we mean?

Anarchistic activists such as David Graeber sometimes separate organizers into two camps: "horizontalists" who favor "anarchist principles of organization, non-hierarchical forms of direct democracy", versus the more pejoratively labeled "verticalists" who "actually like marching around with pre-issued signs and listening to spokesmen from somebody's central committee". [Graeber, 2012, p 27] Graeber invokes a definition of horizontality which emphasizes the lack of formal structural division of power, as one would find expressed through org charts, "bosses", "core organizers", or a "central committee". Horizontal groups such as those described by Graeber find ways to develop coherent behavior without a traditionally conceived "leader". These groups *do* often have at least temporary structural divisions of labor into roles (facilitators, vibes watchers, stack keepers, etc.), but these roles are not intended to afford any greater power to impact the content of a decision – only to guide the process by which it is reached. The form and success (or lack thereof) of these types of

roles are discussed in more detail in chapter 2.

This definition is fairly consonant with the one that Yochai Benkler proposes for “decentralization” in the context of peer production. Benkler identifies in decentralization the basic problem of making the behavior of autonomous individuals coherent:

“Centralization” is a particular response to the problem of how to make the behavior of many individual agents cohere into an effective pattern or achieve an effective result. Its primary attribute is the separation of the locus of opportunities for action from the authority to choose the action that the agent will undertake. Government authorities, firm managers, teachers in a classroom, all occupy a context in which potentially many individual wills could lead to action, and reduce the number of people whose will is permitted to affect the actual behavior patterns that the agents will adopt. “Decentralization” describes conditions under which the actions of many agents cohere and are effective despite the fact that they do not rely on reducing the number of people whose will counts to direct effective action. [Benkler, 2006, p 62]

This more precise definition helps to capture some important aspects of decentralization. The first is the explicit presence of *coherence*. We wouldn’t describe an unorganized collection of unconnected events as “decentralized” or “horizontal” – it’s only when we find a process that exhibits some form of coherence or unity of purpose that we reach for those terms. Decentralization in one dimension necessarily implies centralization (expressed via terms like coordination, coherence, cooperation) on another dimension – whether that dimension is a protocol, a set of rules, a product, or a collective goal.

The problem of designing non-hierarchical decision making processes for groups, then, is the problem of finding mechanisms for coherence which do not inherently favor allocating control to a small number of agents (whether people, servers, etc.), but also encouraging centralization on some other dimension (such as a protocol), thereby enabling coherence. As an example, the Internet Protocol which underlies all communication on the Internet defines a decentralized server infrastructure, but it relies on centralization of the way servers communicate with each other. Wikipedia may have decentralized content sourcing for articles, but contributions are made to a single central database of articles. BitTorrent may have decentralized server infrastructure (no central servers, and multiple competing client implementations), but it has a single common protocol controlled by a for-profit company (BitTorrent, Inc.). Flat companies may reduce the burden of layers of management, but they retain centralized control over the flow of capital. Worker-owned cooperatives might democratize control over the flow of capital, but centralize on protocols or procedures for deliberation. Horizontal

assemblies of activists may lack a central committee to direct them, but they have common protocols for communication, techniques for facilitation, and centralized allocations of resources into meeting places, infrastructure, and actions.

Since decentralizing in one dimension while retaining coherence requires centralizing another dimension, it is possible to decentralize part of a system, but to nevertheless retain centralized control – potentially in opposition to the original motivations for decentralization. A partisan for decentralization who seeks the liberatory potential of being free from authoritative control might be unhappy with the way that Amazon’s Mechanical Turk decentralizes the performance of labor, but retains (and potentially amplifies) the uneven power relationship between bosses and employees, and structurally eliminates the possibility of collective bargaining. In order to evaluate whether a system fulfills the motivations for decentralization, we must evaluate the full spectrum of control used by it, and consider how each aspect of control contributes to the outcomes.

But first, let us consider the mechanisms by which a system can be decentralized.

Mechanisms for decentralization

As defined above, decentralized systems are those which have methods of coherence that distribute control among all of the agents using the system.

There are three ways by which a system can decentralize: *replication*, *federation*, and *subsidiarity*. These different techniques of decentralization may operate simultaneously in different parts of a single system.

Replication is the simplest means of decentralizing – allow people to copy the centralized resource and deploy their own. This allows others to use the resource without the cost of designing it from scratch, and can help to avoid censorship or abuse from centralized control over the original. As an example, the Wordpress blogging engine can be described as decentralized through replication. Blogging is a practice that many people enjoy engaging in, but it would be arduous for each person to develop new blogging software every time they wanted to start a new blog. Freely distributable blog engines like Wordpress allow individual bloggers to host their own blog without duplicating the effort of creating a sophisticated tool. There is no special network or protocol that connects different installations of Wordpress together – once downloaded, each installation stands on its own. This allows full control by the administrators of each installation, but also means that they must develop practices to keep up to

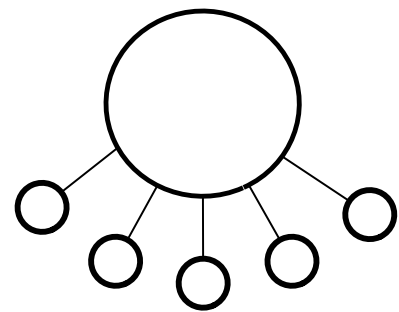


Figure 1.1: A centralized system.

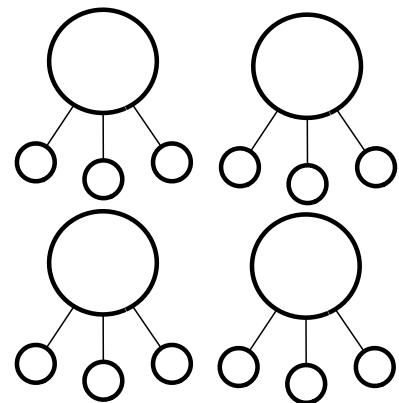


Figure 1.2: A replicated system. The same functions are repeated across multiple unconnected nodes.

date with security patches and upgrades from “upstream” (and due to the broad centralization on a single code base, Wordpress becomes an attractive target for attacks). Similarly, Wikipedia distributes its database of articles under free licenses, allowing other projects to replicate them. The editing of documents, however, is still centralized on Wikipedia’s servers, and therefore projects which replicate the database must regularly seek updates from Wikipedia in order to have the most up-to-date articles.

Federation is a method for decentralization where the site of coordination is shifted from a resource to a protocol. Where replication focuses on duplicating a resource, federation distributes the authority of a resource among multiple nodes, allowing continued networked production without coordination with any centralized source. Ward Cunningham’s “Smallest Federated Wiki” is an wiki system (collaborative user-editable website) which does not rely on a centralized database of articles. Where Wikipedia only considers edits of articles on their own servers to be authoritative, in a federated wiki, anyone can edit articles on their own servers. If any editor likes the changes that other editors have made to their copies of the article database, they can “pull” the changes from the remote database to their own through a common protocol that each wiki implements. Where a central database has a potential for censorship (Wikipedia has to date been censored in China, Iran, Syria, Pakistan, Thailand, Tunisia, and Uzbekistan, among others¹¹), a federated database that avoids the centralized source could be more resilient.

The peer-to-peer file sharing system BitTorrent is another example of a federated system. The protocol is common across the BitTorrent network, but multiple competing software client implementations can all interoperate – any client can join the network as long as it adheres to the protocol. BitTorrent relies on no central servers (unlike preceding technologies like the original Napster, which relied on a central service to coordinate peer-to-peer connections¹²), and is therefore less vulnerable to efforts by authorities of control to shut it down, as there is no single server which can be seized or blocked.

Still, BitTorrent relies on some centralization – the need for easily searchable databases of files available for sharing led to the rise of “tracker sites” such as The Pirate Bay, which were vulnerable to seizure or censorship by governments.¹³ BitTorrent developers have worked around this in part through the addition of federated search functions to the protocol. Even with this innovation, BitTorrent still requires network infrastructure to function, which usually means relying on centralized Internet Service Providers such as Comcast or Verizon, who have the capacity to throttle and block BitTorrent traffic identified through traffic analysis.¹⁴

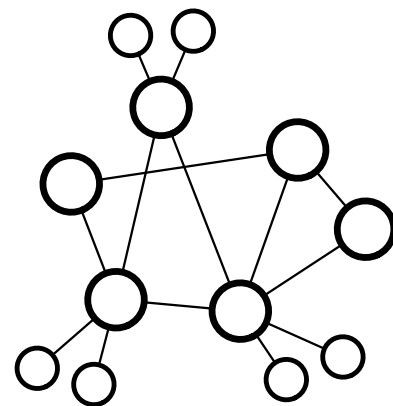


Figure 1.3: A federated system. The nodes interoperate via a shared protocol, and each node is dispensable.

¹¹ Censorship of wikipedia - wikipedia, the free encyclopedia. URL https://en.wikipedia.org/wiki/Censorship_of_Wikipedia

¹² Jeff Tyson. How the old napster worked, October 2000. URL <http://computer.howstuffworks.com/napster.htm>

¹³ Many Torrent Tracker sites have been shut down through seizure of their domain names [Sellars, 2011] – the Domain Name System is one of the centralized components of the Internet. The theoretical limits of decentralization in naming are discussed on page 26.

¹⁴ Milton L. Mueller and Hadi Asghari. Deep packet inspection and bandwidth management: Battles over BitTorrent in canada and the united states. *Telecommunications Policy*, 36(6): 462–475, July 2012. ISSN 0308-5961. DOI: 10.1016/j.telpol.2012.04.003. URL <http://www.sciencedirect.com/science/article/pii/S0308596112000523>

Various projects have attempted to decentralize this basic networking infrastructure, as exemplified by neighborhood mesh networks, which enable peer-to-peer network connections without reliance on central service providers. Mesh networks also decentralize through federation – by centralizing on a mesh protocol¹⁵, access points from a variety of vendors can be assembled to form a network, where (assuming adequate density of nodes) no single node is critical to the network’s function. This same effort to decentralize could be taken further by gaining independence from the power grid through local power generation (centralizing on the “protocol” of frequencies and current characteristics for electricity used by the various devices, but distributing the generating capacity through replication of generators). However, to date, there are still practical limits to such decentralization. One cannot yet reasonably fabricate semiconductors or other electronic components without reliance on economic systems and infrastructure that is ultimately under the control of centralized authorities. Rather than follow resource decentralization to these extremes, it’s more common for a service to function in a hybrid context, where some of its layers of operation work decentrally, while others are centralized.

Subsidiarity is the third method by which a system can be decentralized. Subsidiarity is the principle of reducing the number of coordinating agents to the minimum required for a particular task – in other words, a division of concerns. Projects which implement plugin systems for their platforms express subsidiarity by allowing other developers to retain full control over particular sub-tasks (the tasks performed by the plugins). The Drupal web-based content management system illustrates this to a slight extreme, where even the basic functionality of the system is implemented with pluggable modules. Like federation, subsidiarity relies on protocols for interoperability between components (or at least loose shared conventions, as with the “Unix philosophy” of small, simple, modular parts); but unlike federation, subsidiarity involves different components taking on functionally different roles, rather than distributing the same functional role. Large protest groups such as Occupy Wall Street or the Global Justice movement exercised subsidiarity through the form of working groups, which take on distinct roles such as “media”, “sanitation”, “direct action”, or “anti-oppression”. Boards of Directors exercise the same process through committees.

Each of these mechanisms of decentralization might overlap within a single system – for example, the Wordpress blogging engine distributes through replication, but its ecosystem of themes and plugins distribute through subsidiarity. The BitTorrent client software distributes through replication, but its operation distributes through

¹⁵ Numerous routing protocols are currently in use, though IEEE is working on standardizing a set of protocols the title 802.11s

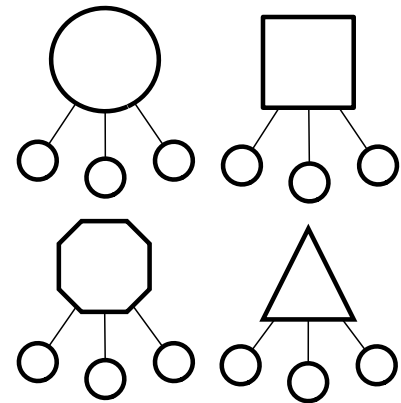


Figure 1.4: A system with subsidiarity. Different responsibilities are distributed across different nodes, which need not be connected.

federation. But in every case, in order to distribute control in one dimension, you must centralize control in another – that centralized part is whatever it is that leads us to describe the distributed system is “coherent” or “coordinated”, whether a shared resource, a standardized protocol, or simply a shared set of values or purposes. When distributing a software project by replication, more people are centralizing on a single code base. To offer a plugin architecture, one must centralize on a set of APIs for interoperability between components. To federate production among multiple servers, there must be a common protocol for the servers to communicate with.

An illustrative paradox in the problem of decentralization is *Zooko’s Triangle*¹⁶, a hypothesized trade-off for decentralized naming, pertinent to the design of Domain Name Systems. The Internet’s Domain Name System is a centralized registry that maps domain names (such as “intertwinkles.org”) to Internet Protocol addresses. As a centralized system that has been used repeatedly as a means of censoring information, there is considerable interest among partisans of decentralization in replacing it with a more decentralized system. A challenge to these efforts is the hypothesis proposed by Zooko Wilcox-O’Hearn that for any naming system, only two of three desirable traits are possible:

- Uniqueness: only one specific entity has a given name.
- Human-meaningful: the name has semantic meaning to people.
- Decentralized: there is no centralized authority determining the meaning of a name.

Decentralized unique names can easily be obtained by choosing random numbers of sufficiently high diffusion. Universally Unique Identifiers, for example, are random identifiers with 2^{128} possible values, giving probabilities of collision among randomly chosen names that are low enough to be ignored. The identifiers (e.g. c622add4-5761-4988-9229-ee81092554d3¹⁷), however, have no human meaning. Human-meaningful, but non-unique names can be made without any centralized registry (for example, names given to babies). Human-meaningful and unique names can be obtained by maintaining a central registry of names (as with the ICANN’s Domain Name System). But there is no way to have a naming system that offers unique, human-meaningful names, with no centralized arbiter.

Aaron Swartz proposed that Zooko’s Triangle could be solved using a variant of BitCoin’s distributed algorithm for generating “coin”.¹⁸ BitCoin operates by a proof-of-work system along with a network protocol for federating nodes, which affords non-duplicable

¹⁶ Mark Stiegler. Introduction to petname systems, June 2010. URL <http://www.skyhunter.com/marcs/petnames/IntroPetNames.html>

¹⁷ Guaranteed to be random.

¹⁸ Aaron Swartz. Squaring the triangle: Secure, decentralized, human-readable names (aaron swartz’s raw thought), January 2011. URL <http://www.aaronsw.com/weblog/squarezooko>

transactions in a distributed system. By leveraging a similar proof-of-work, one could simply “claim” a human-meaningful name, perform a work-unit of arbitrary computation which authenticates the claim, and then share evidence of the work with other nodes in the network, who then mark that name as taken. In this way, a shared ledger of names and owners could be grown without ever having consulted any centralized gatekeeper for names. This proposal has been formalized in the form of the Namecoin protocol, which now offers “.bit” domain names, outside of any control by the traditional domain name system.¹⁹

Attentive readers will note by now that this clever solution does not eliminate centralization: rather, it moves the site of centralization from an authority such as ICANN²⁰ to a protocol for federation of proofs of work along with a central (but replicated) ledger of names. That protocol may enforce just, unjust, or neutral outcomes – one must further analyze the function of control in a protocol in order to evaluate whether it satisfies liberatory motivations for decentralization. In the end, we might well trust an algorithm over a fallible human authority; but we should exercise caution when selecting algorithms to govern our lives.

¹⁹ <http://dot-bit.org>

²⁰ The Internet Corporation for Assigned Names and Numbers, the governance body that controls the Internet’s traditional domain name system.

An analytical framework for control

Systems that strive for non-hierarchicalism can’t reach their goals of liberation through decentralization alone, since that decentralization means some other dimension has been centralized. Similarly, federated or open source software projects don’t automatically give their users greater freedom just because they’ve standardized on protocols for access. Recalling Yochai Benkler’s definition of decentralization (“conditions under which the actions of many agents cohere and are effective despite the fact that they do not rely on reducing the number of people whose will counts to direct effective action”²¹), in addition to basic use of a system, we should consider control over how systems evolve over time in both function and meaning as part of the “effective action” under question. Furthermore, we should examine this from both technical and social dimensions.

To go beyond simplistic conceptions of technical access, I’ve developed the following analytical framework to view a system through four analytical lenses across three functional dimensions of control, as shown in table 1.1. Together these expand into 12 questions about ways in which control is available to all participants or constrained to a few. While we can’t expect any single framework to comprehensively capture all of the complex ways that situational differences result in different group outcomes, my hope is that this framework

²¹ Yochai Benkler. *The wealth of networks: how social production transforms markets and freedom*. Yale University Press, New Haven, 2006. ISBN 9780300125771 0300125771 9780300110562 0300110561

can expand the range of questions that one asks when evaluating the distribution of control in a system.

Analytical framework for control	
<i>Analytical lens</i>	<i>Functional dimension</i>
How do { technical capabilities rules and protocols exclusive meanings and practices structural inequalities }	impact who participates in { access evolution visions, meanings } ?

Table 1.1: Questions to guide an analysis of control.

Analytical lenses are ways of looking at a system. We can consider a system as a technical system which has certain affordances or capabilities, as a group of people with internal meanings, or as a social system embedded in a larger social context – in each case, we are looking at the same system, but emphasizing different ways of interpreting it.

- *Technical capabilities* refer to the affordances and capabilities of the technical structure of a system. For example, if we consider the editing of Wikipedia articles from this lens, we will look at the computer requirements, and the affordances of the editing software, but not the policies or social factors. For software development projects, access to source code repositories and cryptographic signing keys is a technical constraint. For face-to-face meetings, we would consider the physical constraints of the meeting venue. When we consider control from the lens of technical capability, we analyze the ways in which the system affords access, evolution, and creation of meanings by virtue of what is technically possible given the structure of its resources.
- *Rules and protocols* refer to administrative, legal, or social rules for participation, whether stated or unstated. In software projects, this might include licenses, contracts, and procedures for contributing code. In public meetings, this might include rules of order and process. Unstated rules can present particular barriers to access, as they make it difficult for new participants to learn how to appropriately contribute.
- *Exclusive meanings and practices* refer to the internal styles, attitudes, and expectations of members of a group which influence its function, but which may be unwelcome or illegible to others. This might include forms of dress and adornment, diet, language, or demeanor. A business culture that expects formal dress and ivy league diction will be unwelcoming to some; similarly, punk aesthetics and street slang might be unwelcoming to others. While some exclusive meanings and practices might be indicative of structural inequalities (e.g. male, white, or class dominance), oth-

ers are orthogonal to structural inequalities, but are nonetheless alienating.

- *Structural inequalities* refer to the ways that the inequalities present in the wider social context are replicated in a system's internal operations. As an example, a system which requires expensive technology to use replicates financial inequalities. A system where participants frequently use language hostile to women will perpetuate gender inequalities. Systems that lack structures to resist structural inequalities in the social context in which they are embedded will tend to perpetuate them.

Functional dimensions are the basic ways that participants engage with a system.

- *Access* describes whatever is the usual ongoing use of a system. For systems that are based on peer production (such as Wikipedia), access might include editing or contributing as well as reading. For software projects, this might include using the software. For affinity groups, this might include meeting and organizing with others in the group. Control over access concerns structures that allow or prevent different people from using a system.
- *Evolution* describes the process of changing how a system is accessed. For software projects, evolution might mean developing the software further. For a public meeting like an Occupy General Assembly, evolution means changing the procedures of General Assemblies or creating new ones. Control over evolution concerns structures that allow or prevent different people from contributing to structural changes in a system.
- *Vision* describes the creation of the meaning, future directions, and strategies for the evolution of the system. Synonyms for vision include *leadership* and *strategy*. A software project might have its vision controlled by a single person (e.g. a "benevolent dictator"), or open to a wider group (e.g. a "community council"). Similarly, a popular movement might be driven by a charismatic leader, a central committee, or by direct assemblies. Vision is the process of defining what a system *is* (or what it *isn't*), which guides evolution. Control over vision concerns structures that allow or prevent different people from defining and directing the meaning of a system.

To analyze control in a system, answer each of the 12 questions for the system in question. Exhaustive analysis helps to identify less obvious mechanisms of control within a system. This analysis is applied to InterTwinkles in chapter 5 (page 135). To further elucidate

how answers to the questions might identify mechanisms of control, I will now cherry-pick compelling examples of ways that each question has impacted well-known nominally decentralized systems.

How do technical capabilities impact who participates in...

1. Access: The free, paywall-less access offered by Wikipedia differs from every successful encyclopedia that preceded it. Prior commercial encyclopedias used access as their means of monetization – purchase of the physical volumes, or subscriptions to online services, funded the editorial staff which produced the articles. Wikipedia succeeded (where others failed²²) in creating an alternative model where the articles are distributed freely, but the editors are also unpaid, leaving only administrative staffing and server maintenance as expenses, which can be covered through donations. This liberal access to articles in Wikipedia is institutionalized protocologically by the free content licenses under which they are distributed. All articles in Wikipedia are licensed using a “copyleft” license, which ensures that anyone is free to copy and redistribute the articles, but requires that this freedom also be distributed along with it.

The location of access for Wikipedia articles remains largely centralized: almost all people accessing Wikipedia’s content do so using the Wikimedia Foundation’s servers; in contrast to technologies with decentralized infrastructure such as BitTorrent. We could imagine that if the foundation were to turn against the interests of the public, or if a government were to ban access to the foundation’s servers (as China, Iran, Syria, Pakistan, Thailand, Tunisia, and Uzbekistan have at various points), people would lose the ability to access articles. Wikipedia alleviates this somewhat by offering data dumps and APIs which make it easy to *replicate* databases of articles – and several websites, and even some dedicated mobile devices, do so. Still, a strategy of replication requires ongoing maintenance to keep the copy up-to-date. A decentralizing strategy of *federation* which replaces the central server infrastructure with a central protocol for distributed nodes could be more robust, by eliminating a single point of control. Whether such a drastic change in infrastructure could be successful or worthwhile is anybody’s guess; in the case of technologies like BitCoin (a central goal of which is to eliminate any need for centralized trust), or BitTorrent (designed to function in a regime of heavy censorship), the strategy has thus far proved successful. Prior to Wikipedia’s demonstrated success in decentralizing article creation, many would have doubted its viability as well.

2. Evolution: Software developers store code using *version control systems* – databases (“repositories”) of changes in the code that track

²² Benjamin Mako Hill. Almost wikipedia: What eight collaborative encyclopedia projects reveal about mechanisms of collective action | berkman center, October 2011. URL <https://cyber.law.harvard.edu/events/luncheon/2011/10/makohill>

when, how, and why code was added, changed, or removed, as well as facilities for merging changes made simultaneously by different developers. Version control systems help manage both collaboration between different developers, and the changes to code over time. Historically, version control systems were primarily centralized, typically stored on a server maintained by the primary developer of the software project. This arrangement created a relationship of power: even in open source projects that allowed developers to download a copy of the code base, they could only share changes they made to the code by sending them directly back to the owner of the repository, or by “forking” the project (considered hostile) and trying to attract others to the new repository with new history.

In the last decade, developers have created *distributed version control systems* (such as Git, Mercurial, and Darcs) which change this relationship by giving each developer a complete copy of the project’s history, and making merging between different developers’ copies trivial. As a result, it is now possible for developers to maintain parallel repositories free of the control of centralized maintainers, and “forking” is now a trivial and unproblematic operation.

3. Vision: The Debian Project is a free software project which develops the most popular and widely used Linux-based operating system. In 1998, the project adopted a constitution, and subsequently developed a process for voting on resolutions to determine the future of the project (via a “Standard Resolution Procedure”), as well as an election process for project leaders that serve a 1-year term. To support this process, the project relies on mailing lists for ongoing discussion, and a custom-developed voting tool.²³ The communication and voting technology supports a much larger population of participants to contribute to procedural matters than would be practical in the absence of the tools.

A recent resolution passed by the community of developers is a diversity statement, approved June 2, 2012:

The Debian Project welcomes and encourages participation by everyone.

No matter how you identify yourself or how others perceive you: we welcome you. We welcome contributions from everyone as long as they interact constructively with our community.

While much of the work for our project is technical in nature, we value and encourage contributions from those with expertise in other areas, and welcome them into our community.

Out of 951 registered developers, 270 participated; 251 voted in favor of the resolution, and 19 voted against.²⁴

²³ GUNNAR Ristroph. Debian’s democracy. *Online Deliberation: Design, Research and Practice*. Chicago, Illinois, USA: Center for the Study of Language, pages 207–212, 2009. URL http://odbook.stanford.edu/static/filedocument/2009/11/15/Chapter_17._Ristroph.pdf

²⁴ http://www.debian.org/vote/2012/vote_002

How do rules and protocols impact who participates in ...

4. Access: While it wasn't the first to do so, the Occupy movement eschewed a centrally organized strategy for decision making. Many earlier protest movements – such as A.N.S.W.E.R., a major coalition of anti-war groups primarily active in the early 2000's – operated with a central steering committee that directed protest activities (these are groups of the sort panned by David Graeber as the “verticals”). Other groups such as the Direct Action Network, which organized for Global Justice against the WTO and IMF in the late 90's and early 2000's used a more horizontal approach with “spokes councils” representing large numbers of autonomous affinity groups. Occupy, however, raised the ante on horizontal decision making by publicly characterizing themselves as representing “the 99%” rather than a particular partisan subset (striving to do so with a General Assembly model), and further by making the democratic meeting processes a central part of their public activity rather than just a tool for internal organization. Occupy standardized on a set of rules and protocols for how General Assemblies would operate, which enabled anyone who showed up to participate in decision making.

Participation in a General Assembly was still not easy for just anyone – one would still have to find (or worse, found) an encampment, and then take the time to be physically present through long meetings, learn the rules of order including hand signals and proposal processes. Facilitators tried to help democratize access by providing regular instruction on process at the beginning of every meeting. Nevertheless, the meetings could be strenuous – in Zucotti Park, at times exceeding seven hours. Despite these arduous requirements, it is likely that the General Assembly format presented a lower barrier to entry for average people than developing the status and personal relationships necessary to have a seat in a central committee would.²⁵

5. Evolution: The Gnu General Public License is a famous “hack” of the copyright system which aims to make it possible to guarantee that users of a software project will always retain basic freedom over the code. It does so by using copyright law to enforce the rule that any time a changed version of the software project is distributed, the changes made to the software be also made available to the original developers. These legal rules, however, do nothing to ensure that developers will continue a spirit of cooperation and coordination around the code.

An example of where this can harm a project is in Apple Inc.'s handling of the rendering engine for their Safari web browser. Apple used the GPL-licensed “KHTML” rendering engine, originally developed for use on Linux systems, as the basis for the rendering engine

²⁵ Further critiques of the accessibility of consensus-oriented decision making processes are discussed in chapter 2.

in Safari. However, when Apple released their changes as required by the license, they did so with a “code dump” – thousands of lines of code with no comments or explanations, no revision histories, and none of the tools used by programmers to interpret changes of source code. As a result, the labor required to incorporate these changes back into the original KHTML engine approached the complexity of developing the software in the first place, reducing the benefits of decentralization.²⁶ Replication alone as a strategy for decentralization does little to ensure continued decentralization in ongoing production, without cultures or protocols that coordinate ongoing development.

6. Vision: Many software projects – even those that self-describe as “open” and “free” and use copyleft software licenses – are nevertheless explicitly anti-democratic, and guided by a self-appointed “Benevolent Dictator For Life” (BDFL).²⁷ While much of the day-to-day development of the projects may be handled by other developers, the BDFL has the privilege to set development direction, resolve disputes, and veto the inclusion of new features.

Mark Shuttleworth is the founder and BDFL for the Ubuntu operating system, a distribution of GNU/Linux which places particular emphasis on user-friendliness. Ubuntu maintains an explicit governance policy, which sets out a system of governance with a Community Council and a Technical Board, but also enshrines Shuttleworth as the BDFL with the power to direct the work of employees and cast a vote on the council and board.²⁸ The existence of the Community Council and Technical Board, as well as explicit procedures for joining these bodies, greatly increases the democratic accessibility for setting the vision and direction for Ubuntu. Nevertheless, the presence of the BDFL sets a limit to the power that people can exercise through those bodies.

*How do **exclusive meanings and practices** impact who participates in ...*

7. Access: Many housing cooperatives (including those who constitute the user groups in the InterTwinkles field study) pride themselves as inclusive communities who welcome people from all cultural backgrounds, races, genders, and sexualities. However, some cooperative communities also have “weird” cultures. Some examples include a tolerance for public nudity, strict vegetarianism, tolerance for low levels of cleanliness, and open drug use. When a new member considers joining such a house, they may be put off by these practices – either through distaste (e.g. living in less clean spaces), through an incompatibility in culture (not being able to eat food they

²⁶ Joe Brockmeier. Apple and KHTML [LWN.net], May 2005. URL <https://lwn.net/Articles/136538/>

²⁷ A small sampling of notable free software projects which use a BDFL model include: GNU, Linux, Python, Ubuntu, PHP, and Node.js.

²⁸ <http://www.ubuntu.com/about/about-ubuntu/governance>

want to eat), or personal risk (increased likelihood of arrest with the presence of drugs due to past convictions or status as a racial minority).

Some aspects of a cooperative's unique culture may be key components of its identity: for example, concern for animal welfare could be a strong enough shared value of members that they are willing to exclude people who want to eat meat. However, a group should be aware that these exclusive meanings may track with other demographics (such as whiteness, middle class backgrounds, particular religious backgrounds, etc.) which the group wishes to transcend in its values of inclusion.

8. Evolution: "Shit Reddit Says" (SRS) is a subreddit (a topic-focused community on the link aggregator and discussion site <http://reddit.com>), whose topic is pointing out the "shit" that happens in the rest of Reddit, as defined in its topic statement:

Have you recently read an upvoted Reddit comment that was bigoted, creepy, misogynistic, transphobic, racist, homophobic, or just reeking of unexamined, toxic privilege? Of course you have! Post it here.²⁹

While reddit itself is highly impacted by structural inequality, and SRS itself is a counterpoint seeking to fight that trend, here I'd like to focus on one particular aspect of the community that is orthogonal to questions of oppression: the way its community and admins use their design and rhetoric to create a highly exclusive environment. The community uses images and language that are likely to be alienating to many visitors, but also serve to cultivate internal strength and robustness against co-option.

The first rule of the subreddit ("rule x") is that it "is a circlejerk and interrupting that circlejerk is an easy way to get banned. For instance, commenters are not allowed to say 'This post is not offensive' or 'This is not SRS worthy.'" Participants are discouraged from deliberation or discussion about changes to the community. The style and copy text chosen by the administrators, as well as the comments by community members, frequently use a shorthand of idiosyncratic phrases and images. Some examples include: "misandry don't real", meaning misandry doesn't exist; "so brave", calling out comments that promote misogynistic or homophobic views from a rebellious stance; a small image of Ron Paul's head to point out arguments that misogynistic comments are protected by free speech; "shitlord" identifying people who make insensitive comments; and ample dildos.³⁰ These phrases and images provide a way of naming, claiming, and shaming common rhetorical tropes in arguments against feminism, queer liberation, racism, and other forms of structural oppression. SRS successfully uses them to invert the usual power relationship

²⁹ <http://reddit.com/r/shitredditsays>



Figure 1.5: When commenters in SRS break the rules (for example by criticizing or arguing against the validity of others' posts, or using racist, sexist or ablist slurs), moderators ban the users and leave a reply to their post with a colorful "BANNED" message, usually depicting purple dildos ("dildz") forming the letters. Above are three examples.

³⁰ A full list of the stock images available to commenters is at http://www.reddit.com/r/ShitRedditSays/comments/z8ry8/meta_smiley_guide/

found in the rest of Reddit, where complaints about insensitive discourse are often unwelcome.

Because of the loud, brash, illegible style and idiosyncratic rhetoric, SRS may appear hostile even to some who share the community's ethical positions on liberation. However, for those who do appreciate the opportunity to vent and yell in a space free from critique, SRS can provide a stable and welcoming caucus, which resists attempts to change it. We might not want all anti-oppression work to look like SRS, but it does provide a refreshing counterpoint to a sometimes dry and heady norm.

9. Vision: Linus Torvalds, creator and “Benevolent Dictator for Life” of the eponymous Linux kernel, has a famously acerbic personality, with no hesitation to berate, insult, and curse at developers with whom he disagrees. In July of 2013, kernel developer Sarah Sharp complained about this behavior on the kernel development mailing list, which discusses the development and future directions for Linux:

Seriously, guys? Is this what we need in order to get improve -stable? Linus Torvalds is advocating for physical intimidation and violence. Ingo Molnar and Linus are advocating for verbal abuse.

Not *fucking* cool. Violence, whether it be physical intimidation, verbal threats or verbal abuse is not acceptable. Keep it professional on the mailing lists. ³¹

Torvalds defended his action as coming from the best intentions of maintaining the kernel, and as a valid discourse style that he prefers to a faux “professional” style:

If you want me to “act professional,” I can tell you that I’m not interested... The same way I’m not going to start wearing ties, I’m *also* not going to buy into the fake politeness, the lying, the office backstabbing, the passive aggressiveness, and the buzzwords. Because THAT is what “acting professionally” results in: people resort to all kinds of really nasty things because they are force to act out their normal urges in unnatural ways.

For some, Torvalds’ brash style may be refreshing and liberating; for others, it may be off-putting to the point of discouraging participation. While Torvalds postulates that his temper is the only alternative to a back-stabbing passive-aggressive office professionalism, one could imagine a middle ground where people are able to honestly communicate with neither pretense nor hostility.

In contrast to Torvalds’ approach, other free software projects such as Ubuntu³² and Python³³ have Codes of Conduct which ask members to maintain respectful and considerate discourse.



Figure 1.6: Linus Torvalds flipping the bird. Image by *aaltouniversityace*, <https://www.youtube.com/watch?v=MSHbP30pASA>

³¹ Jon Brodtkin. Linus torvalds defends his right to shame linux kernel developers, July 2013. URL <http://arstechnica.com/information-technology/2013/07/linus-torvalds-defends-his-right-to-shame-linux-kernel-developers/>

³² <http://www.ubuntu.com/about/about-ubuntu/conduct>

³³ <http://www.python.org/psf/codeofconduct/>

How do structural inequalities impact who participates in ...

10. Access: Wikipedia has a highly gender-biased population of editors, with as high as 91% male editors as of 2011.³⁴ In the provocatively titled article “Free as in sexist?”, Joseph Reagle argues that structural inequalities left unchecked in the discourse and operation of Wikipedia replicates larger social inequalities. Worse, Reagle argues that the very values of freedom and openness in free culture projects such as Wikipedia can be misapplied as a rationalization for continued inequality:

[These] communities’ openness means that a minority of difficult members (including, for example, a sincere misogynist or an insincere troll) can disproportionately affect the tone and dynamics of interactions... [The] ideology and rhetoric of freedom and openness can then be used to (a) suppress concerns by labeling them as “censorship” and, to (b) rationalize low female participation as simply a matter of women’s choice.³⁵

Without developing systems to combat structural oppression internally, projects and groups are doomed to replicate those structures from the wider social context. The mere presence of values of openness and freedom are inadequate; and can in fact be a shield preventing the correction of abuses.

11. Evolution: Open Source software developer Jessamyn Smith was part of a software development team which she described as generally “clueful”, but nevertheless had a distasteful running joke in the chat room the team used to communicate with each other. The chat room had a “bot” (an automated program that responds to events in the chat room) which would look for phrases that could be taken as a double entendre, and respond with “that’s what she said” (emphasizing the opportunity for sexual innuendo). After confronting the other developers about the general distastefulness of the bot without success, Smith created an intervention in the form of a new bot: whenever someone in the channel (including the original bot) uttered “that’s what she said”, Smith’s bot would respond with a quotation from a notable woman. Smith named her bot “whatshereallysaid”.

This example demonstrates the way that male-dominated culture can creep into a group of developers, making the space unwelcoming to women. The responses of her colleagues to her intervention also demonstrate the resistance to change and awareness that people who are exposed to their unwelcome behavior can elicit. Smith recounted some of the responses:

There have been complaints that we have too many bots in the channel now. There have been complaints about it spamming the channel.

³⁴ Wikipedia:Wikipedians, August 2013. URL <https://en.wikipedia.org/w/index.php?title=Wikipedia:Wikipedians&oldid=566411546>. Page Version ID: 566411546

³⁵ Joseph Reagle. “Free as in sexist?” free culture and the gender gap. *First Monday*, 18(1), December 2012. ISSN 13960466. URL <http://firstmonday.org/ojs/index.php/fm/article/view/4291>

There were several “Make them shut up!” responses. These are not reactions I have seen the other bots elicit, certainly not with such intensity. One person even complained about the name being too long, though to his credit he realized right after he said that that several other people in the channel also have very long handles.

To me, all of this seems like typical geek behaviour: something is making them uncomfortable, and so they attack it on “rational” grounds. Most likely, they aren’t even aware of the gut reaction fueling their logic. Interestingly, the intensity of emotion seemed to carry over into subsequent discussions, including one about women in the Python community... I pointed out that it’s trivially easy to have the bot not say anything: don’t say TWSS.³⁶

12. Vision: In January of 2012, the Occupy Boston General Assembly faced a split over a contentious proposal to adopt a rule requiring people who had been classified as a “level 3 sex offender” under Massachusetts law to identify themselves to the Assembly and obtain special permission to continue to participate. The Massachusetts code defines “level 3” as a designation for people for whom “the risk of re-offense is high and the degree of dangerousness posed to the public is such that a substantial public safety interest is served by active dissemination”³⁷. The proposal was brought in response to concerns by members of the Women’s Caucus, Sexual Assault Awareness working group, and Safer Spaces working group that Occupy Boston was inadequately responding to sexual misconduct that had happened during the group’s encampment at Dewey Square in Boston. When the General Assembly blocked even a watered down proposal that would only reaffirm that Occupy Boston repudiates sexual violence, a third of the General Assembly walked out in protest, expressing that they did not wish to participate in a General Assembly that would not identify as countering systemic violence against women.

At issue in the discussion leading to the rift was tension over two opposing systems of structural oppression: women felt unsafe due to the presence of people with histories of sexual violence, and wanted to ensure that they would continue to feel safe and welcome in the General Assembly. At the same time, a number of members of Occupy Boston disagreed with the way the state Criminal Justice system worked, and in particular opposed the practice of branding people with purported criminal status (often illegitimately) through offender registries.³⁸ Since Occupy Boston’s General Assembly was the site for the creation of future directions, meanings, and leadership within the movement, alienating members who had safety concerns or strong philosophical disagreements limited who could participate in defining the movement.

As a group engaged in prefigurative politics with an intention of

³⁶ Jessamyn Smith. What she really said: Fighting sexist jokes the geeky way!, March 2012. URL <http://geekfeminism.org/2012/03/19/what-she-really-said-fighting-sexist-jokes-the-geeky-way/>

³⁷ Massachusetts Code, Part I, Title II, Chapter 6, Section 178K(2)(c), as of August, 2013.

³⁸ Minutes for the General Assembly, pre-walkout: <http://www.scribd.com/doc/77606108/OccupyBoston-GA-Minutes-1-8-2011>, post-walkout: http://wiki.occupyboston.org/wiki/GA_Minutes_Week_15_-_Jan_8_%282/2%29

challenging deeply embedded systems of injustice, Occupy Boston was left grappling with both the unaddressed cultures of violence towards women, and a model of punitive justice as the only available remedy. While there are ongoing efforts to develop group techniques for confronting sexual assault within a framework of transformative justice (e.g. ³⁹, a zine produced by a Philadelphia based group dedicated to this topic), this work remains at an early stage, outside of the working toolkit for most groups.

Non-hierarchicalism, Leadership, and Protocol in decision making⁴⁰

If the goal of non-hierarchical decision making is to distribute the power to make decisions among all of the members of a group, what is the role of *leadership* within non-hierarchical groups? There is a semantic tendency rampant both inside and outside activist communities to conflate the notions of “leaders” and “leadership” with hierarchy – the notion of imperatively directing action is conflated with the notion of sharing vision and ideas. Within these communities, this can lead to the rejection of cultivated skill and knowledge. In “Freedom is an Endless Meeting”, a history of democracy in American social movements, Francesca Polletta describes this tendency:

Equality has sometimes been interpreted as prohibiting any differences in skills or talents. What group members have viewed as effective leadership at one point has come later to be seen as manipulation. [Polletta, 2004, p 4]

In “Oppose and Propose”, Andrew Cornell’s history of Movement for a New Society (an influential prefigurative social movement active in the 1970’s and 1980’s), interviewee George Lakey describes this as leading to “covert leadership”:

At one point, an organizational development consultant volunteered to work with MNS because it seemed as an organization we were getting sick. She had us do an exercise where she said, “All of you who are leaders in the organization, you go over there.” So like three people, all blushing, go across the room. And she smiled and said, “OK, all of you who do covert leadership, you go over here.” And about a third of the room gets up... So it turned out there was there was this group of covert older male leadership – and this is so traditionally male, too, like we’re holding the family together. – *George Lakey Cornell* [2011, p 73]

Feminist writer Jo Freeman famously observed in “The Tyranny of Structurelessness”⁴¹ that if a group seeks to eliminate all overt structure, it doesn’t end up with equity or liberation. Instead, by

³⁹ Philly Stands Up. A stand up start-up. URL <http://www.phillystandsup.com/PDFS/AStandUpStartUp.PDF>

⁴⁰ Parts of this section derive from a blog post I wrote on this topic in February, 2013: <http://civic.mit.edu/blog/cfd/leadership-in-horizontal-movements>



An Occupy meeting facilitated by an organizer in a tiger suit. Are facilitators merely the hidden tiger of hierarchy? <http://wagingnonviolence.org/feature/what-does-leaderless-look-like/>

⁴¹ Jo Freeman. The tyranny of structureless, 1972. URL <http://www.jofreeman.com/joreen/tyranny.htm>

eschewing internal structure (such as rules for process or order), the group will simply import the external structure (such as white male dominance, elitism, and ideological narrowness) from the social context in which the group operates. Worse, when the group subscribes to values of equality and horizontalism, group members will be reluctant to acknowledge and confront patterns of dominance that emerge. Concern over this trend continues today, as in this observation of a latent “elitism” within the movement from an Occupy participant:

One of the consequences of just how difficult and time consuming participating in the movement became is that key players stopped showing up. Well not exactly; they still showed up, but mostly for side conversations, informal gatherings, and the meetings that planned what would happen at the public meetings. Using social media and social capital, text messaging and chat software, they formed an invisible guiding hand that simultaneously got shit done, avoided accountability, and engaged in factional battles with each other.

This isn't really very different than how powerful elites operate in the real world. But in the real world, leaders are less likely to talk about transparency and horizontalism, and more likely to have to stand for election, hold a title, or at least be subject to being written about. You know what's worse than regular same-old elites? A barely visible elite that denies it is an elite and can't ever be called to account.⁴²

To rescue non-hierarchical groups from leaderlessness, it's necessary to rescue thought leadership and protocol from hierarchy. Skilled charismatic leaders, such as Martin Luther King, Jr. or Mohandas Gandhi, can be of tremendous value to a movement by providing a unified vision and set of values around which people can organize. However, they are also a liability: they concentrate the power of influence into one fallible person, limit the opportunities for new ideas, and can be a “single point of failure” which can end the organization if the leader is arrested, deposed, or killed. But charismatic leadership is not the only way.

Ella Baker, one of the founders of the Student Nonviolent Coordinating Committee (a highly influential civil rights organization founded in 1960) argued against the leadership model expressed by the “prevailing messianic style of the period”⁴³, in favor of a model of leadership that focused on developing capacities and skills in others. Baker described this as “group-centered leadership.”⁴⁴ Instead of expressing leadership through a strong public personality with a single dominating vision, a group-centric leader helps to develop the capacities of other group members as leaders. Group-centered leaders work with groups to develop *group* values and *group* visions, by elevating each member of the group. Later groups formalized this approach into a model for consensus-oriented decision making, and

⁴² clenchner. The 7 sins of occupy, September 2012. URL <http://www.dailykos.com/story/2012/09/20/1134321/-The-7-Sins-of-Occupy>

⁴³ Mumia Abu-Jamal. *We Want Freedom: A Life in the Black Panther Party*. South End Press, January 2004. ISBN 9780896087187

⁴⁴ Barbara Ransby. *Ella Baker And The Black Freedom Movement: A Radical Democratic Vision*. Univ of North Carolina Press, 2003. ISBN 9780807862704

formed what has been described as "leaderless" organizations – organizations that have no single elected representative or charismatic figurehead.

However, this use of the term "leaderless" is challenging, and has foiled people both inside and outside of leaderless organizations, as it makes it easy to conflate decision-making power (imperative direction of a group) with the power to share one's visions, ideas, and values (thought leadership). That mistake leads toward covert leadership or group ineffectiveness. If we think of leadership as a zero-sum commodity, where one person's expression of leadership limits others', we can fall into the trap of stymying any efforts by group members to express their skill and vision. Under a group-centric model, leadership is not a zero-sum game: more leadership results in a more powerful and more capable organization, both on the individual and group level. Rather than answering the question "Who is the leader here?" with "no one"; a group-centric model of leadership allows a group to instead answer "everyone"⁴⁵; an answer which serves as a reminder of the importance of cultivating leadership skill within the group.⁴⁶

Protocol and subversion

The role of "directing group action" within a group can be decentralized by the same mechanisms for decentralization discussed above: the role can be *replicated* (separated into multiple groups, each directed by a single director, each pursuing the same function), *federated* (a single director replaced by a protocol for group direction), or *subsidiaritized* (different functional aspects of a group separated into subgroups, as with committees or working groups). Contemporary non-hierarchical groups regularly use all three strategies, but key among them is federation through the development of protocols for group interaction (the "process" or "rules of order" for a meeting).

Protocols are any procedural scripts that govern interaction, whether explicitly stated or implicitly understood. The procedure for greeting (e.g. "Hi, how are you?", "I'm well thanks. And you?") can be described as a social protocol; similarly, the hand signals used to indicate approval or disapproval in Occupy-style general assemblies are a protocol. Shared protocols reduce the need for reinterpretation and renegotiation of processes with each interaction, allowing interactions to become more efficient. Where well designed and implemented, these algorithms can help to ensure that no single member of the group wields undue power, while still facilitating effective decision making. Where poorly designed and implemented, the process can result in ineffective meetings and inequity.⁴⁷

⁴⁵ More discussion about leaderlessness vs leaderfulness within Occupy: <http://thefutureofoccupy.org/2012/02/is-occupy-leaderless-or-leaderfull-insights-from-working-groups/>

⁴⁶ The free-software enthusiasts might find an analogy with the overloaded meanings of the word *free*, which English defines both as *liberated* and *zero cost*. Just as free-software emphasizes "free as in speech" over "free as in beer", non-hierarchically inclined groups might emphasize "leadership as in organizers" over "leadership as in bosses".

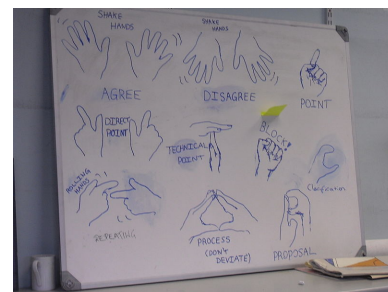


Figure 1.7: Some of the hand signals widely used in consensus-oriented meetings.

⁴⁷ See Galloway [2004] for an in-depth analysis of the relationship between protocol, control, and decentralization.

The development of protocols can be described as the process of finding modes of interaction that are flexible, reusable, and composable enough that they reduce the labor of developing a more circumstantially specific mode of interaction on the spot. Raising one's hand to speak, for example, is a protocological innovation that can improve the function of group communication (wherever groups are too large for free-form conversation, but small enough that you can get through a speaker's list in a timely fashion) – and its efficacy is demonstrated by its almost universal application across a wide variety of domains. The Internet Protocol, which defines the basic communication protocol at the network layer of the Internet, has been an astonishingly successful protocol in enabling an extraordinary breadth of systems which are relieved by virtue of the protocol from redesigning their own networking strategy each time. Groups engaging in decentralized processes are faced with this same task: by what protocols could people communicate, which enable efficiency and effectiveness, but also preserve equality? Chapter 2 will discuss one answer to this, consensus decision making, in much greater detail, though many other protocols and organizing principles are possible. The motivation for the design of InterTwinkles, discussed in Chapter 4, is to explore the design of protocols for online decision making.

Protocols may be enforced algorithmically by computers, as in the case of networked or software protocols, or by the collective will of participants, as in an in-person meeting. They may be explicitly stated or written down as bylaws and policy, or they may be artifacts of the cultural attitudes of members of the group that are left unstated. But regardless of the form or medium, any protocol can be subverted. A protocol may afford a particular purpose, but this *telos* can never be entirely circumscribed by the protocol itself. It remains incumbent upon people using systems to reflect upon and critique the use and subversion of their protocols, in order to continue to further the purposes for which they were adopted. This is why it is necessary for systems that wish to distribute control to consider the functional dimension of vision (the way meanings and strategy are created within an organization), as well as the means for reflection available to participants who are accessing or producing work within a system.

Alternatively, groups wishing to overthrow a protocological regime may find inspiration by considering ways that the ostensible purpose of a protocol could be subverted from within. Table 1.2 shows some examples of protocological subversion.




Figure 1.8: Flocking, the organizing protocol used by the loosely associated network of online activists known as Anonymous, exemplifies a protocol that is not explicitly stated, but rather a function of members' cultural understanding of the meaning of the group.

<i>Protocol</i>	<i>Affords</i>	<i>Subverted by</i>
Alphabet	Universal legibility	Encryption
Conversational norms	Verbal exchange	Chanting / Protest
TCP/IP	Computer message transfer	Denial of Service attack
BitTorrent	File transfer	"Torrent Poisoning"[Cuevas et al., 2010]
Wiki editing	Knowledge aggregation	Wiki spam
Formal consensus	Decision making	Filibuster / belligerence
Guy Fawkes Mask	Reference to Anonymous culture	Dilution of brand

Table 1.2: Methods for the subversion of protocol

Conclusion

Decentralization is a popular subject, and rightly so, given the variety of revolutionary technologies and organizations that have been built at least in part using distributed systems. The language of decentralization, however, lacks precision, leading to the conflation of highly diverse systems under a fuzzy banner of horizontalism. We can achieve greater precision by recognizing the different mechanisms of decentralization (*replication, federation, and subsidiarity*), and by analyzing the functional dimensions of control (*access, evolution, and vision, viewed from the lenses of technical capability, rules and protocols, exclusive meanings, and structural inequalities*).

Non-hierarchical group decision making still functions with systems of control inherited from the protocols by which it operates, whether explicitly or implicitly introduced to the group. Well-designed and well-functioning protocols can increase the equity in a group, but they alone are no panacea, as any protocol can be subverted. Groups must retain systems of reflection both to retain efficacy and to meet their normative goals of shared leadership.

2 Consensus

On November 15, 2011, at 7:30pm Eastern Time, on the evening after their early morning eviction from Zuccotti park, over 2000 Occupy Wall St protesters returned from a day of marching to hold a General Assembly. Without amplification, the crowd laboriously repeated what each speaker said in three waves of “people’s mic”; short 3- and 4-word statements yelled by successive concentric rings of people stretching out more than a city block. One of the facilitators began by introducing members of the facilitation team (who had roles for handling procedural requests, maintaining a list of people who wished to speak, and taking notes). The facilitator then introduced the general assembly process and ideology:

We are a horizontal movement. The cops think that power looks like shouting orders. We do things different here. We use consensus process here. That means we create space to hear as many voices as possible, and seek decisions that are not just supported by a majority, but can be consented to by all. A lot of you have seen this before. We’re going to give a quick recap of some of the hand signals we use...

Over the course of an hour and a half, dozens of speakers arduously made announcements to the crowd, presenting logistical concerns about the camp’s demise, and edifying statements about the importance and meaning of the movement. Following the announcements, the crowd broke into 35 breakout groups for more nimble discussion, considering the prompt offered by the facilitators: “What are the decisions we as a community need to making right now? What are our creative strategies and where are our strengths in your groups?” Each group then discussed, and reported their thoughts back to the main group.¹

Along with tens of thousands of others, I watched this general assembly from afar through live streams broadcast by roving reporters with smart phones. I was amazed that a group of that scale could successfully meet and communicate using the techniques of consensus – techniques familiar to me from over a decade of participation in dozens of cooperatives, collectives, and affinity groups that used a consensus-based process for their governance. The largest successful



Figure 2.1: Frame from a live stream of the eviction-day general assembly, captured by reporter Tim Pool (as “The Other 99”). <http://www.ustream.tv/recorded/18544151>

¹ Minutes for the meeting: <http://web.archive.org/web/20130621192000/http://www.nycga.net/2011/11/nycga-minutes-11-15-2011/>

attempt to use consensus decision making which I had personally experienced was a group of 50 people, most of whom were seasoned practitioners of the form. I'd heard of groups as large as 100 people successfully using a consensus process when the group knew each other well and was guided by professional facilitators, but never an anonymous assembly of thousands.

An expression of consensus process more familiar to me is a typical meeting of the Board of Directors for Boston Community Cooperatives, a small non-profit housing co-op, on which I served for three years. Meetings unfold as follows: at the scheduled time, the six to nine board members gather in the living room of the non-profit's sole property. The facilitator is a single person, responsible for taking stack as well as guiding progress through the agenda, while a second person keeps minutes of the meeting. Responsibility for both tasks rotates throughout the group between meetings. The facilitator writes out the agenda on a small hand-held white board. The points in the agenda have been collected by the meeting's facilitator before the meeting using the group's mailing list and in-person discussions (most of the board members share residence in the organization's only co-op), though at the start of the meeting, members can request that additional items be added to agenda if there is time.

The meeting commences with a round of "check-ins" (introductions aren't necessary, as every member of the group knows each other well already). Each person in turn talks about how things are going in their life, especially anything contributing to their emotional state which could influence their participation in the meeting. Often, a humorous or thoughtful ice-breaker question (e.g. "What would your roller derby name be?") is used to help to loosen people up and inspire creative thinking. Following check-ins, the facilitator guides the group through the agenda. Each point is discussed, with the facilitator keeping "stack" (an ordered list of people who want to speak). Sometimes, instead of simple discussion, the facilitator will suggest various structured brainstorming or deliberative processes more suitable to particular tasks. Minutes record any decisions made, as well as action items for next steps. Agenda items range from the mundane (e.g. report-backs on events members have attended) to the substantial (e.g. re-financing mortgages and working to acquire new houses). With the agenda complete, the group "goes around" once more to "check out" – each person in turn reflects on the meeting and what the group has (or hasn't) accomplished. While less spectacular than a large public assembly, this expression of small-group consensus constitutes the regular decision making practice of thousands of organizations around the world.²



Figure 2.2: The meeting space for Boston Community Cooperatives (the author served on their board for 3 years, and is on the far right).

² Considering housing alone, the Fellowship for Intentional Communities directory lists over 1338 housing communities that identify as using a "consensus" decision making process. <http://directory.ic.org>

While large assemblies such as Occupy's and a small co-op board meeting both call their process "consensus", and they share superficially similar techniques (the use of facilitators, the use of hand signals to communicate efficiently, the intentions for inclusion, language about "blocks"), comparing them is like comparing the floor of the U.S. House of Representatives to a small corporate board room, just because they both use Robert's Rules. The differences in scale, in membership, in intended outcomes, in buy-in, and in relationships between participants are vast enough to consider the two to be fundamentally different activities.

On any scale, the techniques of direct democracy are relatively unfamiliar to most people living in contemporary democratic societies, which emphasize large-scale representative structures and simplistic mechanisms like voting rather than inclusive deliberative processes. Even those who *have* experienced direct democracy through large-scale horizontal organizing or small-scale meetings may not have experienced both – leading to a tendency in critics and proponents of particular deliberative techniques to generalize their successes or failures too broadly. The term "consensus" within activist groups has come to represent not just a set of techniques, but a whole constellation of values, meanings, and movement affiliations that can cloud attempts to refine or adapt technique.

For designers interested in supporting consensus decision making practices online, this muddiness presents a challenge. New deliberative techniques must be learned, tried, and evaluated by groups to see if they're worth the trouble – but without narrowing down the form and constraints of the particular meaning of "consensus" one is designing for, one risks over-generalizing the lessons learned, or designing features that have the wrong scale and affordances. In this chapter, to lay the groundwork for this type of design, I will lay out a taxonomy of seven different meanings for "consensus" as a method for decision making, and identify the key differentiating factors for each. I will then identify some of the core values that motivate direct democratic approaches in general, which can help to form a foundation from which new designs for meeting technique can be built. I will zero in on what I call "affinity consensus" – the type of consensus decision making practiced by small affinity groups – and discuss its history, techniques, principles and challenges. This understanding forms the basis for the design of the InterTwinkles software platform, which will be discussed in chapter 3.

The Meaning of ‘Consensus’

Consensus is a loaded term – one that has a superficially simple meaning as the general of will of the group (as in, “the consensus of the group is to go out for Chinese food”), but a much more complex set of meanings in use as a decision making *method*, as in “we do our decision making *by* consensus.” Different subcultures refer to highly different practices with this word – practices which are meaningfully different enough that we might prefer to have different words to describe them. Despite this, debates about the efficacy or value of “consensus” as a decision making method usually involve a wrangling for authority about what consensus *is*, as though there were a single ideal type which monopolizes the correct usage of the term (e.g. ³, ⁴, ⁵, ⁶). Questions as basic as whether the decisions require unanimity (and what that means); who is allowed to participate; whether formally prescribed processes are required, forbidden, or optional; and whether decisions are binding all vary among different expressions of “consensus process”.

Before asking the questions ‘Is consensus worthwhile?’ or ‘What is a good consensus process?’ (or worse, ‘How do we design new consensus processes?’), we should start with the simpler question: ‘What does consensus refer to?’⁷ In my survey of consensus as a method of decision making, I’ve identified seven functionally different forms of consensus decision making, each of which has different enough processes, participants, and outcomes, that they are worth differentiating semantically. To do so, I’ve assigned each an appropriate adjective, which I will subsequently use when referring to different types.

1. **Corporate consensus** (or *nemawashi*) is the informal practice of building agreement and support across different parts of a formal power structure in order to avoid organizational conflict. When a manager talks about “building consensus” among their employees for decisions that they have already made (or which they could make regardless of whether they successfully built support), the manager is using this form. One expression of this is found in conflict-averse Japanese business culture, where the term *nemawashi* describes a ubiquitous form of persuasion. Back-room discussions at all levels of a company seek to head off any dissent before formal decision making meetings, so that the meetings can proceed as theatrical displays of unanimity. Corporate consensus is found where formal power differentials exist, and the achievement of buy-in is beneficial to the organization, but isn’t needed to determine the outcome of a decision. The power to make decisions

³ Howard Ryan. Blocking progress: Consensus decision making in the anti-nuclear movement, 1983. URL <http://www.connexions.org/CxLibrary/CX6187.htm>

⁴ Justine Alexandra Roberts Tunney. Occupiers! stop using consensus!, February 2013. URL <http://occupywallst.org/article/occupiers-stop-using-consensus/>

⁵ Randy Schutt. Consensus is not unanimity, December 2007. URL <http://www.vernalproject.org/papers/process/ConsensusNotUnanimity.pdf>

⁶ Seeds for Change. Consensus decision making. URL <http://seedsforchange.org.uk/consensus.pdf>

⁷ Rhetorically, I’m seeking to cover the same type of ground that Abend [2008] covers in his problematization of the meaning of ‘theory’ within sociology. That we both arrive at seven variant meanings of our respective terms is coincidental; that my rhetorical strategy follows his is tribute to the clarity of his writing.

is ultimately still in the hands of the management.

2. **Scientific consensus** is a preponderance of agreement among researchers of a particular field about the correctness of a particular theory (e.g. evolution, climate change, or the standard model of physics). As a *process*, scientific consensus is just the aggregate practice of science, with components like peer review – but there is no formal or universally accepted metric for when a scientific position becomes a “consensus position”⁸. Rhetorical usage of the phrase is largely based on the desire to back one’s argument with the authority of *science* where political contestation expresses uncertainty. In the absence of political contestation, one rarely talks about “scientific consensus”. In contemporary times, the claim that the Earth is round is not a matter of scientific consensus, but a matter of common sense. But if a sufficiently strong political faction of flat-earthers were to arise with the potential to impact policy, we would expect discussion of the scientific consensus around the roundness of the Earth to emerge.

3. **Standards consensus** is the process used by standards bodies such as the IETF, W3C, and IEEE⁹ to agree on specifications for interoperability in computer systems. The process is sometimes described as a “rough consensus” (the IETF appends “and running code”¹⁰). Stakeholders represent a roving cast of industry participants who hash out technical designs for protocols on the basis of technical and business needs (as well as ulterior motives such as getting more of one’s patents baked into a protocol to maximize royalties). Participation is not open to just anybody; participants must have demonstrated technical acuity or investment in the field to be allowed to participate. The consensus process is a deliberative process based on written documents and revisions of those documents; the final product is a tightly specified technical document. Where standards are contentious, processes might get bogged down and fail to complete, or to complete without the buy in of substantial parts of the community.¹¹ The adoption of standards is non-binding – no one is forced to implement a particular standard, except to the extent that they desire to remain interoperable with other implementers of the same standard. Competing standards can and often do operate in parallel.

4. **Consociationalism, or factional consensus**, is a form of representative governance where members of minority factions are guaranteed representation within a decision making body, which operates on the basis of unanimous consent of the representatives (that is, each representative is given the power to veto). This form of

⁸ Uri Shwed and Peter S. Bearman. The temporal structure of scientific consensus formation. *American Sociological Review*, 75(6):817–840, December 2010. ISSN 0003-1224, 1939-8271. DOI: 10.1177/0003122410388488. URL <http://asr.sagepub.com/content/75/6/817>

⁹ Acronyms: the Internet Engineering Task Force, World Wide Web Consortium, and Institute of Electrical and Electronic Engineers, respectively

¹⁰ Scott Bradner. IETF working group guidelines and procedures, September 1998. URL <https://tools.ietf.org/html/rfc2418>

¹¹ As an example, one of the early authors of the OAuth2 specification quit in protest and advises against its use; nevertheless, the standard was “approved” by IETF. [Hammer, 2012]

consensus is an answer to the particular challenge of governance where significant minorities have deep ideological differences with a majority or other minorities. Rather than risk quashing the minority with majoritarian proportional representation, each pre-defined faction is granted veto power, regardless of its size.¹² Consociationalism is a fundamentally representative, (rather than direct) democratic form; it is characterized by a highly structured, formal process between well defined factions.

¹² Sidney John Roderick Noel. *From power sharing to democracy: post-conflict institutions in ethnically divided societies*. McGill-Queen's Press - MQUP, 2005. ISBN 9780773529489

5. **Mob consensus** is the ephemeral and rough “consensus” achieved by unstructured groups that temporarily cohere into a particular unified action. Flocking behavior, the movement of unorganized protest marches, and the decision making process (or lack thereof) used by Anonymous are all examples of this form. Mob consensus is characterized by a lack of structure and a quickly changeable set of participants. A decision is ‘made’ when enough participants converge on a course of action to exhibit coherence. There is no power to veto or block; the only recourse for those who disagree with an action is to abstain from acting, and to persuade others to also abstain.
6. **Assembly consensus** is the form of mass consensus used by Occupy protesters. The assembly consists of anyone who decides to show up; but unlike mob consensus, the group uses a structured, formal process. The consensus process used in peer production projects like Wikipedia is a species of assembly consensus (we might call it “commons consensus”), where the group is primarily concerned with the management or production of a resource. Assembly consensus is characterized by an open membership, directly democratic, egalitarian process which is formally structured. The power to ‘block’ or ‘veto’ in assembly consensus often diminished as the size of the group increases – very large groups might use a fallback to a supermajority where deliberation fails to find a synthesis of views. In smaller groups, each member might have the power to block.
7. **Affinity consensus** is the form of small-scale group consensus used by cooperatives, collectives, and affinity groups. Different groups have different degrees of formal structure in their process, but each participant knows the others well, and membership is closed and relatively slow to change, as compared to assembly or mob consensus. This techniques used may be aesthetically similar to those of assembly consensus, but are functionally very different due to the closed membership, which allows the group to build personal relationships and trust over time. *Formal consensus*, a term

popularized by C.T Butler¹³, is one expression of this type of consensus; however, I use the term “affinity consensus” as a broader type which also includes less formal groups’ process, which Butler would exclude from his definition. In affinity consensus, each member of the group has the power to block decisions.

¹³ C. T. L. Butler and A. Rothstein. *On conflict and consensus: A handbook on formal consensus decisionmaking*. Food Not Bombs, 1991. URL http://theanarchistlibrary.org/library/C.T._Butler_and_Amy_Rothstein_On_Conflict_and_Consensus_a_handbook_on_Forma_Conensus_decisionmaking.pdf

Table 2.1 compares each of these seven forms with four distinguishing questions:

- Is membership open? That is, can anyone participate, or must one achieve some status within the community before participating?
- Is the process egalitarian? That is, will any participant’s opinion be given equal weight with respect to others’ opinions?
- Is the process formal? That is, does the process have well-defined rules and procedures?
- Are outcomes binding? That is, are you free to ignore an outcome with no consequence other than a loss of coordination (non-binding), or are there further consequences (binding)?

Consensus type	Open membership	Egalitarian	Formal process	Binding
Corporate	No	No (hierarchical)	No	Yes
Scientific	Yes	No (meritocratic)	No	No
Standards	No	No (meritocratic)	Yes	No
Consociationalism	No	Yes	Yes	Yes
Mob	Yes	Yes	No	No
Assembly	Yes	Yes	Yes	No
Affinity	No	Yes	Yes	Yes

Table 2.1: Seven types of consensus, and some of their distinguishing characteristics.

Among activists, the greatest source of confusion is in conflating *assembly* and *affinity* consensus (and occasionally *mob* consensus) – perhaps due to their shared historical roots. Additionally, for many activists, the term ‘consensus’ isn’t just a referent to the decision making practice – it signals identification with a larger constellation of meanings and values. In “Consensus: What it is, what it isn’t, where it comes from, and where it must go”, an essay reflecting on the use of consensus in activist movements, Andrew Cornell argues that the usage of consensus is not merely a description of a deliberative process:

Consensus functions as a synecdoche—a part rhetorically standing in for a greater whole. In this case, the whole that consensus stands

in for is a participatory, egalitarian, self-determining movement ... an autonym to a whole range of social institutions and ways of behaving.”
[Khatib et al., 2012, p 333]

By invoking “consensus”, one is claiming affinity with the civil rights movement, the feminist movement, the global justice movement of the 90’s and 2000’s, anti-war protests in the 2000’s, and alternative economic institutions like cooperatives and worker’s collectives. The term also signals an attitude towards direct democracy that is contrasted with representative structures like those typically found in organized labor unions. Critics of consensus within activism might seek to merely argue about the efficacy of particular meeting techniques, but find themselves arguing against the conflated weight of movement identity and ethics. Cornell posits that embedded within *consensus* are three distinct motivating beliefs – “an ethic of agreement, an ethic of participation, and an ethic of non-conflictual decision making”. He argues that practitioners should distinguish between these beliefs in order to avoid “muddled strategic thinking”, allowing more careful design of particular decision making strategies within groups.

Motivations for consensus

A brief history of affinity and assembly consensus

Multiple thoroughly researched histories of the use of consensus in activist movements have already been written¹⁴, so I will only provide a brief, U.S.-centric summary here.

While the traditions of consensus arguably trace back to pre-colonial tribal groups in North America and elsewhere, religious traditions in the early US stretching as far back as the 16th century (particularly the Quakers), and radical leftists in Europe in the 19th century, contemporary historians often begin discussions of current practices of consensus decision-making in activism with the Civil Rights movement. Popular images of the early Civil Rights movement were characterized by larger-than-life, “messianic” movement leaders who managed their organizations in a top-down, hierarchical fashion. Inspired to organize with more horizontal structures, groups such as the Student Non-violent Coordinating Committee formed and began to experiment with group-centric leadership and inclusive decision making. In the late 1960’s, the Students for a Democratic Society popularized the notion of “participatory democracy”, and flirted with (and eventually abandoned) consensus-oriented decision making practices, but nevertheless increased visibility of the idea of direct democratic organizing. In the 1970’s, feminist groups began

¹⁴ For a discussion of the development of the ethos of participatory democracy from the SNCC through the Direct Action Network, see “Freedom is an Endless Meeting” by Polletta [2004]. For a more succinct take with emphasis on the Movement for a New Society, see Andrew Cornell’s writing, particularly “Consensus: What it is...” in Khatib et al. [2012], but also “Oppose and Propose” [Cornell, 2011]. For a viewpoint more oriented towards assembly consensus, including discussion of the history of democracy in the US and before colonization, see “The Democracy Project” [Graeber, 2012]. For a viewpoint focusing on the roots in the Black Freedom Movement, see [Ransby, 2003]. For a discussion of horizontal organizing and its relationship to Marxist and anarchist traditions of leftist thought, see “The Horizontalists” [Marcus, 2012].

to critique the institutional oppression characteristic of centrally organized activist groups, and together with anti-nuclear activists, began to formalize meeting techniques drawn from Quaker traditions. In the late 1970's, the Movement for a New Society began conducting nation-wide trainings in consensus process, and popularized the idea of small affinity groups and collectives using intentional democratic structures. In the 1990's through the early 2000's, the Direct Action Network and other activists within the Global Justice movement further developed the organizing model of clustering independent, consensus-oriented affinity groups into a "spokes-council", a representative hub structure for coordinating independent groups. Rounding out this short history, the Occupy Wall Street encampment and ensuing movement in 2011 institutionalized the use of general assemblies and *assembly consensus*.¹⁵

Democracy and self-organization

If this dissertation so far reads as a survey of broad, fuzzy, and emotionally laden political terminology, allow me the liberty of adding one more: *democracy*. I don't wish to refer to the myriad layered meanings that come with contemporary "democratic countries" and "making the world safe for democracy", but rather the simpler root concept: democracy as the consultation of stakeholders for decisions that affect them. Whether it is invoked for practical or normative reasons, the value of consulting stakeholders is at the core of what consensus decision making is, even in its most watered-down formulations – and the way consensus approaches consultation distinguishes it from other democratic forms.

For most, "democracy" calls to mind "voting" and "majority rules". But in voluntary associations, majority voting is a less compelling means to coordinate action. Considered reductively, as a member of a minority, why would you ever go along with a majority with which you disagree? If the association is loose and voluntary enough, a minority opinion holder would just defect from the group. For that reason, consensus is the *de facto* operating principle for loose groups: everyone does what they want to, and coordinates to the extent that it is convenient and mutually agreeable. Where the majority has the power to compel you through the force of violence, "majority rules" makes more sense – you go along with the majority because they have the power to harm you if you don't. But coercion needn't be as obvious or overt as violence. In the more messy reality of human social interaction, relations that aren't overtly violent (including social status, friendship, the desire to be part of a group and a culture, hopes and dreams for a better future, and mutual de-

¹⁵ At the same time as these developments in the US, consensus and horizontalism grew in other parts of the world. In the 1990's, the Zapatista uprising in Oaxaca, Mexico formalized the use of *carajoles*, community political centers where members of a local municipality gather for local decision making, as well as coordination with neighboring communities. In 2001, Argentinian communities developed a movement for *horizontalidad*, taking over closed factories to repurpose them as democratic worker's collectives, with a strong emphasis on shared leadership and interpersonal relationships. This movement is documented in "Horizontalism: Voices of Popular Power in Argentina" [Sitirin, 2006].

pendence) can be powerful persuasive forces. There is a continuum between violent coercion and autonomous independence; the vast majority of human interactions are somewhere in between. Consensus as a formal set of decision making practices operates within that spectrum as an attempt to retain effective coordination, while minimizing coercion.

On practical grounds, consensus is a default in cases where there is little coercive force to bind people together (we see this expressed in *assembly consensus*, *mob consensus*, and *standards consensus*). David Graeber provides this example:

Imagine the city is about to destroy some cherished landmark and someone puts up posters calling for people to meet in a nearby square to organize against it. Fifty people show up. Someone says, okay, "I propose we all lay down in front of the bulldozers. Let's hold a vote." So 30 people raised their hands yes, and 20 people raise their hands no. Well, what possible reason is there that the 20 people who said no would somehow feel obliged to now go and lay in front of the bulldozers? These were just 50 strangers gathered in a square. Why should the opinions of a majority of a group of strangers oblige the minority to do anything—let alone something which will expose them to personal danger?¹⁶

¹⁶ David Graeber. Some remarks on consensus, February 2013. URL <http://occupywallst.org/article/david-graeber-some-remarks-consensus/>

Practically, a coordinating strategy which does not ask anyone to do anything they do not agree to is the only way to work with loose assemblies that have little history or reason to be bound together. But as a group grows and solidifies over time, gains power, and acquires resources, people might feel greater identification with the group and each other, and begin to feel a greater loss from defecting from the group (lost potential, lost hopes for the future, lost friendships, lost community, lost resources, etc.). The more people feel a deep personal connection to a group, the more the group has the power to coerce them to take a majority position with which they disagree. At this point, the group gains the potential to exercise coercion over its participants. Groups that wish to reduce that coercion must cultivate skillful strategies to arrive at mutually agreeable outcomes.

In the case of people's jobs and housing, there is immense potential for coercion. Bosses and landlords can make extraordinary demands, because the personal loss incurred by losing one's home or job is so great. *Affinity consensus* is used in groups such as worker's collectives and housing cooperatives as a means of reducing this coercive potential, while maintaining the group coordination necessary to function.

The Three Tyrannies

Writers about consensus processes often invoke one of three different “tyrannies” to describe coercive arrangements that the methods seek to overcome: the tyranny of the majority, the tyranny of the minority, and the tyranny of structurelessness.

One of the earliest uses of the phrase “tyranny of the majority” is found in John Adams’ 1794 “A Defence of the Constitutions of Government of the United States of America”. Adams identifies both the problem of unchecked minority power, and the “absurd” solution of giving minorities the power to veto as its only alternative, within a single-assembly government:

In a single sovereign assembly, each member, at the end of his year, is only responsible to his constituents ... not to the constituents of the minority who have been overborne, injured, and plundered... There is, in short, no possible way of defending the minority, in such a government, from the tyranny of the majority, but by giving the former a negative [(veto)] on the latter, the most absurd institution that ever took place among men... As the major may bear all possible relations of proportion to the minor part, it may be fifty-one against forty-nine in an assembly of an hundred, or it may be ninety-nine against one only: it becomes therefore necessary to give the negative to the minority, in all cases, though it be ever so small. Every member must possess it, or he can never be secure that himself and his constituents shall not be sacrificed by all the rest.[Adams, 1794, p 290-291]

Within Adams’ quote we can find both the “tyranny of the minority” and the “tyranny of the majority”. An unchecked majority which has the power to coerce the minority can trample it; at the same time, a minority which is given an unchecked power to veto the majority can result in obstructionism. In the looser forms of consensus like *mob consensus* or *scientific consensus*, the lack of any coercive power to bind any participant to the will of the group avoids the tyranny of the majority. But in more formal forms of consensus like *assembly* and *affinity*, groups usually have procedures for allowing participants to raise blocks against courses of action to which they disagree.

In order to avoid allowing blocks to become a tyranny of the minority, groups usually attempt to narrow their use through more nuanced instructions about when it is or isn’t appropriate to wield that power. For example, Occupy Boston used the following definition of a block (which facilitators read out loud, verbatim, every time there was a call for blocks in the process):

Blocks are an extreme measure, only used when a member feels a proposal endangers the organization or its participants, and/or the person feels so strongly about his/ her position that s/he will walk away from Occupy Boston.¹⁷

¹⁷ From <http://wiki.occupyboston.org/wiki/Facilitation%26DecisonMakingGuide>

Furthermore, after someone raised a block against a proposal, Occupy Boston would require the blocking participant to explain their reasoning, and then take a series of votes to see if the block was legitimate. First, the facilitators would ask the assembly if it believed that the reasoning correctly fit the definition of a block. If 75% of the assembly believed that it did, the facilitators would then ask if the assembly supported the block – and only if 10% of the assembly did so would the block be upheld.

The third tyranny that democratic groups face is the tyranny of structurelessness – a concept first advanced from the perspective of feminist activism by Freeman [1972], and corroborated in laboratory studies of groups by Bettenhausen and Murnighan [1985]. Groups never exist in isolation; they exist within larger social context which includes structural oppression such as racism, sexism, classism, and more. Where a group fails to pay attention to these dynamics, it will fall into patterns which replicate social dominance within the group, and thus become less democratic. Groups that wish to fight against this trend need to take time to reflect upon their process and culture, to develop individual consciousness around feminism, queer liberation, anti-racism, and other struggles against structural oppression. To institutionalize anti-oppressive practices, groups can adopt *structures* in the form of explicit procedures or protocols, as part of an effort to rewrite their habitual social scripts. Decision making processes that are compatible with formal structures (such as *assembly*, *affinity*, or *standards* consensus) can include mechanisms like rotation of facilitators, progressive stacks, and communication norms designed to counteract oppressive habits.¹⁸ Such protocols are not a replacement for individual consciousness-building, but rather a component of it: groups can develop consciously engineered habits to improve each member's awareness through structures that help it to constantly push their limits toward liberation.

These structures are by no means fool proof: as discussed in chapter one (page 41), formal protocols of any kind can be subverted – and a formal process of consensus is no exception. Groups still need to build in processes of reflection that allow them to adapt their protocols to support internal liberation. But even without overt subversion, groups will face an inherent tension between liberatory structures they seek to build and the institutionalized structures of oppression from the wider social context. Heteropatriarchy, white supremacy, imperialism, and other oppressive structures have been institutionalized at massive scales over centuries, and aren't easily phased out – groups wishing to resist these forces internally will need do so in active resistance to dominant external structures.

¹⁸ This topic is discussed in more detail in chapter 1 (page 38) with regard to non-hierarchicalism; the same applies to small-scale democracy.

Critiques of consensus

Consensus is not the only directly democratic form going, and despite its ubiquity in contemporary activism, is far from universally loved. The following are some of the major critiques of consensus as a fulfillment of democratic ideals. Most of these critiques speak specifically to particular forms of consensus; in each case I will identify the forms of concern in the critique.

Consensus as a force for conservatism

A facilitation trainer I spoke to during the course of my research recounted a story where she was working with a worker's collective which was struggling with their consensus process. The group had around 40 members, and relied on a consensus process to decide on changes to policy, new business strategies, and other regular decisions. However, some of the members had developed an interpretation where they felt that it was appropriate to block courses of action with which they weren't fully comfortable, or with which they had minor concerns. As a result, the group developed a highly risk averse business strategy, expressing a highly conservative attitude toward taking on new projects or changing policy.

Many consensus theorists consider the process to be based on preservation of the status quo until decisions modify it – the current policy, actions, and past decisions of the group are the “current consensus” until a new formal decision can be reached, which creates a new status quo. As long as there continue to be blocks to new decisions, the status quo remains in effect. Some trainers attempt to fight this trend by emphasizing that blocks are *extreme measures* that should only be used when one has a fundamental disagreement of principle with an action. Another strategy is to formalize a process of “consensus minus one”, or “consensus minus two”, where the group approves proposals even when one or two people “block” them.

Consensus selects for the most stubborn

A process that depends on the good will of participants to consider each others' interests can be undermined by obstinate participants who refuse to adapt their positions. An example from *commons consensus* (a subtype of *affinity consensus*) is in the fight over the Wikipedia article for the “Institute for Cultural Diplomacy”. In late 2012, the article was nominated for deletion due to potential violations of Wikipedia's policies on conflict of interest, advertising, neutrality, and notability – in short, the article was written by members of the ICD, read as an advertisement for the ICD, and seemed

to exist only to boost the reputation of the ICD. Multiple Wikipedia editors tried to improve the article, but faced resistance from ICD editors who wanted to maintain the organization's image. Ultimately, the article was nominated for deletion. Following its removal, the ICD threatened to sue Wikipedia editor Benjamin Mako Hill, who had been involved in the process. Hill wrote:

In Wikipedia, debates can be won by stamina. If you care more and argue longer, you will tend to get your way. The result, very often, is that individuals and organizations with a very strong interest in having Wikipedia say a particular thing tend to win out over other editors who just want the encyclopedia to be solid, neutral, and reliable. These less-committed editors simply have less at stake and their attention is more distributed.¹⁹

Following the lawsuit threat, Hill decided to steer clear of any further participation with the issue of ICD's article – he felt he had been intimidated out of the process. This case is an extreme example, but it demonstrates how the imbalance in attention from committed editors with ulterior motives can harm the never-ending deliberative process of the commons.

Outside of open-membership groups (such as *standards* and *assembly* consensus), stubbornness is less invisible, but can still be problematic. A stubborn member in a group practicing *affinity consensus* may cause visible harm that the group can work around, but if the group lacks the skill and tools to challenge or remove them, the group may still be stymied. The group may need to cultivate skill in negotiating with difficult personalities, or removing harmful members from the group, both of which are more difficult if the group has ethics of inclusion that are unmatched by skill in negotiating personality conflicts.

Consensus and the ineradicable diversity of moral belief

Consensus depends on group members sharing at least some fundamental principles and values – minimally, a group must share the value of remaining coordinated as a group; in most cases, groups must share a wider array of interests. In a piece critical of Occupy Wall Street, The Economist asserted that a consensus-oriented process was only possible because participants self-select:

Because the participatory democracy of OWS is an ideological endeavour, it can avoid the hard problem of liberal society: the ineradicable diversity of moral belief and the impossibility of consensus. Consensus-based communes composed of individuals who opt in specifically because they already agree with the commune's founding values can work precisely because the people who would make consensus

¹⁹ Benjamin Mako Hill. The institute for cultural diplomacy and wikipedia, March 2013. URL <http://mako.cc/copyrighteous/the-institute-for-cultural-diplomacy-and-wikipedia>

impossible—people with very different opinions and values—stay away.²⁰

A similar argument comes from Chantal Mouffe, who argues in *Deliberative Democracy or Agonistic Pluralism* that democratic theory needs to “acknowledge the ineradicability of antagonism and the impossibility of achieving a fully inclusive rational consensus.”²¹ Mouffe argues for a model which she describes as ‘agonistic pluralism’, where partisan factions place questions of power and antagonism at the center of their practice. Rather than striving to come to agreement with those with divergent interests, agonists conceive of those with whom they disagree as adversaries whose ideas are to be openly combated.

A formal model of these concerns is found in the game theoretical subdiscipline “Cooperative Game Theory”, which generalizes the search for Nash equilibria by examining the effects of coordination within coalitions of subgroups of players. This field seeks to find models of collective behavior based on a concept of immovable differences in belief.²²

Critiques of consensus from the ineradicable diversity of moral belief mostly apply to models of consensus that work at large scales, but which also have binding outcomes. Non-binding models of consensus such as *standards consensus*, *scientific consensus*, or *mob consensus* aren’t impacted, as partisans who disagree can simply “agree to disagree” and move on. In binding models (such as *affinity consensus*), the self-selection of participants works as an effective means of protecting the group from wildly divergent principles (though it does limit the model’s application to problems involving national-scale governance). Large groups with binding outcomes will be the most impacted – when an *assembly consensus* group crosses a threshold from a loose affiliation to a group with deeper personal investment, if it has not done the work to define core principles which identify the group (and thus exclude those who disagree), it will face difficulty.

Consensus is fragile

The plethora of accounts of poor consensus process with negative group outcomes, and the “no true Scotsman” caliber claims of consensus proponents that negative outcomes were a result of a misapplication of the form, leads to the critique that consensus is an inherently fragile process which is difficult for groups to apply correctly. Consensus trainers emphatically warn practitioners against what they describe as misconceptions about the meaning of consensus. Each of the following might seem to be a common-sense

²⁰ The Economist. Leaderless, consensus-based participatory democracy and its discontents, October 2011. URL <http://www.economist.com/blogs/democracyinamerica/2011/10/occupy-wall-street-3>

²¹ Chantal Mouffe. *Deliberative democracy or agonistic pluralism*, 2000

²² Thomas Ferguson. *Game theory*. URL http://www.math.ucla.edu/~tom/Game_Theory/Contents.html

understanding of the meaning of consensus, but each is considered wrong by trainers:

- Consensus is unanimity²³, or compromise, or analogous to voting with a 100% threshold.²⁴
- Blocking is a “no” vote, and signals the end of discussion.²⁵
- Consensus is a set of procedures.²⁶
- Consensus works without formal structure.²⁷

Consensus proponent David Graeber argues that one of the challenges in the contemporary exercise of direct democracy is that “Americans, for all their democratic spirit, mostly [have] absolutely no experience of democratic deliberation.” [Graeber, 2012, p 194] If his view is correct, we might find that the challenges groups have with their deliberative processes are the growing pains of a nascent culture of participatory democracy. But it may also be that consensus as conceived by groups attempting to implement it is a fundamentally fragile form of direct democracy, which is difficult to teach and difficult to “get right”. We might benefit from the development of less fragile protocols.

Consensus and scarcity

Consensus is fundamentally a process for negotiating the use of a centralized, scarce resource (whether that resource is a property, or just the capacity for coordinated action in a group). But as the discussion in chapter 1 illustrates, we could decentralize this resource and be free from any need to negotiate around its use. This decentralization could occur through replication, federation, or subsidiarity – but in any case, there would no longer be a need to seek agreement; each participant would be empowered to do what they wish with their corner of the resource. A post-scarcity housing cooperative might have a separate house for each member that wants it (with post-scarcity labor maintaining it); a post-scarcity worker’s collective would allow each worker to do whatever work they liked and none that they disliked; a post-scarcity General Assembly would eliminate groups’ limits in energy, time, and collective will as limits to effective political action. Rather than striving toward forms of deliberation that can coordinate use of a scarce resource (of whatever kind), we could devote attention to eliminating that resource’s scarcity, and therefore free ourselves from any need to coordinate.

This argument is rather fanciful and speculative; it is likely that we will need to do substantial negotiation around resources (whether friendship, love, housing, work, or protest) for a good long time

²³ Randy Schutt. Consensus is not unanimity, December 2007. URL <http://www.vernalproject.org/papers/process/ConsensusNotUnanimity.pdf>

²⁴ Seeds for Change. Consensus decision making. URL <http://seedsforchange.org.uk/consensus.pdf>

²⁵ Tree Bressen. The top 10 most common mistakes in consensus process and how to avoid them, October 2011b. URL <http://treegroup.info/topics/Top-10-Consensus-Mistakes.pdf>

²⁶ David Graeber. Some remarks on consensus, February 2013. URL <http://occupywallst.org/article/david-graeber-some-remarks-consensus/>

²⁷ C. T. L. Butler and A. Rothstein. *On conflict and consensus: A handbook on formal consensus decisionmaking*. Food Not Bombs, 1991. URL http://theanarchistlibrary.org/library/C.T._Butler_and_Amy_Rothstein__On_Conflict_and_Consensus__a_handbook_on_Forma_Conensus_decisionmaking.pdf

to come. But the post-scarcity viewpoint can help us to question whether we really need as much deliberation, and to consider the alternative of eliminating scarcity where feasible.

Consensus as protocol

Non-hierarchical organizing works by centralizing some coordinating function (whether that's a shared vision, a shared identity, or a protocol) in place of a leader who would direct the action of a group (see chapter 1, page 23). The techniques used by groups for consensus decision making can be described as a collection of protocols which can help them to coordinate their activity, develop ideas, and increase trust and understanding in one another, while maintaining values of inclusiveness, equality, and non-coercion. By formalizing these techniques, one can reduce the need for groups to perpetually re-invent their practice, and increase overall efficiency. Techniques which prove effective can be shared among different groups, allowing progress in the cultural skill of communication and organizing.

Consensus techniques are not the only efforts to protocologically increase the efficiency of group deliberation. One body of literature considers deliberation framed as “problem solving”; generally conceived for small groups engaged in tasks that do not involve contestation or negotiation. See “Techniques for Structured Problem Solving”²⁸ for an excellent and thorough catalog of this approach.

Robert's Rules of Order²⁹ is another particularly well known formalization which was developed for parliamentary debate (see page 151 for an abbreviated summary of the rules). Robert's Rules are wildly popular – and are “officially” used by 85% of all organizations in the United States³⁰, though most organizations only loosely hew to its prescriptions. Many organizations choose Robert's Rules when constructing their bylaws because they feel compelled to specify a formal process; many fewer actually hold themselves to it when conducting meetings. Groups may adopt “consensus” for similar reasons – among the cooperative houses surveyed for the InterTwinkles study, each claims consensus as its decision making process, even if they haven't taken time to consider the rationale, and use the process with varying degrees of formality. Unlike Robert's Rules, however, consensus is far less precisely defined – there is no single manual which defines a definitive “consensus process”, but rather a whole constellation of techniques.

One of the challenges in adapting Robert's Rules to different group and meeting contexts is the low degree of *composability* of its protocols: the set of rules stands as a monolithic operations manual for a deliberative body. By contrast, the techniques of consensus are

²⁸ Arthur B. Vangundy. *Techniques of Structured Problem Solving*. Springer, April 1988. ISBN 9780442288471

²⁹ Henry M Robert and Sarah Corbin Robert. *Robert's rules of order newly revised*. Da Capo Press, Philadelphia, Pa., 2011. ISBN 9780306820205 030682020X

³⁰ Jim Slaughter. *Parliamentary procedure in the 21st century*

atomic and composable enough that groups can appropriate those that they find valuable for their process, ignoring others. This lends consensus its immediate aesthetic recognizability (e.g. groups visibly using hand signals with a facilitator taking stack), even though the form, constraints, and process used by any particular group might differ drastically.

Techniques of Affinity Consensus

In the design of InterTwinkles, I have taken an approach of developing structured communication tools inspired by the techniques of consensus (particularly those drawn from *affinity consensus*). Each tool is intentionally designed to execute a particular composable component of a larger consensus-oriented process, without constraining the full form. To round out this discussion of consensus, I will now describe several of the key techniques. A full catalog of the techniques listed in the sidebar, along with instructions and recommendations for use, can be found in the appendix (page 155).

Facilitation: a facilitator helps to direct the process of discussion in a meeting. At its most basic, facilitation might include simply maintaining a stack (keeping track of who wishes to speak, and calling on them sequentially); though facilitators are often tasked with maintaining the agenda, keeping time, summarizing discussion, choosing structured discussion formats (such as go-arounds, fishbowls, world cafes, etc.), counting votes, drawing out people who are participating less, and keeping the group on-topic and on-task. Good facilitation is a skill, learned both through formal training and experience in meetings.

Ideally, facilitators will be neutral participants in a discussion – given the power they have to frame and direct discussion, biased facilitators can undermine the values of non-hierarchicalism in a group. For this reason, many groups will actively seek an outside facilitator when tackling issues that are deeply divisive within the group, when none of the group’s own members could be considered sufficiently impartial. Even for regular non-contentious work, most groups rotate the role of facilitation among different members of the group to distribute this power (as well as the practical skill of facilitation) more equally.

During the course of my research, I sat in on a meeting of a maker space (a member-owned and operated shared workshop), where the group was considering a major reorganization of their policies. Aside from myself, all of the meeting participants were members of the space. The facilitator of the meeting was deeply invested in the topic under discussion, and for the first 45 minutes of the meeting,

A short list of some of the widely used techniques in consensus-oriented meetings. Detailed descriptions of how each of these techniques works can be found in the appendix (page 155)

Meeting Phases:

- Orientation
- Check-in
- Announcements
- Agenda
- Breaks
- Check-out

Facilitation tools:

- Stacks
- Progressive stacks
- Clarifying questions
- Points of process
- Direct responses
- Proposals
- Friendly amendments
- Straw polls
- Discussion summaries
- Tabling
- Quorum

Hand signals:

- Twinkles
- Approval / disapproval
- Block
- Wrap-it-up
- New proposal
- Point of information
- Direct response

Roles:

- Facilitators
- Note takers
- Timekeepers
- Vibes watchers
- Shepherds
- Buddies

Formats:

- Icebreakers / fire starters
- Open stack
- Go-around
- Popcorn
- World cafe
- Fishbowl
- Dotstorm
- Spectrogram
- Solipsist’s meeting
- Neighbor interviews
- Break-out groups

was developing an unhealthy pattern of responding with a counter-argument to every comment that others in the group made, as well as frequently interrupting people to argue about what they had just said. When I noticed this trend, I offered to “keep stack” in the meeting. After some initial resistance to the idea from the facilitator, others in the meeting agreed. By enforcing the protocol that each person would raise their hand to speak and only speak when called on, the interrupting and back-and-forth arguing was preempted, and progress in the meeting accelerated considerably. Following the meeting, I overheard members of the group comment about the drastic difference in productivity before and after my intervention.

Hand Signals: One of the most visible and recognizable techniques widely used in consensus decision making is a language of hand signals for non-verbal communication. Among some of the more widely used hand signals are:

- *Twinkles:* Groups use twinkles (after which this dissertation project is named), a signal derived from American Sign Language’s sign for “applause”, as a signal of approval which is less audibly disruptive than clapping or verbalizing (e.g. “hear, hear”). Within the Occupy movement, aiming one’s fingers downward when twinkling was used to indicate disapproval. Some other variant signals for approval include knocking in the air with one’s fist (derived from American Sign Language for “yes”) or snapping one’s fingers (still audible, but quieter than applause).
- *Speaking turn types:* When requesting a speaking turn from a facilitator, many groups use a set of symbols to indicate a “type” of speaking request, which might influence the stack order. A “clarifying question” is a request, without editorializing, that the last speaker clarify something about what they just said, indicated by a “C” shape with the hands or a question-mark sign with the index finger. A “point of information” indicates some non-editorial, factual point directly pertinent to the discussion, indicated by an index finger pointing straight up. A “point of process”, which could include a suggestion for how the group’s deliberation should proceed, is indicated by two hands held in a triangle shape.
- *Proposals:* A request to make a formal proposal might be indicated by a “P” shape with two hands, or some other variation. A symbol used in one creative group in which I participated (and which subsequently spread to three other groups when I shared it with them, who loved its absurd name and whimsy) is the “proposal llama”, with hands in a shadow-puppet animal head shape.

- *Wrap-it-up*: A request to finish discussion and move on (or for a long-winded speaker to finish their thoughts), indicated by making a rolling motion with both index fingers.
- *Votes*: When testing for consensus to a proposal, or calling for opinions on a straw poll (a non-binding vote gauging opinion), groups indicate their opinions with a variety of hand signals. “Thumbs up” might indicate approval, and “thumbs down” for a block or disapproval. Some groups get much more precise – one with which I worked uses a full “Likert scale” with a five position hand signal spanning strong agreement to strong disagreement: thumbs up, palm up, palm sideways, palm down, thumbs down. Groups also may have signals for abstention or “stand asides” (allowing the proposal to proceed, but neither blocking nor signaling approval).

The particular form and semantics of hand signals varies from group to group; teaching the group’s hand signals is a regular part of training for new group members.

Proposals: Not every task consensus-oriented groups engage in results in a proposal – sometimes, groups simply wish to brainstorm, generate ideas, or share information. But proposals are the primary currency of decision making work in consensus-oriented groups. A proposal typically consists of a succinct statement about a course of action, opinion of the group, or new rule. Groups adopt proposals after a “test for consensus”, which consists of the facilitator asking for a show of hand signals to indicate the degree of support for a proposal. If members of the group object strongly enough, the proposal is either reconsidered through further discussion and amendment, or dropped. Consensus-oriented groups have developed a variety of flowcharts formalizing their proposal process – the sidebar shows one of the most generic and popular, though several other variants are presented in the appendix.

Perhaps the most critical aspect of a proposal process is having a well-defined way to determine when the group has finished considering a particular option, and is ready to move on. In the absence of some ritual or signal that finalizes a proposal, it is never clear whether consensus has been reached or not.

Discussion formats: The most common discussion format is a standard “stack” format discussion, where each person raises their hand to speak and is called on in turn by the facilitator. But depending on the needs of a particular issue, a facilitator might suggest a host of other discussion formats. Here are several examples:

- *Go-arounds*: every member of the group speaks, in the order in which they’re seated, without interruption – e.g. “passing the

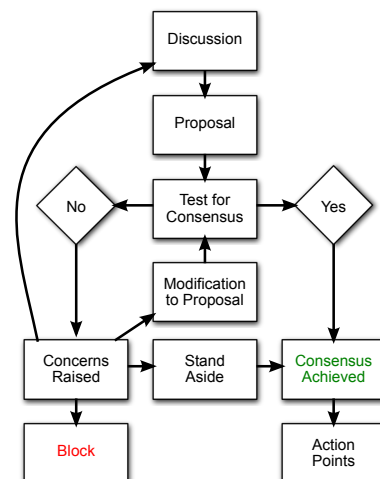


Figure 2.3: A generic proposal process flowchart. A larger rendering of this, along with several variants, is in the appendix on page 145.

conch”.

- *Popcorn*: the facilitator stops calling on people; instead, people just shout out short statements when they think of them.
- *Progressive stack*: people who speak the least, or who are the least empowered, are given preference in speaking order.
- *Fishbowl*: A subset of the group sits in a small circle in the middle of the room, and discusses an issue. The larger group listens without commenting. This is similar to a panel discussion, but where the members of the panel are directing their attention toward each other rather than the audience.
- *Breakout groups*: the group divides into smaller sub-groups for more nimble discussion. Small groups report a summary of their discussion to the larger group.
- *World cafe*: The facilitator sets up several stations around a room, each given a topic which is a component of a large problem or project, and places poster paper and markers at each. The group divides into a sub-group for each station. Each sub-group visits each station in a rotation, adding ideas to the poster paper.

Education: A critical component of the practice of any formal method of consensus is educating new members about the process. This becomes more critical the more quickly membership turns over in the group. At the Occupy Boston general assembly, facilitators would give an introduction to the hand signals, proposal process, and roles of facilitators at every meeting, so that any newcomer to the group could at least have gotten the gist of the process. In order to avoid a concentration of power in facilitators (or a loss of skill should a facilitator leave the group), groups must make an effort to train each member in facilitation practices. Since members of consensus-oriented groups will often be members of other groups which also practice consensus, an attitude towards popular education also enables effective techniques learned in one group to be shared with others. In this way, the composable techniques of consensus process that have the most effectiveness can be recombined with the processes and techniques already in use.

For those readers interested in more detail about the variety and range of meeting techniques, see the appendix (page 155).

Conclusion

The word “consensus” refers not just to a state of the group or an attitude of how to get there, but a whole constellation of meanings and

associations. There are at least seven functionally distinct types of decision making process which are referred to as “consensus”, which express differences in the openness of membership, egalitarianism, formality, and degree to which approved decisions have binding authority on the participants. Beyond these differences in type, within the domain of activism, the word also carries associational weight which connects a group to a history of horizontal activism stretching back to the 1950’s.

Despite this diversity in approach, there are a variety of well-defined, composable meeting techniques used within consensus-oriented groups, which can be regarded as protocols for efficient communication. Consensus-oriented groups (particularly those practicing *affinity consensus*) express skill in structuring their communication with types, symbols, and formats appropriate to various ideation, negotiation, and decision making tasks. By considering the function, intention, and appropriation of these protocols, we can gain inspiration for the design of decision making protocols for different media.

3 *Decision Making with Computers*

Man-computer symbiosis is an expected development in cooperative interaction between men and electronic computers. It will involve very close coupling between the human and the electronic members of the partnership. The main aims are 1) to let computers facilitate formulative thinking as they now facilitate the solution of formulated problems, and 2) to enable men and computers to cooperate in making decisions and controlling complex situations without inflexible dependence on predetermined programs. In the anticipated symbiotic relationship, men will set the goals, formulate the hypotheses, determine the criteria, and perform the evaluations. Computing machines will do the routinizable work that must be done to prepare the way for insights and decisions....¹

From the early days of digital computing, people have dreamed of ways that computers might contribute to human communication and decision making. In 1962, Douglas Engelbart described the way that computer systems might assist in collaboration within workgroups, kicking off the start of research into *computer supported cooperative work* (though that name wasn't coined until 1984).² In 1968, Licklider and Taylor explicitly explored the potential for computers as a tool for deliberation in groups³, presaging the development of Group Decision Support Systems (GDSS). From the first explicit GDSS study in 1982 through the 1990's, academic interest in GDSS resulted in a huge outpouring of research into the design and function of systems aimed at supporting and shaping decision making processes – though this research, despite its volume, was overwhelmingly narrow in focus, almost entirely consisting of single-shot studies with ad-hoc groups of undergraduate students in “decision support rooms”.⁴ As a result, the lasting contributions from the field are largely in framing, aspirations, and theory, rather than functioning systems or meaningful evaluations.

While InterTwinkles shares a heritage with GDSS as an effort to support group decision making with computers, it differs from most of that literature in the methodological approach to design and development (a participatory design process targeting real-world groups), the target user groups (real-world, non-hierarchical, consensus-

¹ J. C R Licklider. Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics*, HFE-1(1): 4–11, 1960. ISSN 0099-4561. DOI: 10.1109/THFE2.1960.4503259

² Douglas Engelbart. Augmenting human intellect: A conceptual framework. Technical report, October 1962. URL http://sloan.stanford.edu/mousesite/EngelbartPapers/B5_F18_ConceptFrameworkInd.html

³ Joseph CR Licklider and Robert W Taylor. The computer as a communication device. *Science and technology*, 76(2):1–3, 1968

⁴ Jerry Fjermestad and Starr Roxanne Hiltz. An assessment of group support systems experiment research: Methodology and results. *Journal of Management Information Systems*, 15(3):7–149, 1998. ISSN 07421222. URL <http://search.proquest.com/docview/218925443/abstract?accountid=12492>

oriented groups), and the technological context (web-based, mixed synchronous and asynchronous, dispersed). Nevertheless, it is valuable to consider the strategies that have been tried, and noting especially the lessons and failures from what has come before. In this chapter, I will provide an overview of current and past approaches to supporting group decision-making with computers. I will start with a survey of theoretical viewpoints concerning what constitutes “decision making” as an activity, and proceed with an overview of research methods appropriate for the study of groups. I will summarize prior work from GDSS, as well as the best related contemporary systems.

Foundations for the study of group decision making

Contemporary with early work in GDSS, researchers in social psychology developed comprehensive analytical frameworks to guide research into groups in general. One seminal work in particular, “Groups: Interaction and Performance” by McGrath [1984], proves especially valuable for its insightful representations of typologies of tasks, meetings, research methods, and group types. McGrath’s typology of tasks has been hugely influential, and was used to classify study tasks for almost all subsequent GDSS research. Equally valuable – but largely ignored by subsequent GDSS research – are McGrath’s typologies of groups and research methods. I hope you permit me the liberty of reproducing adaptations of no less than four of McGrath’s diagrams regarding group study, due to their exceptional clarity and perspicuity (figures 3.1, 3.3, 3.5, and 3.6). I will use these diagrams as anchor points for these three foundational questions:

- What is decision making?
- What factors influence decision making in groups?
- What are methods for the study and evaluation of group technology?

What is decision making?

Just as “voting” alone is inadequate as a description of democracy (it omits critical questions, not the least of which is “who selects the options?”), “choosing” alone is an inadequate description of decision making. Any decision will depend on myriad relationships, ideas, options, and constraints that face the group. In order to design

systems to support decision making, researchers have constructed models for how these activities function.

The first, simplistic model GDSS theorists use to answer this question is the *rational model*, which conceives of group decision making as a process where group members “optimize their choices on the basis of careful specification of the facts and refinement of their understanding of the probable consequences of their available options.”⁵ GDSS systems designed toward this model tend to be based around information retrieval and logical formalization. As Kraemer and King [1988] note with chagrin, this model “specifically excludes the baffling nonrational or quasi-rational behaviors individuals often exhibit.” While no rational person would believe that organizations or individuals actually operate under an exclusively “rational” model, some designers believe that they *ought* to, and design with the intention of nudging group participants toward more rationalized frameworks of decision making (e.g. ⁶ for a recent example). This model is strongly discordant with the prevailing literature on group facilitation, mediation, and consensus building, which emphasizes the importance of emotional understanding and relationship building among members of a group, rather than just rational information transfer.

A contrasting model at the opposite extreme is the *garbage can model*, which considers organizations as “collections of choices looking for problems, issues and feelings looking for decision situations in which they might be aired, solutions looking for issues to which they might be an answer, and decision makers looking for work.”⁷ This inversion of the relationship between “solution” and “problem” emphasizes the fact that organizations with explicit mechanisms for decision making can feel an urge to use them, and thus devote time and resources to this activity. Under the garbage can model, “decision support” would consist of providing more effective avenues for airing those grievances, emotions, and choices looking for problems that are the most pressing in a group, while striving to avoid burdening the group with the “attractive nuisance” of extra meetings and process.⁸

Another approach to characterizing decision making is to consider it as a collection of different *tasks*, each of which might have different goals. This approach is compatible with the garbage can model, but with an emphasis on the task rather than the agent seeking to perform it. This approach is particularly consonant with the design of formal structures for communication, as one can design new protocols to support particular task types. While numerous models of task types exist, the most pervasive is McGrath’s circumplex of tasks (figure 3.1). The typology divides group tasks into four different ma-

⁵ K.L. Kraemer and J.L. King. Computer-based systems for cooperative work and group decision making. *ACM Computing Surveys (CSUR)*, 20(2):115–146, 1988

⁶ Christopher Fry. Justify: A web-based tool for consensus decision making. February 2012. URL <http://web.media.mit.edu/~cfry/justify/Justify-IUI-Submitted.doc>

⁷ Michael D. Cohen, James G. March, and Johan P. Olsen. A garbage can model of organizational choice. *Administrative Science Quarterly*, 17(1): 1, March 1972. ISSN 00018392. DOI: 10.2307/2392088. URL <http://www.jstor.org/discover/10.2307/2392088?uid=3739696&uid=2&uid=4&uid=3739256&sid=21102485258521>

⁸ One anecdote from my research came from a long term cooperative organizer and facilitator, who recounted that his current residential co-op operates as a “do-ocracy”: anyone willing to do work gets to decide how that work is completed (for example, choosing colors to paint a room). Only when others in the co-op step up to share the work are they given a say in the matter. This approach reduced the burden of extra decision making, resulting in fewer “problems” to be solved through group decisions. The flip side of this is that in the absence of avenues for release, grievances and conflicts can percolate and build, causing disproportionate harm.

job types (*generative, choosing, negotiating, and executionary*), which are further divided across two axes: *conceptual vs behavioral*, and *conflict vs cooperation*. Conceptual tasks are primarily based in considering ideas and concepts, whereas behavioral tasks involve action. Cooperative tasks involve problems and work where every one's interests are aligned; conflictual tasks involve those where different participants have diverging interests. Between the four major types and two axes, we end up with eight task types:

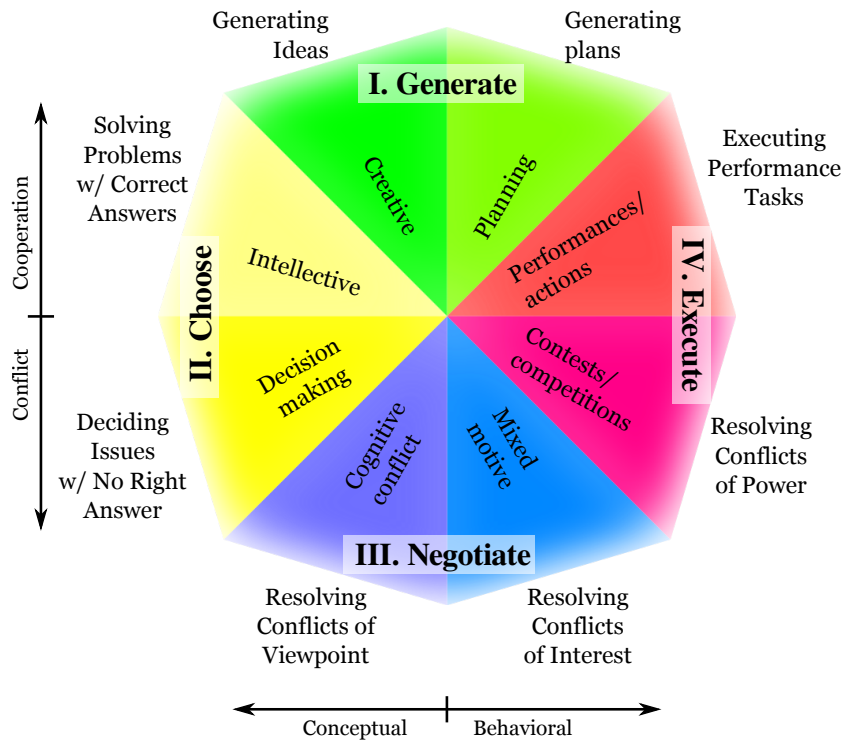


Figure 3.1: Joseph McGrath's circumplex of task types. Graphic adapted from original in [McGrath, 1984]

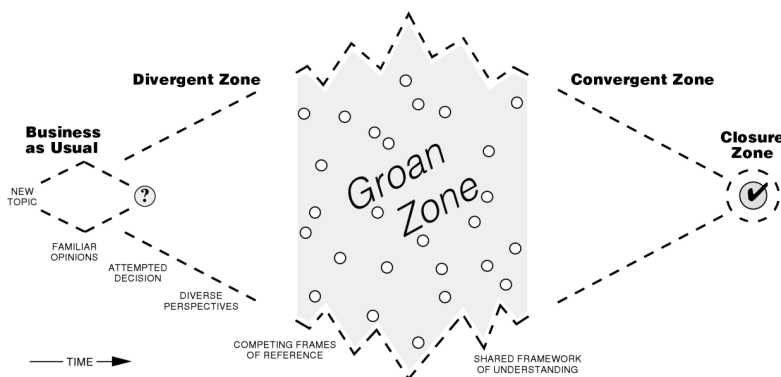
1. **Generating tasks:** These tasks involve coming up with new ideas; they are entirely cooperative (rather than conflictual), as they operate by accumulation rather than reduction of possibilities. "Conceptual" generation includes creative brainstorming, whereas "behavioral" generation involves generating possible plans.
2. **Choosing tasks:** These tasks involve decision making and calculation; they are entirely conceptual (rather than behavioral), as they only involve decisions rather than action. Cooperative choosing involves solving problems where there is an objectively right answer; conflictual choosing involves decision making where there is no objectively right answer.
3. **Negotiative tasks:** These tasks involve resolving differences be-

tween members of a group; they are entirely conflictual (rather than cooperative), as they are specifically about differences. Conceptual negotiation involves resolving conflicts of viewpoint, whereas behavioral negotiation involves resolving conflicts of interest.

4. Executionary tasks: These tasks involve performing and doing work; they are entirely behavioral (rather than conceptual), as they are specifically about action. Cooperative executionary tasks includes doing work together (e.g. stuffing envelopes); conflictual executionary tasks includes contests or competitions (e.g. wrestling matches).

This typology has been cited in hundreds of studies about group decision making, and is one of the key frameworks for differentiating concocted tasks imposed on groups by researchers studying group behavior with decision support systems. Particular group processes, such as the Nominal Group Technique⁹ (the group ideation technique on which the InterTwinkles tool “Dotstorm” is loosely based) might combine several task types – for example, a generating phase to collect as many ideas as possible, followed by a choosing phase to pick one or a few from among them.

However, one should be careful to avoid equating a *prescriptive process* such as the Nominal Group Technique with a *descriptive analysis* of what groups actually do. Figure 3.2 shows a common chart used in facilitation guides to describe the stages of decision making – first, a divergent stage with ideation, followed by a “groan zone” of uncertainty, and a convergent zone of decision making and closure.



⁹ André L. Delbecq and Andrew H. Van de Ven. A group process model for problem identification and program planning. *The Journal of Applied Behavioral Science*, 7(4):466–492, July 1971. ISSN 0021-8863, 1552-6879. DOI: 10.1177/002188637100700404. URL <http://jab.sagepub.com/content/7/4/466>

Figure 3.2: Group decision making stages from [Kaner, 2011]

Marshall Poole’s early research contradicted the claim that such “stage theories” were accurate descriptions of real-world group decision processes – he found that actual group processes were less

clearly defined, and slipped more readily between stages in a non-linear fashion.¹⁰ Anecdotally, the primary value I have found for such descriptive diagrams (as opposed to prescriptive diagrams such as the flowcharts in the appendix, which suggest a strategy for process) is to help calm anxious group members in moments of ambiguity, by providing vocabulary with which to name challenging moments in a meeting.

What factors influence decision making in groups?

There is more to a meeting than just the proximity of bodies in space and time. Meetings are impacted by the social structure of the group prior to their interaction (hierarchical or horizontal, homogeneous or heterogeneous membership, factions within the group, etc.), the characteristics of individuals (dominant and submissive personality traits, knowledge, skill, etc.), the task or situation (any of the quadrants in the typology above), the physical setting (the room, white boards, computer systems, etc.), and structural inequalities present in the wider social context.

Figure 3.3 shows the relationships between different elements contributing to a decision making process as mapped out by McGrath from 1984. Each of the entities, and each of the arrows connecting them, is a potential area for study – as Desanctis and Gallupe note in their “Foundation for the Study of Group Decision Support Systems”:

In examining decision processes and outcomes related to GDSS use, an astounding number of issues become worthy of study. Researchers must come to grips with what issues, among the many of relevance, they should address. Put another way, the question becomes, “what do we study?”¹¹

¹⁰ Gerardine DeSanctis, Marshall Poole, Ilze Zigurs, and other Associates. The minnesota GDSS research project: Group support systems, group processes, and outcomes. *Journal of the Association for Information Systems*, 9(10), October 2008. ISSN 1536-9323. URL <http://aisel.aisnet.org/jais/vol9/iss10/6>

¹¹ G. Desanctis and R.B. Gallupe. A foundation for the study of group decision support systems. *Management science*, pages 589–609, 1987

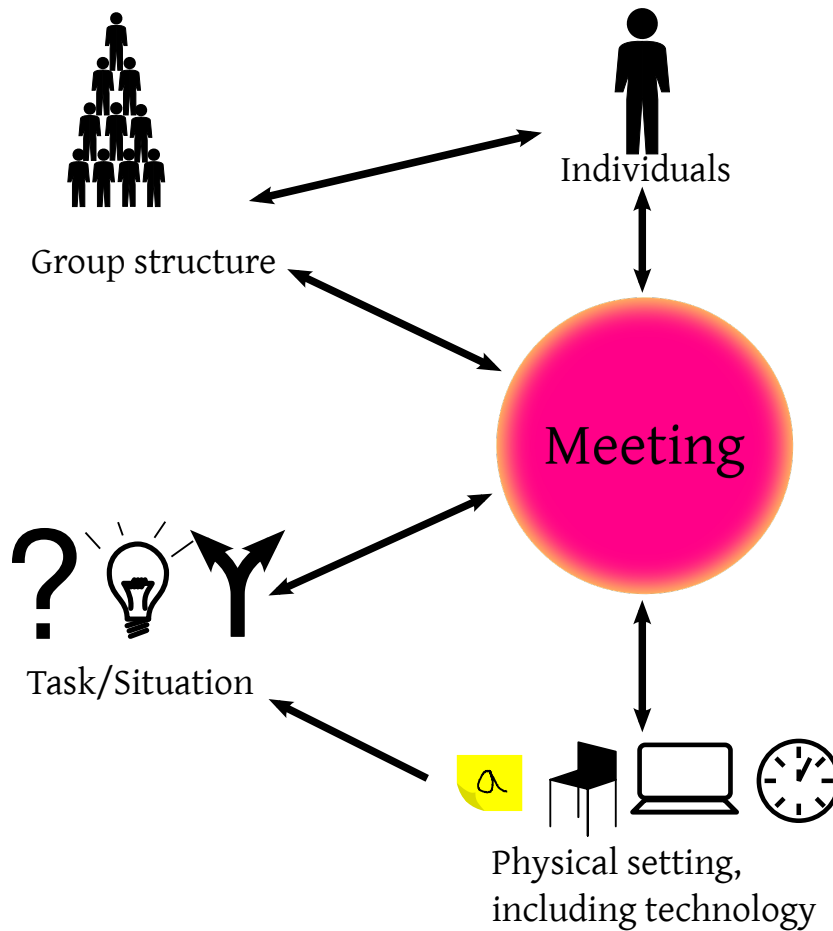


Figure 3.3: Contributing factors to the function of a meeting. Graphic adapted from original in [McGrath, 1984]. The original splits the "Meeting" into two distinct components (a "Behavior Setting" and an "Acting Group"). For simplicity and clarity I have combined them; their distinction is not necessary for the purposes of the analysis here.

This presents a plethora of potential avenues for study, but as described below (page 77) just a few of the study paths have been far more heavily trodden within the context of Group Decision Support research, leaving ample fertile ground for future research.

One example of such a research area of particular interest to democratic groups is considering how structural inequality interacts with meeting contexts. Figure 3.4 shows a map of different factors contributing to conflict in groups, including scarcity, urgency, ego, rigid outcomes, different backgrounds and cultures. This map was developed in the context of a workshop put on by the Anti-Oppression Research and Training Alliance, which considers group conflict from a lens of intersectional privilege and oppression. From this perspective, there is a more fluid continuum between factors that one would consider “individual” concerns, and factors that derive from structural concerns such as people’s conception of their roles in interactions with each other and the group. Technology designed for group use has the potential to contribute to changes in these structures – whether by amplifying inequalities or reducing them.

What are methods for the study and evaluation of group technology?

A study technology might seek to improve any of group outcomes, individual outcomes, or task outcomes. An improvement in group outcomes might include an increased sense of trust, understanding, or solidarity among group members, a greater sense of cohesion and capability, or a more positive image of the group among members. Individual outcomes could include improvements in the knowledge or skill of individual members of the group. Task outcomes might include more efficient or effective performance of particular tasks. In the history of GDSS research, various studies have sought to optimize each – and the instruments and methods used vary depending on the target. To study group outcomes, one might interview or survey group members to establish a baseline of some group characteristic, and repeat the test after use of the system (for example, [Carasik and Grantham \[1988\]](#) used a semantic differential scale, with self-reported numeric scores for whether members of the group felt the work team was “good” vs “bad”, “harmonious” vs “dissonant”, “strong” vs “weak”, etc.). To study individual outcomes, one might use similar assessments or aptitude tests directed at individual members. Assessing task performance is difficult in the case of complex real-world tasks that do not repeat; but by using laboratory tests with concocted tasks and a larger number of groups, one can establish statistically whether the GDSS meaningfully impacts group



Figure 3.4: A map of factors contributing to conflict in groups, including scarcity, urgency, ego, rigid outcomes, different backgrounds and cultures, fear, trauma, and more, taken from a training by the Anti-Oppression Research and Training Alliance, a group that provides training and support to consensus-oriented groups. All of these factors can influence meeting outcomes in real-world groups. Allied Media Conference, 2013.

performance for the particular task type studied.

Logistically, it is between difficult and impossible to arrange for natural groups (groups that exist outside the context of the research) to engage in controlled experiments, given the breadth of variables which can impact the outcomes (including variation in the group’s communication and decision making process reflexively impacted by participation in studies, which confounds iterative trials with the same group). For this reason, many researchers seek to eliminate variables arising from group structure by using concocted groups (for example, undergraduate students recruited to participate in a study). To further eliminate variation, researchers will usually impose concocted tasks rather than relying on highly variable natural tasks. Some standardized tasks include hidden profile tasks¹² and the desert survival problem¹³. Figure 3.5 shows a matrix of natural and extant groups, with natural and imposed tasks.

¹² Garold Stasser and Dennis Stewart. Discovery of hidden profiles by decision-making groups: Solving a problem versus making a judgment. *Journal of Personality and Social Psychology*, 63(3):426–434, 1992. ISSN 1939-1315(Electronic);0022-3514(Print). DOI: 10.1037/0022-3514.63.3.426

¹³ J. Clayton Lafferty, Patrick M. Eady, Alonzo William Pond, and Human Synergistics. *The Desert Survival Problem: A Group Decision Making Experience for Examining and Increasing Individual and Team Effectiveness: Manual*. Experimental Learning Methods, 1974

Composed Groups			Natural (Extant) Groups		Broad Band of Activities	"Natural" Tasks Endogenous to Group
Restricted interaction		Unrestricted interaction	Limited Term	Very Long-Term		
Restricted channels	Restricted process	Face to Face Ad Hoc				
		Natural Groups →	Expeditions (e.g. space crews)	Embedding systems (e.g. families)		
	Concocted Groups →	Systems tests (not using an existing crew)	Task forces (e.g. study commissions)	Standing crews (e.g. sports teams; work crews)	Limited Band of Activities	
Quasi Groups ↓		Mock studies (e.g. mock juries; artificial families)	Crew practice/training studies		Simulations	
Structured communication channels studies (e.g. restricted modality studies)	Restricted communication process studies (eg communication strategy training studies; brainstorming)	Ad hoc laboratory groups	Crew tests (not on system tasks)		Free Performance Form	Imposed Tasks
Highly stylized constrained task & communication (e.g. PDG; coalitions; communication Nets; Delphi)	Structured task & restricted communication process studies (e.g. NGT)	Structured task studies (e.g. SJT; MAUA; concept attainment task)			Restricted Response Form	

Figure 3.5: Matrix of natural vs concocted groups and tasks. Adapted from McGrath [1984]

It is not possible in any single study to maximize all possible study outcomes – that is, one cannot maximize both observations of the in-situ function of a system, and generalized observations of human behavior. Researchers may choose to engage in laboratory

experiments in contrived settings, field studies or experiments in natural systems, analysis of simulated or theoretical environments, or non-behavioral surveys of judgment or opinion. No choice is inherently more “correct” than any other; however, the epistemological outcomes of each will vary. A system designed for laboratory experimentation cannot be expected to reliably work in the field, but it will enable precise measurements of behavior. By contrast, a field study can result in lessons applicable to the diffusion of group decision support systems in real-world contexts, but will be beset by the particular configurations of the study context and generalize less easily. Rounding out our tour of McGrath diagrams, figure 3.6 presents a circumplex of different study types involving groups, which highlights the optimizations of each.

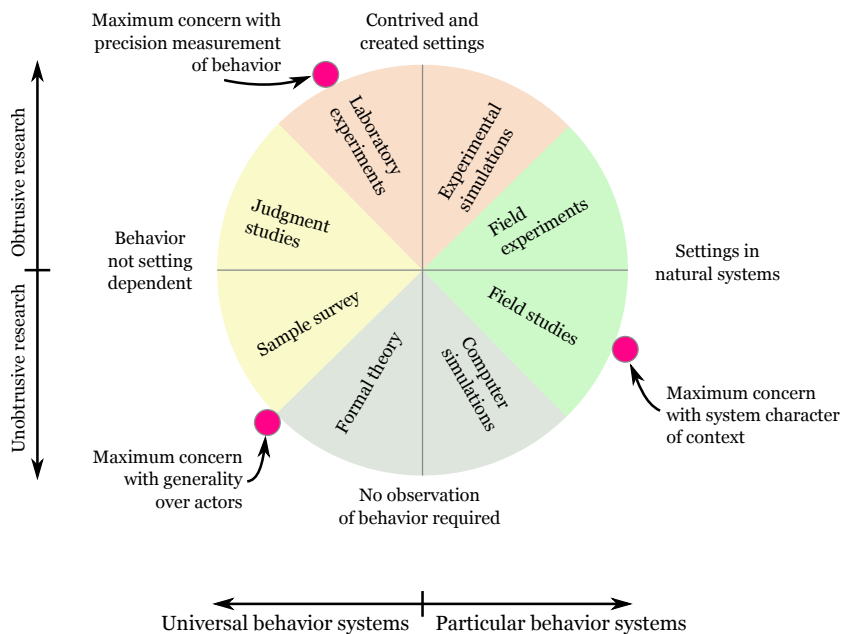


Figure 3.6: Methods for the study of groups. Adapted from [McGrath, 1984]

For researchers interested in design for groups which practice *affinity consensus*, field studies are the most appropriate approach, due to the particular constraints around which affinity consensus functions. Laboratory decision support studies with composed groups and imposed tasks do not investigate any effects caused by long-lasting relationships between group members, concerns of the group’s sense of identity and purpose, or issues of integrating a group decision support system into existing group constraints such as policies, politics, or mundane practical details of use context. Until it is proven robust and functional with those constraints, a GDSS for such groups has not been fully evaluated.

Once a system is in place and in wide use, it becomes possible to optimize it through controlled experimentation (e.g. A/B testing, where two variant interfaces are presented to different subsets of system users, and if one outperforms the other, it becomes the default for all users). This style of testing, however, is unlikely to provide generalizable results for other systems. As an example, Microsoft's "Bing" search engine famously performed A/B testing on the color of blue used for hyperlinks in the search results, and found one particular shade to result in substantially higher click rates.¹⁴ However, this result can't be taken to indicate a general preference for that shade of blue across all sites: it depends on the relationship of that color to the other design elements and colors in the page, the similarity or distinctiveness from colors used by competitors (particularly Google), and the wide scale practices and popularity of different colors in different contexts. To indicate general preferences for color (maximizing concern over actors), one would have to perform experiments across a wide variety of systems over time.

¹⁴ Ina Fried. Behind Bing's blue links, March 2010. URL http://news.cnet.com/8301-13860_3-20000623-56.html

After encountering challenging inconsistency in early experiments with GDSS – even in laboratory settings which sought to control as many variables as possible – GDSS researchers (particularly the research team at the University of Minnesota) were faced with a need to develop more refined analytical methods which adequately responded to the complexity of the systems they were studying. Introducing a novel technology to a group depends on the existing structures and personalities in the group, the characteristics of the technology, the context of training, and more – a level of research complexity more familiar to organizational sociologists than technologists.

Fittingly, Poole, Seibold and McPhee turned to sociology to adapt *structuration theory* as a method for analyzing organizational change following the introduction of a new technology. Structuration is an analytical lens developed by Anthony Giddens in [Giddens, 1993] (first published in 1976) and [Giddens, 1986] to help explain how it can be possible for social "structures" to persist, without resorting to classical objectivist ontologies of disembodied immovable forms. Giddens proposes that social structures persist only in the form of memory traces in individuals, but come into being when they are re-performed by agents who believe them to be the best course of action in any given moment.

Two notions are key to structuration theory: first, social structures are born of a "duality" – they are "both the medium and the outcome of the practices which constitute [social] systems."¹⁵ Structures come into being when people perform them; and by re-performing them, their persistence in peoples' understanding of appropriate

¹⁵ Anthony Giddens. *New Rules of Sociological Method: A Positive Critique of Interpretative Sociologies*. Stanford University Press, 1993. ISBN 9780804722261

social action changes. The second key notion is that agents are not feckless automatons driven by structures they cannot control – on the contrary, each individual actor in society is a de-facto social theorist, reasoning and willfully choosing how to act, and which structures to perform:

All competent members of society are vastly skilled in the practical accomplishments of social activities and are expert ‘sociologists’. The knowledge they possess is not incidental to the persistent patterning of social life but is integral to it.¹⁶

The result of this theory is a useful explanation for how social structures can come into being and become powerfully dominant (long-lived persistent structures become “institutions”) – when enough participants continuously re-perform a certain structure (such as family roles, boss-employer relationships, belief in the value of money, etc.), they develop understandings and expectations of how one ought to perform in these circumstances that reinforce scripts and habits they have previously learned. In particular, this can include structures of dominance, including heteronormativity, paternalism, sexism, classism, etc. But in addition to this, structuration theory points to a way out: since structures “exist” only in the form of memory traces and frequent re-performance, they can be changed by convincing agents to perform differently.

Sewell [1992] developed structuration theory further to enumerate some of the ways in which social structures can change; providing potential inspiration for activists interested in creating social change. Sewell identified five key mechanisms for change:

- *The multiplicity of structures*: Agents can choose between any of several different possible structures to perform in a given domain.
- *The transposability of schemas*: Schemas (memory traces of structures) from one domain can be applied to another domain, using analogical or metaphorical reasoning.
- *The unpredictability of resource accumulation*: Even deeply entrenched institutional structures don’t always result in the same material consequences – so over time, agents might find those structures less appealing.
- *The polysemy of resources*: The physical artifacts that result from the performance of structures can be interpreted in different ways by different agents. Technologies, tools, and systems can be appropriated to different ends.
- *The intersection of structures*: Where different structures overlap and intersect, agents might be able to move between roles within them.

¹⁶ Anthony Giddens. *The Constitution of Society: Outline of the Theory of Structuration*. University of California Press, January 1986. ISBN 9780520057289

Group communication technologies have a particularly strong relationship with social structures, as their design will enable certain types of interactions, and disallow others. Details such as the graphic design, the language used for different components and roles, and what types of communication are made easy or difficult can powerfully reinforce certain behaviors and discourage others. This does not mean that such technologies will instrumentally cause social change – the group must still choose to use the technology, and might still find ways to appropriate it contrary to the designer’s intention. From a structural perspective, a group technology must be analyzed both with respect to the “spirit” of the technology (the intentions of the designer and users’ interpretations of those intentions), the appropriation of the technology by the group, and any reflexive changes to the group brought on by performing the structures implied by the technology.

Poole and DeSanctis formalized this approach as *Adaptive Structuration Theory*, which has subsequently been used by numerous studies of organizational adoption of technology.¹⁷ AST provides an analytical lens for describing the reflexive interplay between technologies and social structures in organizations. The analysis might use ethnographic methods, surveys, or micro-sociological techniques like conversational analysis to analyze behavior of participants in a group, with a focus on the enacted structures that group members actually perform, rather than the technology’s inherent form as intended by its designers. The generalizability and precision of an AST analysis will be a function of the methods used and diversity of systems considered.

History of GDSS

In 1985, a team of researchers at the University of Minnesota (including Geraldine DeSanctis, Marshall Scott Poole, Ilze Zigurs, Brent Gallupe, and others) formed a research group to formalize the study of Group Decision Support Systems. For the next 20 years, this group dominated the field, producing dozens of articles, book chapters, and dissertations, and forming a foundational body of work.¹⁸ DeSanctis and Gallupe’s 1987 “Foundation for the Study of Group Decision Support Systems”¹⁹ provided a basis for the vast majority of subsequent research, delineating many of the primary concepts and organizational principles for the field.

One of the important contributions of this work is DeSanctis and Gallupe’s taxonomy of different GDSS types. They describe three “levels” of GDSS, which correspond to different levels of intervention into the function of a group: Level 1 consists of tools which

¹⁷ Geraldine DeSanctis, Marshall Poole, Ilze Zigurs, and other Associates. The Minnesota GDSS research project: Group support systems, group processes, and outcomes. *Journal of the Association for Information Systems*, 9(10), October 2008. ISSN 1536-9323. URL <http://aisel.aisnet.org/jais/vol9/iss10/6>

¹⁸ Geraldine DeSanctis, Marshall Poole, Ilze Zigurs, and other Associates. The Minnesota GDSS research project: Group support systems, group processes, and outcomes. *Journal of the Association for Information Systems*, 9(10), October 2008. ISSN 1536-9323. URL <http://aisel.aisnet.org/jais/vol9/iss10/6>

¹⁹ G. DeSanctis and R.B. Gallupe. A foundation for the study of group decision support systems. *Management science*, pages 589–609, 1987

introduce new channels or opportunities for communication (for example, video chat, shared documents, messaging systems, etc.); level 2 provides formal structures within a channel (for example, defining protocols by which people communicate for more efficient decision making, etc.); and level 3 performs reasoning or content-based contributions to the process (for example, automated just-in-time information retrieval, content-aware mediation and facilitation, etc.). Table 3.1 shows Desanctis and Gallupe's explication of these different levels of support, with example support features.

Task Purpose	Task Type	GDSS Level	Possible Support Features
GENERATE	Planning	Level 1	Large screen display, graphical aids
		Level 2	Planning tools (e.g., PERT)
	Creativity	Level 1	Anonymous input of ideas; pooling and display of ideas; search facilities to identify common ideas, eliminate duplicates
		Level 2	NGT, Brainstorming
CHOOSE	Intellective	Level 1	Data access and display; synthesis and display of rationales for choices;
		Level 2	Aids to finding the correct answer, e.g., forecasting models, multi-attribute utility models
	Preference	Level 1	Preference weighting and ranking with various schemes for determining the most favored alternative; voting schemes
		Level 2	Social judgment models; automated Delphi
		Level 3	Rule-based discussion emphasizing equal time to present opinion
NEGOTIATE	Cognitive conflict	Level 1	Summary and display of members' opinions
		Level 2	Using social judgment analysis (SJA), each member's judgments are analyzed by the system and then used as feedback to the individual member or the group
		Level 3	Automatic mediation; automate Robert's Rules
	Mixed Motive	Level 1	Voting solicitation and summary;
		Level 2	Stakeholder analysis
		Level 3	Rule base for controlling opinion expression; automatic mediation; automate Parliamentary procedure

Table 3.1: DeSanctis and Gallupe's (1987) Example GDSS features to support six task types, across 3 "levels" of GDSS. "Level 1" tools introduce new channels or opportunities for communication, "Level 2" tools structure channels for more effective process, and "Level 3" tools perform reasoning or content-based contributions to a process.

By this analysis, the successful GDSS tools which are now ubiquitous features of contemporary groups' toolkits are primarily level

1 tools, which provide new channels for communication – these are the “groupware” or collaborative productivity tools like shared document editing, video conferencing, email, etc. Tools which impose a particular process on group action, such as issue or ticket systems, calendaring systems, or question answer services, could qualify as “level 2”. There are fewer contemporary and historical examples of “level 3” systems – this, according to DeSanctis and Gallupe, is as it should be, as “level 2” systems are not yet well developed: “We propose that research into the design and use of GDSS should proceed in an iterative manner, beginning with Level 1 and Level 2 systems and advancing to the study of Level 3 systems after some understanding of the needed features and impacts of lower level systems has been achieved.”²⁰

Controlled experiments

Following the publication of DeSanctis and Gallupe’s *Foundation*, a veritable flood of studies of Group Decision Support Systems were published. Despite the quantity, the studies present a remarkable similarity in approach. In a meta-study of over 200 Group Support Systems experiments, Fjermestad and Hiltz [1998] describes the worrying paucity of breadth. Highlights include:

- 95% used ad-hoc groups composed for the study, with no prior history.
- 90% use students as subjects, mostly undergraduates.
- 73.5% employ groups for only one decision making task (thus offering no chance to develop familiarity and skill with using the system).
- There is a low degree of attention to details of the tools and procedures used: “most experiments seem (mistakenly) to assume that all GSS are a standard ‘package’ that will have the same effect.”
- 91% used face-to-face systems (primarily “decision support rooms”).
- Over half use one of three software packages in decision support rooms (monolithic packages developed by the University of Arizona, the University of Minnesota, and the New Jersey Institute of Technology), and little attention is paid to differences between features and interfaces.

Most remarkable from the results of this meta study is the apparent pervasiveness of an attitude among decision support researchers that

²⁰ G. Desanctis and R.B. Gallupe. A foundation for the study of group decision support systems. *Management science*, pages 589–609, 1987



Figure 3.7: A “Group Decision Support Room” at the International Scientific Research and Development Institute in Amsterdam, Netherlands. (2002)

the particular details of the software and task support are unimportant – this despite ample evidence that difficult-to-use interfaces are significant contributors to failures. This attitude is apparent in the studies’ tendency to describe their results not in terms of the efficacy of particular designs, methods, or implementations, but in terms of the efficacy of “Group Support Software” in general – a rhetorical style reminiscent of an overly-broad patent application. As a software designer who agonizes over design details as minute as the color of a button, this tendency is baffling. Furthermore, the designs employed were generally substandard, even for the era, as Fjermestad and Hiltz note:

Unfortunately, the scope of this body of work and its external validity/generalizability for “real” problem solving groups is weaker than would be hoped. And, despite the relative recency of this body of work, most of it was done on what is already “outmoded” technology, given the increasing persuasiveness of graphical user interfaces (GUI) such as Netscape™ and of hypertext and hypermedia systems embedded in the World Wide Web...²¹

The University of Minnesota’s own self-survey of their research group’s contribution to the field is far more charitable.²² It emphasizes the research value in using the same system longitudinally – specifically, the *Software Aided Meeting Management* decision support room, which researchers used both in the lab, and installed for studies with extant groups at the Internal Revenue Service and Texaco Corporation. To their credit, DeSanctis, Zigurs and others within the research program did publish studies which suggest methods for evaluating design elements of decision support systems in a more complete way (e.g. ²³ and ²⁴); nevertheless, their research program’s results remain (as they freely admit) vulnerable to the idiosyncrasies of the platform, and stuck in the room.

While Fjermestad and Hiltz recommend more diversity in GDSS studies, I remain skeptical that merely expanding experimental coverage into other corners of a classification matrix is sufficient. The combinatorial explosion of confounding factors in GDSS design simply grows too quickly to exhaustively study each characteristic using a method of controlled trials. In addition to task type and the three “levels” used to distinguish GDSS tool types, DeSanctis and Gallupe add distinctions of group scale (“larger” vs “smaller”), and proximity (“face-to-face” vs “dispersed”). Later analyses added temporal characteristics (“synchronous”, “asynchronous”, or “mixed”), and the degree of freedom offered for arbitrary communication (“restrictiveness”). Considering just these characteristics (and regarding “restrictiveness” as binary), we already have 192 configurations – nearly as many as all the GDSS experiments of the nineties.

²¹ Jerry Fjermestad and Starr Roxanne Hiltz. An assessment of group support systems experiment research: Methodology and results. *Journal of Management Information Systems*, 15 (3):7–149, 1998. ISSN 07421222. URL <http://search.proquest.com/docview/218925443/abstract?accountid=12492>

²² Gerardine DeSanctis, Marshall Poole, Ilze Zigurs, and other Associates. The minnesota GDSS research project: Group support systems, group processes, and outcomes. *Journal of the Association for Information Systems*, 9(10), October 2008. ISSN 1536-9323. URL <http://aisel.aisnet.org/jais/vol9/iss10/6>

²³ Gerardine DeSanctis, James R. Snyder, and Marshall Scott Poole. The meaning of the interface: A functional and holistic evaluation of a meeting software system. *Decision Support Systems*, 11(4):319–335, May 1994. ISSN 0167-9236. DOI: 10.1016/0167-9236(94)90079-5. URL <http://www.sciencedirect.com/science/article/pii/0167923694900795>

²⁴ Ilze Zigurs, Bonnie K. Buckland, James R. Connolly, and E. Vance Wilson. A test of task-technology fit theory for group support systems. *SIGMIS Database*, 30(3-4):34–50, September 1999. ISSN 0095-0033. DOI: 10.1145/344241.344244. URL <http://doi.acm.org/10.1145/344241.344244>

Now if we add differences in group type (ad-hoc vs natural), task endogeneity (imposed vs natural), group structure (hierarchical, factional, horizontal), facilitation style, training, and any of the other features implicated as factors in GDSS, we can quickly convince ourselves that a strategy of exhaustive analysis is unsound. Dividing GDSS study features into ever finer distinctions in order to explain significant variations simply repeats the fallacy of logical positivism, learned long ago by other complex fields: in the face of high complexity, one simply can't control all the variables. Instead, we need a different approach – such as design using non-positivistic methods (e.g. iterative, participatory, and auteur); testing with real-world groups; and iteration to optimize within the context of practicing groups.

Designs

As hinted above, despite the mountain of research into group decision support systems, it is surprisingly difficult to find accounts which provide design details of the systems under study. To round out the historical discussion of GDSS, I will discuss a few notable designs here.

The **Software Aided Meeting Management system (SAMM)** was developed at the University of Minnesota, and was the basis of the Minnesota research group's entire canon of research (including 29 refereed articles, 34 book chapters and proceedings publications, 11 doctoral dissertations). The system was designed for face-to-face, synchronous use in groups of 3 to 16 people. In the retrospective summary of their research program, the Minnesota researchers emphasize the "spirit" toward which the system was designed – the intentions and values that the designers hoped would be imparted on groups using the system. This included:

- (a) participatory decision making guided by rational discussion; (b) democratic, shared leadership; (c) efficient use of group resources; (d) confrontive, constructive conflict management; and (e) an informal, safe climate for the group.²⁵

Functionally, the system was a text-mode, menu-driven interface that provided mechanisms for defining agendas; brainstorming and gathering ideas; evaluating ideas through ratings, ranking, votes, and weights; performing stakeholder analysis, multi-criteria decision analysis and problem formulation; and recording thoughts, moods, scratch-pad ideas, and minutes. While the system began as a strictly lab-based research tool, researchers subsequently installed instances of the SAMM system in corporate environments for field trials.

The **Electronic Information Exchange System (EIES)** was a dis-

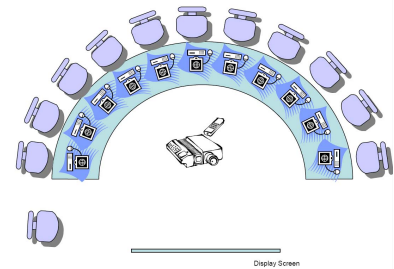


Figure 3.8: Graphic showing the setup for the SAMM decision support room. [DeSanctis et al., 2008]

²⁵ Gerardine DeSanctis, Marshall Poole, Ilze Zigurs, and other Associates. The minnesota GDSS research project: Group support systems, group processes, and outcomes. *Journal of the Association for Information Systems*, 9(10), October 2008. ISSN 1536-9323. URL <http://aisel.aisnet.org/jais/vol9/iss10/6>



Figure 3.9: Image of a server powering New Jersey Institute of Technology's EIES system in the 1980's. From http://wikiworld.com/wiki/index.php/EIES_History

tributed communication system developed by the New Jersey Institute of Technology starting in 1977.²⁶ Prior to the existence of a public Internet, NJIT began developing email, chat, and discussion group functionality using donated Unix servers with dial-up modems. The EIES system and its successors (especially EIES2) comprised the basis of the majority of GDSS research into distributed systems. More than a single group support package, the system was an extensible environment for composing interactive applications around a structure of typed data objects which could be modified through pre-defined actions. Using this system, researchers implemented various well-known group decision making processes including the Nominal Group Technique (a phased group ideation and choosing method) and the Delphi Method (an asynchronous, iterative consensus building process based on expert analysis). Studies using EIES2 employed any of these particular “activities”. Jim Whitescarver describes the design of EIES2:

EIES2 provides a comprehensive environment for electronic mail and computerized conferencing. The self organizing properties of the communications database help members cope with large amounts of on-line information. Conferencing coupled with a hierarchical item response structure organize communications into discussions and sub-discussions.... Activities within EIES2 integrate application support with communications.... A simple menu system with consistent screen layouts provides extensive on-line help for new users.²⁷

The Coordinator was an effort to build a group communication platform that worked on the basis of structuring group communication according to speech-act theory. While its authors did not describe it as a “decision support” application (and they actively critiqued the framing of “decision making” as an ontology for group action), it nevertheless fits easily into DeSanctis and Gallupe’s typology of GDSS systems, and stands out from its contemporaries in two ways: first, there is a thorough case study of the system’s usage in field trials with real groups performing endogenous tasks. Second, it employs a unique communications model which sets it apart from the highly similar group communication models used by GDSS systems.

Based on Carlos Flores and Terry Winograd’s theory of language and action²⁸, The Coordinator provided an electronic messaging system for workgroups which forced users to assign “types” to their messages, which allowed the system to automate the enforcement of workflows. Carasik and Grantham [1988] describe a case study deployment of the system in a work group at Pacific Bell. To use the system, as with other contemporary dispersed applications, users had to dial in to a modem several times per day, and manually run a

²⁶ For a thorough history of the early development of EIES and its goals, see [Hiltz, 1993]

²⁷ Jim Whitescarver. Electronic information exchange system II, February 2002. URL <http://web.archive.org/web/20020220171051/http://www.njit.edu/njit/Department/CCCC/eies.html>

U.S. Patent May 4, 1993 Sheet 11 of 40 5,208,748

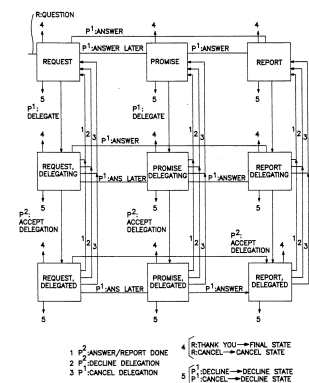


FIG. 13

Figure 3.10: Diagram from the patent for “The Coordinator”, showing the mechanism for delegating requests to other users. Flores and Winograd sought to commercialize the system [Carlos Flores et al., 1993]

²⁸ Fernando Flores, Michael Graves, Brad Hartfield, and Terry Winograd. Computer systems and the design of organizational interaction. *ACM Trans. Inf. Syst.*, 6(2):153–172, April 1988. ISSN 1046-8188. DOI: 10.1145/45941.45943. URL <http://doi.acm.org/10.1145/45941.45943>

command in order to check mail. Due to the heterogeneous hardware environment with incompatible computers, special-purpose PCs had to be installed just for the use of this application. Reports from the six-week trial were strongly negative:

Comments ranged from “worse than a lobotomized file clerk” through “this doesn’t fit the way we work” to “I learned it in spite of the interface.” Management urging was the motivation for continued use. At the completion of six weeks, the users voted to discontinue use of The Coordinator as a work group tool and adopt PROFS [(IBM’s host-based email system)] as a common electronic communications method.²⁹

Despite the manager imposing its use, “he did not personally use the software, he used his secretary to communicate with the group by way of The Coordinator.” This study demonstrates the danger of imposing a structured communication format incompatible with the group’s way of working, as well as the importance of careful attention to user interface.

²⁹ R. P. Carasik and C. E. Grantham. A case study of CSCW in a dispersed organization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '88, pages 61–66, New York, NY, USA, 1988. ACM. ISBN 0-201-14237-6. DOI: 10.1145/57167.57177. URL <http://doi.acm.org/10.1145/57167.57177>

Contemporary systems

While interest in research that calls itself “Group Support Systems” and “Group Decision Support Systems” dropped off by the end of the 1990’s, there are ample contemporary examples of web-based software tools that can support decision making activities of various types. Broadly, these can be classified into three areas:

- Groupware, or collaborative project tools
- Forums with voting
- Large-scale e-democracy systems

Contemporary systems tend toward light-weight, single-purpose tools that exist within a highly heterogeneous communications environment, using the web and email as common protocols. The web has replaced bespoke networked frameworks like EIES, and also obviated the need for monolithic tool packages that justified the high cost of implementation by trying to do everything. The result is more atomic, composable, and flexible tooling.

Groupware and collaborative project tools

We may not think of them as decision support tools by contemporary standards, but project management tools, issue trackers, and group calendaring and scheduling tools would all fit the taxonomy of group decision support tools from the 1980’s (specifically, “level 2” tools supporting planning, executionary, and intellectual tasks). I will

highlight two examples here which are particularly relevant to the design of InterTwinkles: *Trello* and *Doodle*.

Trello is a project management tool developed by New York's Fog Creek Software. The tool is loosely based on the *Kanban* system, a work flow and scheduling system developed by Toyota in the 1950's, which organizes work items into columns (representing state) and cards (representing tasks). The abstraction of cards and columns proves to be both highly intuitive and flexible, allowing adequate structure to help in organization, but is adequately unrestrictive to support a breadth of metaphors. As of January 2013, Trello has over a million users³⁰ – the site operates by a freemium model, where users can use the service for free, but can obtain more features by paying. Notable features of Trello include near-real-time collaboration (all user actions update in near-real-time on all connected computers, without locks or refreshes) and a highly intuitive, web-based interface which presents a very low barrier to entry.

Doodle is a light-weight scheduling tool, which presents a very simple poll for users to indicate what times they are free to meet. An initiator creates a scheduling poll, and chooses a set of possible meeting dates, and then emails a link to any potential participants. Participants can indicate their availability for the dates given. The tool is exceptionally simple – but this simplicity is its greatest advantage. Recognizing that scheduling is basically a simple information sharing problem (each participant needs to share their schedule, pertinent to the particular event), Doodle provides a simple-as-possible way to share that information. Once the information is shared, the “correct answer” is obvious – in that sense, scheduling could be described as an “intellective” decision making task. While calendaring systems could be configured to automatically share this information, this would require all participants to centralize on a single calendaring protocol. Furthermore, in order to preserve social flexibility in the notion of “availability”, such a system would still need a user's confirmation of whether to indicate availability for a given day (e.g. I might be “free” for a meeting with a close friend at a somewhat inconvenient time, but “unavailable” for an acquaintance with whom I'm willing to expend less effort). The extra complexity of developing unified protocols has little return in reduced data entry.

Crabgrass is a social networking tool targeted towards activists and organizers, developed by the Riseup Collective.³¹ Crabgrass was conceived primarily as a social networking tool, but also offers a variety of different group collaboration tools, including discussion forums, wikis, and todo lists. It offers fairly sophisticated modeling of group structures, including committees (subgroups), councils (committees with special powers), and networks of groups, with

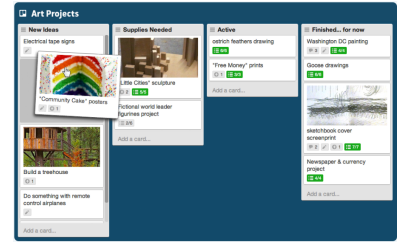


Figure 3.11: Image of a trello board, from <http://trello.com>.

³⁰ <https://trello.com/1m>

		July 31, 2013				
		9:00	11:00	2:00	4:00	8:00
Tom			✓		✓	
Paula		✓	✓	✓		✓
Chris			✓		✓	✓
		<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Figure 3.12: Image of a doodle board, from <http://doodle.com>.

³¹ <https://we.riseup.net/crabgrass/about>

careful attention to privacy and security as a primary concern. For decision making, Crabgrass offers simple ranked and approval voting tools. While Crabgrass garnered attention in 2009 and 2010 as an activist-friendly alternative to Facebook³², the project has not seen substantial development since 2011.

Forums with voting

There are a plethora of contemporary “decision making” systems which consist of some variation on a discussion forum with voting mechanisms. The most notable is *Loomio*.

Loomio is a discussion and proposal voting site developed by Enspirial, a New Zealand software collective. The site provides users with a group-specific space to create “discussions” and “proposals” for voting. The interface is very simple – the complete feature-set can be described as linear discussion, proposals, voting, “likes”, and notifications. While there are many other forum-plus-voting tools, Loomio stands out in the high design quality, the effort to educate users on its function, and the high degree of attention to user experience. The company prioritizes user education and documentation as much or more than technology.

Large-scale e-democracy

A tangential but related area to group decision making (and one that often comes up in discussion, as people associate online decision making with voting) is large-scale e-democracy systems, designed to enable large scale (100’s and more participants) voting and public deliberation. Notable surveys of the field include [Hacker and Dijk \[2000\]](#) and [Davies and Gangadharan \[2009\]](#). “Digital Democracy: Issues of Theory and Practice” (2000) is an edited volume of theoretical concepts in online democratic systems. Essays within the collection consider the early history of computers as communication media, theories of the public sphere in relation to digital democracy, theories of large-scale public debate, models of large-scale online deliberation, and case studies of voter guidance systems, government-driven messaging using the Internet, and online public debate. “Online Deliberation: Design, Research, and Practice” (2009) is a collection of case studies of systems for civic engagement, political dialog, “electronic rule-making”, and large-scale moderated policy discussions. The volume grew out of a 2005 Stanford conference on deliberation.

While this field primarily considers deliberation which, by virtue of scale and membership style is functionally and practically unrecognizable from affinity consensus (which is the focus of this thesis), I will describe some notable projects that tread the boundary between

³² Marcus Foth, Laura Forlano, and Christine Satchell. *From Social Butterfly to Engaged Citizen: Urban Informatics, Social Media, Ubiquitous Computing, and Mobile Technology to Support Citizen Engagement*. MIT Press, 2011. ISBN 9780262016513

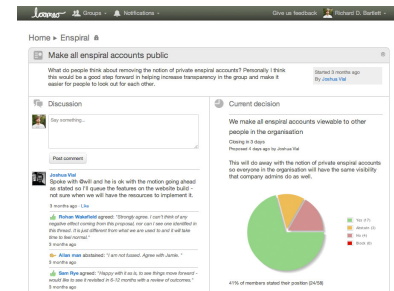


Figure 3.13: Image of a loomio board, from <http://loomio.org>.

large-scale and smaller-scale groups here: *ConsiderIt* and *Community PlanIt*.

ConsiderIt³³ is a platform for pro- and con-style deliberation developed by researchers at the University of Washington, developed for white-label usage in a variety of contexts (notably, it was used by the Washington state “Living Voters Guide” as a way to collect opinions on ballot measures). The system constrains comments into three types: issues, “pro” statements, and “con” statements. Individual users can build their own personal pro/con lists from the ideas generated by others, as a way to help make individual decisions about an issue.

Community PlanIt is a collaborative game platform which communities can use to engage in planning tasks for their neighborhoods. Organizers of particular games collaborate with local governments and organizations to back winning plans with real funds. Players complete “missions” to earn points, which they can then allocate as virtual coins to different ideas, helping to collaboratively choose project and development ideas. Unlike the other projects described in this section, which exist entirely online and can be used electively at any time, Community PlanIt instances only exist for the duration of particular events, and are only open to residents in the effected communities.

The Debian Project is an online community that builds an enormously popular free operating system. Since the community-led project first ratified its constitution in 1998, it has operated as a democratic body with elected leadership and a “Standard Resolution Procedure” that uses direct democratic voting to approve resolutions. While custom voting software is used to finalize some resolutions, the bulk of the deliberative process within the community takes place over email mailing lists (which average 50-100 messages per day) and online chatrooms (IRC). Key to the function of the democratic process is the well-specified procedures and policy, including a constitution, explicit procedures for referenda, and a membership process that ensures that new members of the project have the necessary skills, philosophical inclinations, and dedication to the project.³⁴

Conclusion

When I began research into GDSS, I read some of the aspirational writing from DeSanctis and Gallupe and was amazed at the clarity of purpose and well charted map for how the field could flower into usable tools – all the more so because I saw such a strong parallel between the structures suggested by these frameworks and the techniques for consensus decision making with which I had first-

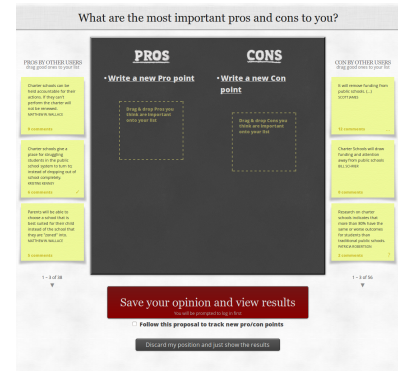


Figure 3.14: Image of a ConsiderIt board, from <http://wash.livingvotersguide.org>.

³³ Travis Kriplean, Jonathan Morgan, Deen Freelon, Alan Borning, and Lance Bennett. Supporting reflective public thought with considerit. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work, CSCW '12*, pages 265–274, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1086-4. DOI: 10.1145/2145204.2145249. URL <http://doi.acm.org/10.1145/2145204.2145249>

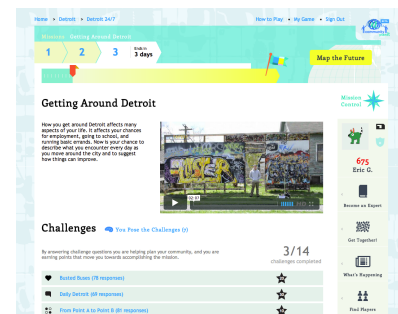


Figure 3.15: Image of a Community PlanIt game in Detroit. <https://communityplanit.org>.

³⁴ GUNNAR Ristroph. Debian’s democracy. *Online Deliberation: Design, Research and Practice*. Chicago, Illinois, USA: Center for the Study of Language, pages 207–212, 2009. URL http://odbook.stanford.edu/static/filedocument/2009/11/15/Chapter_17._Ristroph.pdf

hand experience. Having now spent a substantial period of time with this body of research, I think the answer to my early question, “Whither the GDSS?”, is clear: withered away; lost to a strangely myopic and overly positivist view, divorced from concerns of design, and mired in the technology of 1988.

The theoretical contributions of Desanctis and Gallupe [1987], McGrath [1984], and DeSanctis et al. [2008] provide valuable insights into framing research with groups and the design of decision support systems. The handful of real-world case studies of novel communications tools such as Carasik and Grantham [1988] provide instructive examples and inspiration for future designs. But the research perspective of the majority of the field leaves little to inform the design of new systems; I suspect, due to a lack of concern for user-centered design. My belief is that for systems as complex as group decision making, one can only obtain meaningful results by trialing the tools with real-world groups – and that the most meaningful evaluative metric is the groups’ choice to adopt or reject the design.

The technological landscape today is also markedly different from the heyday of GDSS. Reports from studies in the late 1980’s complain about the need to install custom PCs, and to ask users to dial in to group-specific modems multiple times per day in order to use the study systems.³⁵ Today, we can simply use the web and email, and it runs on the phone in your pocket. Contemporary systems also show a trend towards atomic, composable tools that perform one communication task well, designed to function within the heterogeneous communication environments of email and social media – this trend can relieve us from another round of monolithic, aging, do-everything tool packages.

³⁵ R. P. Carasik and C. E. Grantham. A case study of CSCW in a dispersed organization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '88*, pages 61–66, New York, NY, USA, 1988. ACM. ISBN 0-201-14237-6. DOI: 10.1145/57167.57177. URL <http://doi.acm.org/10.1145/57167.57177>

4 Design of InterTwinkles

Despite the advances in ease of access to computers and networking in the last 20 years, a central challenge remains for the design of tools to support non-hierarchical group decision making: a group must *choose to adopt* the tools as a group, and subsequently *choose to use* them as part of their ongoing work. The “field of dreams” attitude to technology design (if we build it, they will come) is unreliable for elective democratic systems: a democratic system, without its *demos*, is not actually a democratic system – as DeSanctis et al. [2008] remarks, “A voting feature does not, for all practical purposes, exist for a group if the group never considers or employs it.” Unlike target groups of field research in past GDSS, which had managers who could mandate that a particular new tool be adopted by a group, InterTwinkles’ groups must elect to do so of their own volition. Consequently, all of the mundane practical details which make a system hard or easy to use become as critical to its function as the theoretical models of decision making process it embodies.

To that end, I undertook InterTwinkles as a design research project to build a democratic system that groups would be interested in using. In this chapter, I will describe the process of participatory and iterative design that led to the platform’s current form, including discussion of key features and design principles – in particular, concerns relating to maximizing the modularity of components and autonomy of users. In the next chapter, I will describe the field trials we conducted with six groups over a period of three months, and a structural analysis of the groups’ adoption of the tools, and then subject the designed platform to the same critical analysis I’ve presented for others in this dissertation by analyzing the ethical fit and systems of control that the platform embodies.

Throughout the progress of this project, I have been a native participant in the groups that use these tools. During the field study, I also contracted with three research assistants who were living in groups that were participating in the study. This close perspective affords a potential for heightened precision in descriptive analysis. Employing community members as research assistants presents

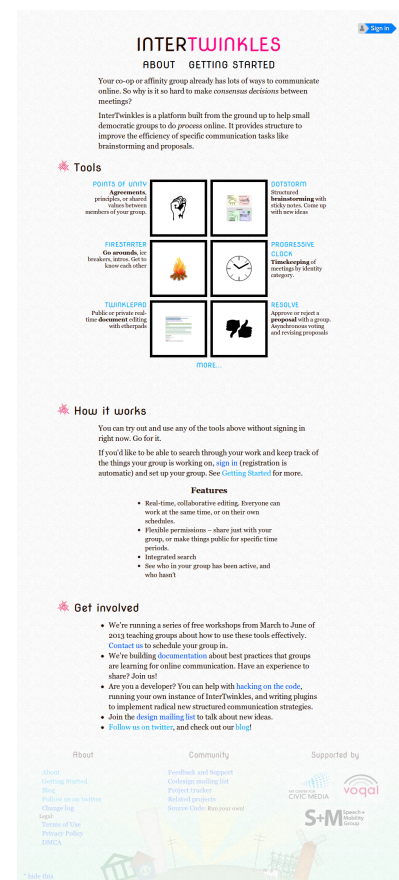


Figure 4.1: Front page of InterTwinkles, as of July, 2013.

a practical approach to participatory research where community resources are limited.¹

First steps: email-based decision making

In 2011, based on the motivations discussed in the introduction, I started playing with the idea of building tools to improve the between-meeting decision making processes in non-hierarchical groups to which I belonged. The medium that has been ubiquitously used in every formal group I've been in for the last 15 years for all asynchronous between-meeting work is email. From my own observations, I noted the following issues in attempts to discuss group concerns and make decisions over email:

- Discussions were often dominated by one or two people, who might argue a point of minutia back and forth, leaving little room for others to contribute.
- People would lose attention quickly; though it was never clear who was still participating and who was now ignoring the thread.
- It was never clear where the group was in its progress toward a decision – proposals weren't clear; votes and quorums were ambiguous; it was never clear when a matter was finished; etc. In short, there was no structure.

These issues are identical to the issues that plague poorly facilitated in-person meetings, but are made more acute by the asynchronous medium and absence of body language. Like many researchers of GDSS before me, I believed that adding even a modicum of structure to this channel to guide discussions toward concrete proposals and resolutions could be a step forward. However, from my experience with introducing new web-based group technology (e.g. wikis and calendaring systems – I deployed my first wiki with non-hierarchical groups in 2003), I knew that a central difficulty in introducing a new technology is getting people to use it. This is particularly important in a participatory decision making process where full participation is a normative requirement – even the most technology-averse participants must be comfortable.

Consequently, I set out to design a tool which would be based on email, to meet groups where they already were. Under the working name “byconsens.us”, I created a special-purpose listserv, which a group could use in place of its current mailing list, with the intention of helping people to facilitate decision making processes.

The listserv functioned in two modes: regular discussion, and structured decision-making. In the regular mode, it worked like any

¹ For discussion of the role of native participation in ethnographic study, see [Spersneider and Bagger, 2003], [Gatson and Zweerink, 2004], and [Maanen, 2011]

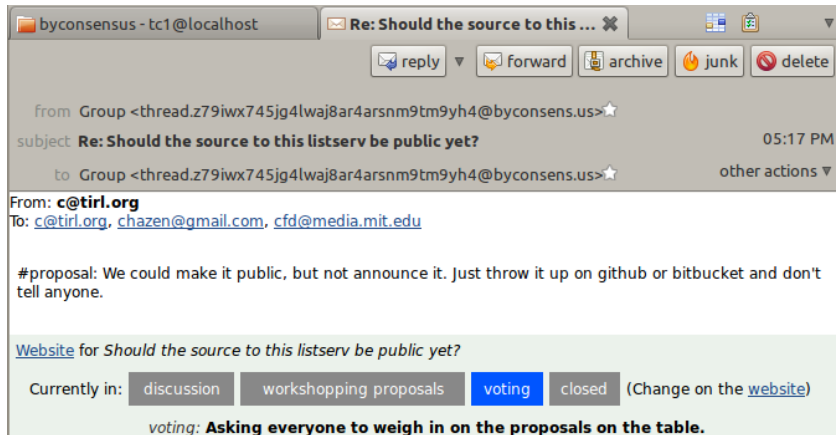


Figure 4.2: An email message, with context (the pale green box at the bottom) appended by the listserv, displaying the current stage of the process.

other mailing list; but in the structured mode, it added new functions. First, it created a dedicated webpage for each decision-making thread, which provided an alternate way to view the entirety of a discussion. Second, it appended a bit of information contextual to the current decision at the bottom of each email, indicating a prescriptive “stage” in a decision making process (one of “discussion”, “workshopping proposals”, “voting”, and “closed”), as well as any proposals or votes that were current. Third, it offered a way to take proposals out of the back-and-forth discussion of a thread, and instead make them objects which appeared on the web page and in the message footer context.

The “stage” of the proposal was simply an indicator which could be changed by any of the participants by clicking on a link on the website. The intention was to allow facilitators of an asynchronous discussion to proactively decide to change the current stage when they believed that the group had reached clarity – for example, once a proposal had been advanced, one might change the stage from “discussion” to “workshopping proposals” in order to encourage participants to begin stating their ideas in the form of proposals. The stage could be easily switched back and forth, giving groups flexibility to return to discussion or back off from voting if the jump to proposals was premature.

To allow smoother handling of proposals, they were pulled out of the discussion to a webpage (and added to the context box in emails). This made modifying proposals much simpler. Rather than copying and pasting the text of someone’s proposal from an earlier reply and modifying it – or worse, just suggesting modifications as replies in a threaded discussion – proposals could be manipulated directly on the web page. The current version of a proposal was always apparent from the footer of each email and the summary on the webpage.

Should the source to this listserv be public yet?

cfd@media.mit.edu 1 minute ago

What are the pros and cons of making this codebase public now? It's kinda rough around the edges still, for sure. It also contains test cases that include encoded URLs that log people in for rills.

The code is hosted on bitbucket (which recently started supporting git, and also allows private repos for free -- they charge for collaborators, not repos).

-charlie

c@tirl.org just now

#proposal: We could make it public, but not announce it. Just throw it up on github or bitbucket and don't tell anyone.

c@tirl.org just now
Activity changed to: voting

c@tirl.org just now
c@tirl.org voted yes for proposal 1

cfd@media.mit.edu just now
cfd@media.mit.edu voted yes for proposal 1

Process stage: discussion workshopping proposals voting closed

♣ Asking everyone to weigh in on the proposals on the table.


#proposal #anonymous Send message »

cfd@media.mit.edu | [threads](#) | [logout](#)

3 participants in this thread:
[c@tirl.org](#) [chazen@gmail.com](#)
[cfd@media.mit.edu](#) [remove edit](#)

Proposal 1

#proposal: We could make it public, but not announce it. Just throw it up on github or bitbucket and don't tell anyone.



[show totals](#)

Still waiting on votes from:
chazen@gmail.com

Figure 4.3: Web page backing an email thread in the early byconsens.us listserv.

To allow an entirely email-based workflow, the system used “hash-tags” to execute specific actions. An ordinary email discussion on the listserv was converted to a decision-making discussion by adding the tag “#byconsensus” – with the thought that group members could add a line in the course of an ordinary discussion: “Why don’t we decide this #byconsensus?”, causing the listserv to begin appending context to the thread and to create the webpage to back it. Proposals could be created by Adding “#proposal” to one’s message. An opinion could be sent anonymously (sender information would be stripped from the email) by adding the hashtag “#anonymous”.

I tested this system informally, using it with sympathetic friends and also with members of the housing cooperative in which I was living. Some key outcomes of those trials made me doubt the validity of this design’s approach:

- The email-based workflow was confusing. Users struggled with hashtags, and used them incorrectly or not at all. Consequently, proposals weren’t properly created or handled.
- The lack of any standard behavior for reply quoting among different email clients and users presented substantial barriers to correctly appending the contextual information about state without duplicating it. Worse, idiosyncratic quoting behavior could reproduce hashtags, with unexpected results.
- Users’ habitual patterns of hitting “reply” when starting new threads made algorithmically determining whether something was a new discussion very difficult. Neither the subject heading, nor the email headers (e.g. ‘reply-id’) were reliable indicators of which thread a message belonged to. Users would often accidentally continue a previous thread when they intended to start a new one, or start a new one when they intended to reply to an old one.

People understood and were able to use the web-based portion of the tool much more effectively, but the substantial hurdles in making the email-based experience work made me doubt that it was worth making email the primary medium for the tool. As a highly unstructured medium, email has evolved to support a wide variety of behaviors; imposing a structure onto it means fighting not just users’ ingrained behavior (which is substantial enough), but also the behavior of an unbounded multiplicity of client applications. While it might have been possible to train users to make changes to their workflow in order to meet the constraints of the system, the whole point of an email-based system in the first place was to reduce the need for major changes to users’ workflow. Instead, it seemed more prudent to work with the far richer affordances of

a web based system, and train users to switch contexts to the web when a discussion warranted a different structure.

Based on this experience, I went back to the drawing board for how to best approach the design. To that end, I began an intentional participatory design process with people from cooperatives and consensus-oriented groups around Boston.

Participatory design

At the beginning of 2012, in order to tackle this question more directly, I began a concerted effort to create a codesign project around the design of tools for consensus. Participatory Design² is a well-known approach to design which involves users throughout a design process; codesign is a particular take on Participatory Design with an emphasis on participation at the earliest stages of creation, often with formal relationships between designers and participant organizations.³ It is an approach which is particularly well-matched for design for consensus-oriented groups, because the groups' ethical orientation towards inclusion and participation are matched by the techniques of design. Under the helpful guidance of Sasha Costanza-Chock, I initiated a research program with a series of design workshops, while simultaneously beginning software development on a number of the components which would become InterTwinkles. Eric French Monge, a visiting student, helped with some of the early participatory workshops and ideation.

Our initial goal was to identify one or a few committed community partners with whom we could forge a formal relationship, complete with a Memorandum of Understanding – an approach that codesign researchers have found to be invaluable in creating strong relationships between researchers and communities which respect the interests and needs of each. After close to six months of effort to do so, we failed to identify a partner who had the time and capacity to formalize a relationship for the project. Many individual members of groups were interested, but we could not mobilize groups as a whole to commit to being partners in a design process. As a result, we shifted strategies to an iterative design process, relying more heavily on my own native understanding as a member of multiple consensus-oriented groups.

Though we were unable to secure codesign partnerships for the duration of the design process, support from the Media Ideation Fellowship later in the project enabled us to hire three part-time research assistants during the course of the field study from within the study groups. By paying community members as co-researchers, we were able to gain native insight throughout the field study and

² Clay Spinuzzi. The methodology of participatory design. *Technical Communication*, 52(2):163–174, May 2005

³ Elizabeth B.-N. Sanders and Pieter Jan Stappers. Co-creation and the new landscapes of design. *CoDesign*, 4(1):5–18, 2008. ISSN 1571-0882. DOI: 10.1080/15710880701875068. URL <http://www.tandfonline.com/doi/abs/10.1080/15710880701875068>

evaluation, while compensating for the burden of time required for participation.

Workshops

In March of 2012, we conducted three design workshops to lay the foundations of understanding for designing tools.



Figure 4.4: Whiteboard from the first design workshop. Detailed notes from the writeup are available at http://project.intertwinkles.org/projects/consensus/wiki/Mar17_Design_Workshop.

The first workshop included members of democratically-oriented groups and cooperatives in the Boston area. Focus in the workshop was on identifying the channels that people had available to them for communication, the affordances of them (particularly, what sorts of messages and communication seemed most appropriate to each), and the protocols and rules that people already used to structure them. We also brainstormed technologies that might be able to improve communication and understanding in groups.

The second workshop was conducted with activists and members of a community space in Providence, Rhode Island. The workshop focused on defining consensus, its benefits, its meanings, and its drawbacks for the members, followed by an ideation session around what sort of features might go into online tools to support consensus.

The third workshop was conducted with members of the Board of Directors of NASCO, a bi-national nonprofit that supports housing cooperatives across the US and Canada.⁴ At this workshop, we took a deeper dive into considerations about structuring communication. We played a game of “Moon Talk”, a live action role-playing game

⁴ I currently serve on the NASCO board, and had done so prior to this workshop, but was between terms at this particular meeting.



Figure 4.5: Chalkboard from the second design workshop. Detailed notes from the writeup are available at http://project.intertwinkles.org/projects/consensus/wiki/Mar25_Design_Workshop.

developed to motivate thinking about communication in structured environments (described below), and developing a set of principles for respectful email communication using an early prototype of Ten Points, a predecessor to the InterTwinkles tool “Points of Unity” (described below, page 108). At this meeting I also collected constructive input on wireframe mockups of some potential InterTwinkles tools.

Workshop games

During the initial participatory design phase of InterTwinkles, we developed two games designed to motivate thinking about online communication and decision making: *Flame War* and *Moon Talk*. We tried these games with a variety of groups, and found them to be valuable tools for inspiring critical discussion around communication processes.

Flame War is a card game for 4-6 players, which is based on group decision-making over email. The game consists of “Goal Cards”, which have different types (“Decide something”, “Schedule something”, “Share info”), and point values that vary depending on the difficulty of the goal. Game play consists of starting a group “goal”, and then playing “email” cards on that goal which can either support the goal, or “flame” it. If too many flames are played, the goal is discarded; players can win points from the goal by playing the email that crosses the needed threshold of support.

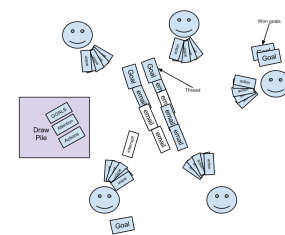


Figure 4.6: Diagram of game setup for Flame War. Full instructions and cards available at http://project.intertwinkles.org/projects/consensus/wiki/Flame_War.

The value of this game as a group educational tool is almost entirely in the *flavor text* (the messages on the cards that do not contribute to the game's mechanic), which highlight a variety of familiar constructive and destructive email behaviors. The design intention is to encourage groups to discuss and think about the ways that their email behavior is conducive to or detracts from effective group decision making. The hope is that using this game might open up in-person communication around the usually solitary pain of dealing with dysfunctional email threads.

Moon Talk is a facilitated live action role playing game for 10 or more players (the mechanic breaks down if there are too few people). The game is based on a scenario: all of the players are on a moon base, wearing moon suits and moon helmets. They can only communicate over a single-channel radio that they have in their helmets – and only one person can talk at a time. And, oh no!, a problem has arisen: mice have gotten into the food supply, and they're eating all the food. The group must now decide what to do about it. But first: they must schedule a time to meet.

Each player is given a piece of paper with three times written on it; these are the *only* times that player is available to meet. The task of the group is to find the one time when the largest number of group members are available. The distribution of times have been generated such that one time has 70% of the members, and other times have lower percentages – thus there is a single best time, but it doesn't work for every group member, and other times are close. As a result, with a large group of players, it will not be easy or obvious to determine the single best time without some attempt to do so procedurally.

The facilitator asks everyone to blindfold themselves (or just close their eyes), and to raise their hand when they wish to speak. Only one person may speak at a time. If all goes as intended, the group will fail to determine the correct time within the allotted period (3-5 minutes is sufficient). Next, the facilitator asks the group to take off their blindfolds, and to discuss a protocol as a group for how they will solve the problem. They're given the opportunity to add a new communication channel, "clapping", which they can assign the meaning they want to. The group is given new times, and asked to repeat the exercise with the new channel and protocol.

Design principles

Based on input from the workshops and our own experience in consensus-oriented groups, we developed a set of potential failure modes for this design project, as well as a set of principles to act as



Figure 4.7: The moon. Full game-play instructions and generator for schedules is available at <http://moon.intertwinkles.org/>.

guideposts to avoid those failures. To this end, we identified five major failure modes and nine guiding principles:

- *F1: Digital divide*: The goal of this project is to improve the way that groups can use consensus processes. But if a group adopts a tool which structurally excludes some members of the group, this is a huge fail. The other side of the coin is that digital technology could enable *more* participation – limiting consensus process to in-person meetings can structurally exclude people who don't have the time or money to travel to a meeting. The most insidious digital divide failure mode would be a tech-savvy subset using the tool heavily, and thus structurally excludes a particular set of members. Principles derived from this include:
 - P1: Using the tools, participation should increase or remain constant.
 - P2: Levels of participation should be visible to users of the tools.
 - P3: Control of the tools should be accessible to all participants, regardless of their degree of technical sophistication.
- *F2: Bad scope*: Is it better to build something that is highly suited to a particular purpose, or something that is more general? Software developers are frequently warned against “premature generalization”; at the same time, inadequate generalization can lead to lower utility and adaptability. The takeaways are:
 - P4: Don't be an island. The tool should integrate with existing systems.
 - P5: Don't be everything to everyone. If the needs of different potential users are too different to reconcile, pick one, and do it well, rather than doing many things poorly.
- *F3: Framing failure*: Framing failures are related to scope failures, but focus more on the problem statement than the proposed solution. The abundance of past efforts at group decision support systems – most of which have failed to gain user acceptance – gives pause: Is an online decision support system something that there is a real need for? To avoid failures of framing:
 - P6: Ground design in existing practice. Preference what people actually do (and how they think about what they actually do) over what we might imagine is better.
 - P7: Develop a theory for why something failed, and address the reasons, before building something similar.

- F4: *Redundant design*: In addition to past tools for group decision support systems, there are a plethora of existing sophisticated tools for general communication and collaboration – including shared documents, mailing lists, social media networks, etc. One should be careful to avoid recreating this work without a good reason.
 - P8: Don’t build Facebook. The tool should not attempt to replace any existing functionality, without a very good reason for doing so.
- F5: *Inadequate user involvement*: We were unable to recruit any participants who were willing to sign on to a project as a full codesign partner complete with a memorandum of understanding; nevertheless, there were a large number of interested but diffuse participants without formal relationships. The concern this raises is that we might end up failing to adequately involve people in the design process, and thus end up with something that fails to meet users’ needs. A principle to avoid this problem:
 - P9: Don’t build anything without a clear target user, who is participating in the design.

An evaluation of the design outcomes in light of these failure modes and principles is discussed below on page [122](#).

Iterated platform development

Following the initial participatory design phase, I began an iterative design phase to create a robust platform that would meet the usability standards for regular use in the target communities. To do this required addressing a large number of seemingly “paper cut” style minutia; though these issues were non-trivial once subjected to scrutiny. They included authentication, event tracking, notifications, as well as technical considerations of the server model. Throughout, I sought to design toward the requirements of activist groups that required autonomy from centralized government or corporate services. In this “design notebook” style section, I will describe the major hurdles and design choices.

Server architecture

The InterTwinkles server architecture went through three major phases: first, an initial phase when all the tools stood alone as entirely independent services; second, a mixed phase, where the tools were independent services (running on separate servers), but shared

a common authentication system and group membership infrastructure; and third, an integrated platform where each of the tools is a plugin, but does not function independently of the whole. This evolution reflected the tension between the desire to maintain loose coupling and a decentralizable architecture, and user needs that required integration between components in order to be fulfilled.

While historical GDSS tools tended to come in monolithic packages with multiple different communication paradigms and group interactions combined into a single installation (justifiably, given the high cost of custom hardware necessary to operate systems of the day), with the ubiquity of the Internet and web-based devices, we can now create much more modular single-purpose tools which will still be easily accessible to all group participants. This carries multiple advantages: it allows for much more rapid advances for individual tools, since each can be changed without harm to the others. It allows for designs that support natural *composability* – each tool can be used alone, but can also be arranged in different configurations with other tools to support particular group needs, just as a group might mix different facilitation techniques in a single meeting. A modular architecture also helps to minimize the dependence that users have on a single vendor to maintain a central toolkit. This expression of *subsidiarity* also allows contributions from different developers to the suite tools to support different interactions.

With this approach in mind, while conducting the participatory design workshops, I began to develop the first tools that became part of the InterTwinkles platform: *Dotstorm*, *Ten Points* (later renamed *Points of Unity*), and *Progressive Clock*; each operating as a standalone independent service. Each operated without any authentication, utilizing “secret URLs” as the strategy for privacy – the same strategy used by other familiar tools such as Doodle and Etherpad⁵ (both solid examples of independent tools supporting a focused group collaboration task).

The downside of independent tools, of course, is the need to replicate infrastructure. One piece of infrastructure in particular became problematic: authentication. Many early users of the tools requested a way to automatically index all of the activities which they had completed with the tools. Indexing, however, would scuttle the privacy offered by secret URLs, unless users could authenticate before being allowed access to indexes and search. This was complicated further by the group context: each activity completed with the InterTwinkles tools is a *group* activity, which is most meaningfully owned by a group, not an individual. To allow for meaningful group-based actions (such as adding and removing group members, or viewing new activities started by your group’s members) required creating a

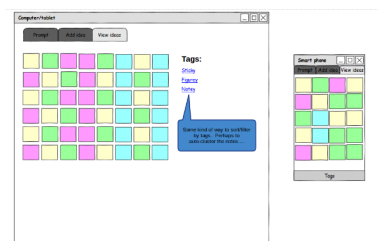


Figure 4.8: Early design mock-up of Dotstorm, developed in part with workshop participants.

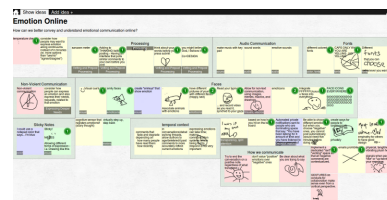


Figure 4.9: Early stand-alone version of Dotstorm.



Figure 4.10: Early stand-alone version of Ten Points.

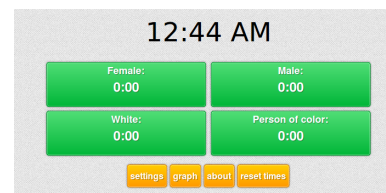


Figure 4.11: Early stand-alone version of Progressive Clock.

⁵ Doodle (<http://doodle.com>) is a simple online scheduling service. Etherpad (<http://etherpad.org>) is a free software real-time collaborative document editor.

representation of a group and its members within the access control system. But asking users to create and maintain a representation of their group multiple times, once for each tool, is far too onerous – that aspect demands centralization.

In order to address this, I developed a central “group server” which would handle the tasks that were common needs for all the InterTwinkles tools, and which particularly benefited from centralized function: authentication, group memberships, search, event logging, and notifications. This server exposed an API which allowed separate services running on separate servers to provide these features to users without having to re-implement them. While this solved the need for centralized functions like a group model, it introduced a considerable amount of complexity, which slowed development and complicated testing. Furthermore, this decreased the *replicability* of the entire infrastructure, which put an uncomfortable degree of central control on the service. To install the group-server based infrastructure, one would need to configure and expose multiple different components, including mapping subdomains. This made installation substantially more arduous than the standard for self-hostable web applications such as Wordpress. To ease this, I developed installation scripts that would provision virtual servers with all necessary components installed – but even that still limited self-hosting to people who could afford to provision a virtual server. I suspected that this complexity would decrease the likelihood of replicated installations and additional developers contributing to the project. I wanted InterTwinkles to be a service which people could use without trusting me – and therefore, wanted to make the architecture easier for other people to install on their own servers.

Finally, I iterated the infrastructure a third time, to use a single stand-alone server, but which implemented each of the communication tools as plugins. This preserves a similar degree of subsidiarity as the previous iteration (one can still write plugins without having to replicate the central functions of authentication, groups, search, events, and notifications), but simplifies the architecture, so that one doesn’t need to have a complex infrastructure with multiple servers and mapped subdomains in order to run your own implementation. While this decreases the atomicity of each tool, they remain similarly composable.

It would be preferable in the long run to have a federated “group service” which different communications tools could consume. Existing social graph paradigms (including overlapping reciprocal “friend” relationships with Facebook, personal “circle” or “aspect” views with Google Plus and Diaspora, and unidirectional “follow” models with Twitter or StatusNet) do not include the fully connected small-group

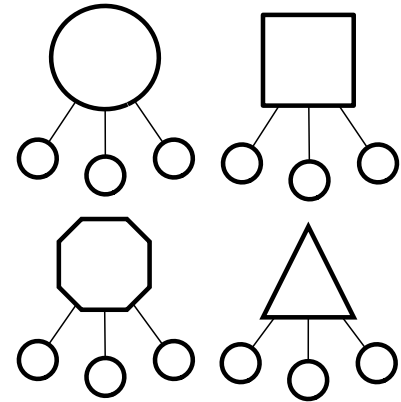


Figure 4.12: The first iteration of the InterTwinkles architecture: an infrastructure of pure subsidiarity – each tool stood completely independently (small circles represent users, large shapes services).

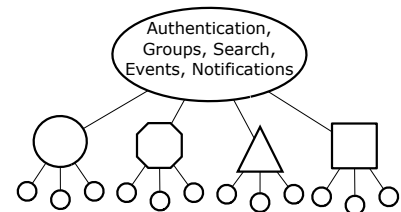


Figure 4.13: The intermediate iteration of the InterTwinkles architecture: central group server handling authentication, groups, search, events, and notifications; but distinct servers handling different communications tools (e.g. Dotstorm, Points of Unity, etc.)

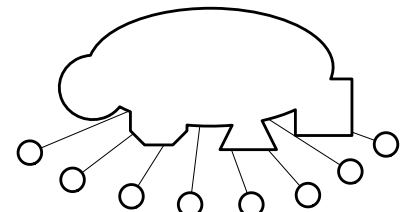


Figure 4.14: The final iteration of the InterTwinkles architecture: central server with plugins. Less architecturally modular, but far easier for developers to install, making decentralization through replication more likely.

or “workgroup” model which would be required to support affinity groups. As a result, developers wishing to write group-centric applications are stuck re-implementing basic group infrastructure, and groups are stuck re-creating representations of their membership for each service they want to use – and due to the difficulty of convincing members of a group to sign up for a service, this affords a high degree of vendor lock-in. To support activist’s needs for autonomy, such a service would need to be federated to avoid reliance on the good will of a single corporate provider or government to provide this service.

Group and privacy model

Authentication remains optional for all of InterTwinkles’ tools – signing in offers a “progressive enhancement” which enables notifications, event logging, search, indexing, and privacy strategies beyond secret URLs. Rather than requiring authentication from the beginning, users can start activities without signing in, and later “claim” the activity for their group after authenticating. This approach allows for maximum flexibility – one can use the tools anonymously, for a quick one-off, or one can require all of the members of one’s group to sign in before viewing or participating.

Every activity shares a simple but flexible sharing model, which is group-centric. First, the activity might belong to a group, or have no group. With no group, the activity is shared on a secret URL basis – anyone with the URL can view or edit the activity. Once an activity is claimed by a group, it defaults to being private to that group, so that one must be logged in as a member of the group to edit or view it. If the group wishes less restrictive access, they can make it publicly viewable or editable, or add specific outside people who can have access to edit or view it. In addition, the group can choose to open access for a short period of time (e.g. an hour, a day, a week), so that the group can ease restrictions for a particular event (such as a public meeting), but avoid long-term exposure to search engines and vandals.

The group model for InterTwinkles assumes that members of a group will have a high degree of trust in each other – in this respect, it is designed to favor the needs of *affinity consensus* groups over the needs of *assembly consensus* or other types. However, this approach gives substantial advantages to enabling full participation: group members can take actions on each others’ behalf. That way, a group member that is away from their computer (or prefers not to use the computer) can still be represented in InterTwinkles activities. The assumption is that since group members know each other, they will

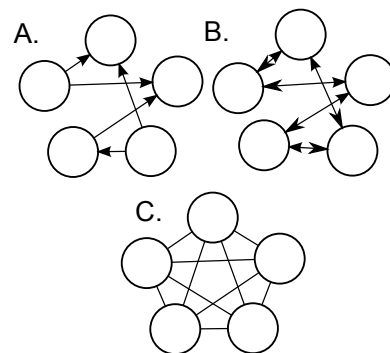


Figure 4.15: Models of a group used by various social networks – circles represent people. (A) Twitter uses a non-reciprocal “follow” model. (B) Facebook uses a reciprocal, but still not fully-connected “friend” model. (C) Affinity groups have a fully-connected “group” model, as they collectively organize around a central resource (such as the coordinated action of the group).

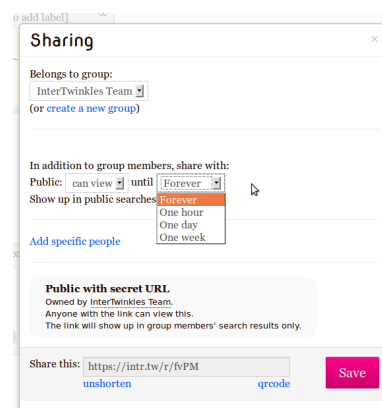


Figure 4.16: InterTwinkles sharing settings dialog.

hear about it if someone takes an action on another's behalf without their permission. To preserve transparency, the event logging records both the user who took the action and the user on whose behalf it was taken.

Administration of membership in the group is available to any group member – there are no “moderators” or privileged user accounts. Using the event and notification systems, all group members are notified if the membership of the group is changed (for example, a member is added or removed).

Event logging

Data privacy is a substantial concern to many activists – the proliferation of analytics and surveillance tools on the web, either in the form of advertisers who sell users' browsing habits to increase advertising revenue or governments who wish to infiltrate and control activist groups, has led to an arms race between privacy protecting systems and fingerprinting tools that seeks to thwart them. At the same time, event logging has the potential to be of tremendous benefit to the users of applications. The “quantified self” movement has seen an explosion of interest in personal data tracking – we could imagine a similar movement as a “quantified group” which is more aware of its performance. The problem with event logging and analytics isn't so much that data is being collected, it's *for whom* and *to what end*, and who's in control. In the affinity group context for which InterTwinkles is designed, there is tremendous collective value in knowing what other members of your group have been doing with respect to your collective activities. If I create a proposal, it's very helpful to know who in my group has viewed it. If I create a brainstorm, it's helpful to know if a particularly creative member of my group has had a chance to contribute. In asynchronous environments, without this type of event logging, each action is as a shot in the dark; one is never sure if others have noticed or care.

For this reason, InterTwinkles includes a sophisticated event tracking framework – but one which exposes all of its data to groups, and is optimized for legibility to them. There are several contexts in which event summaries are shown: on the group's dashboard, where one wants to see all the events pertaining to that group; in an activity summary email, informing group members of things their group has worked on recently; and in a summary of events specific to a single activity.

With such a variety of contexts, each of which would summarize events semantically differently, this has the potential to become burdensome to tool developers. To maximize its benefit, the event

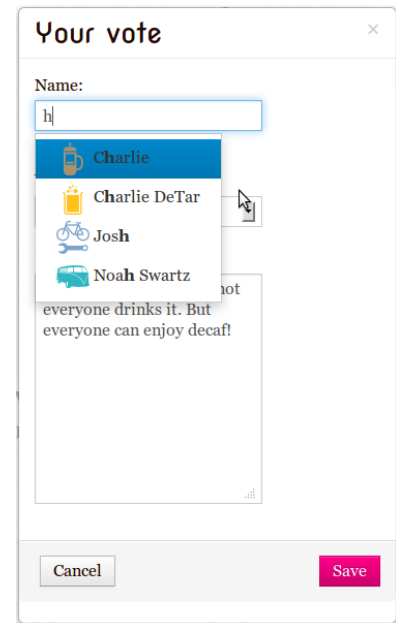


Figure 4.17: Voting dialog for *Resolve*, which allows group members to vote on each others' behalf.

system needs to support arbitrary event types from tool authors (for example, “added a new drawing”, “changed a vote”, “retired a point of unity”, or “foobared a dingle on a whizbang”), but to not burden developers with anticipating all the grammatical contexts in which such events might be used. To solve this, I implemented a simple event grammar, which asks the developer to specify five parts of speech for each event, allowing them to be combined relatively simply:

- Entity: a noun-phrase for the object of the event.
- Aspect: a component or qualification of the entity.
- Collective: a noun-phrase for a collection of events of this type.
- Verbed: a past-tense verb describing the event.
- Manner: a qualification of the verb.

With each tool plugin defining functions to provide these parts of speech for each event type, it became relatively straight-forward to generate accessible, natural language event summaries which organized the events in different ways.

Security

Given that InterTwinkles’ target user communities includes activists whose activity is sometimes targeted by government surveillance, security is a chief concern – but one that currently involves compromises which make the system unsafe for activist groups participating in direct actions or civil disobedience. As a result, when groups contacted me asking about the security of InterTwinkles for contestational political work, I advised them against using it.

In 2013, the long-held fears and assumptions of many activists that the US Government engages in broad-scale surveillance of its citizens were confirmed by revelations about programs such as PRISM and XKeyscore. These programs involve both passive data collection (in which government agencies monitor and record all traffic passing through various interchanges), and active collection in which the government coerces private service providers into giving access to their users’ data.

The InterTwinkles server encrypts all data in transit between users and the server using HTTPS (Transport Layer Security), configured with “Perfect Forward Secrecy” so that a compromise of the SSL certificate will not compromise past. This provides substantial, but not fool-proof, protection against passive surveillance. A man-in-the-middle attack which uses a forged certificate could allow a

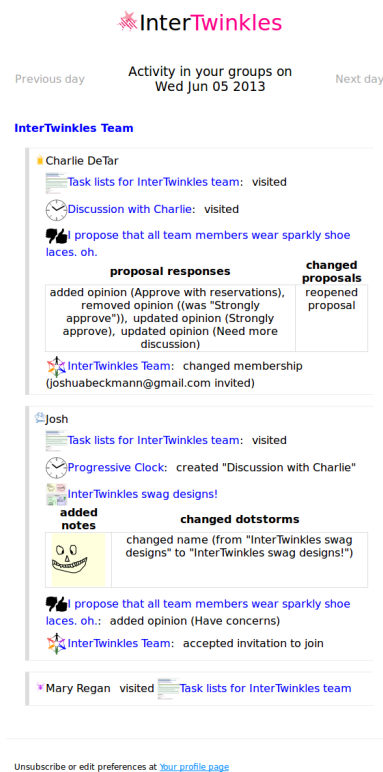


Figure 4.18: Example activity email, showing a summary of everything everyone in one’s groups did in the last day across all InterTwinkles tools.

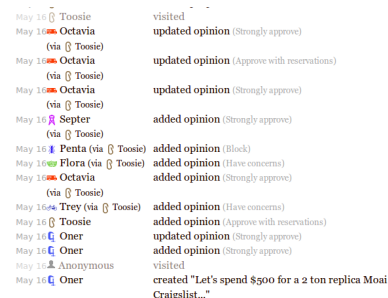


Figure 4.19: Events for a single activity, from that activity’s page.

committed attacker to eavesdrop on data in transit between users of InterTwinkles and its server. This type of attack is unlikely to be used by a government except in extraordinary cases, as it is easy to detect – and if it were detected, would likely result in a shutdown of the certificate authority that issued the forged certificate. Nevertheless, any weakness of SSL as a protocol also impacts InterTwinkles' data in transit between users and the server.

It is not currently practical for group-oriented server-based systems such as InterTwinkles to be “host-proof” (that is, to encrypt all data end-to-end so that operators of the server do not have the technical capability of viewing data). This means that event data, activity data, and anything else recorded by groups using the tools could be divulged at a later date if server operators were subpoenaed. In addition to the risk of a direct subpoena to operators of InterTwinkles, the server uses Mozilla's Persona service for authentication, which means that in principle Mozilla could be coerced to giving access to particular users' accounts.

InterTwinkles' codebase is open source and published on the Internet, which is a double-edged sword, security-wise. While it means that friendly programmers could find bugs to improve security, malicious attackers can also look for bugs which they can exploit. It is possible that bugs exist which could give attackers access to users' data. The software is also designed to be as easy to install and run as possible, so users could in principle download and operate their own instances of InterTwinkles, for which they control all of the data. While this has the potential to reduce the attractiveness of a single InterTwinkles server for attacks, it also relies on users to develop the expertise to maintain secure servers and keep software up to date.

Tool profiles

The intention of the InterTwinkles platform is to support a growing, pluggable collection of different communication protocols, as a way of increasing groups' communication capability, and encouraging experimentation around different strategies for reaching agreement in asynchronous contexts. Described below are the initial six tools developed for the platform. Each is based on an adaptation of tools that people use in face to face consensus-oriented meetings, but adapted to the needs of an asynchronous context.

The communication tools are not, and not intended to be, comprehensive. Human facilitators are still needed to guide a group's process when using the tools. Also notably absent from all of the tools present here is any facility for free-form discussion. This is an intentional omission: every group with which I've worked already

has a mailing list which they use for this purpose, and the presence of another discussion forum would only serve to divide attention. Instead, each tool is designed to be an atomic and composable component of a larger discussion and decision making process, which can be referred to from discussions in other fora.

All of the tools are real-time collaborative. Any edits made by any member of the group are immediately available to all other members, without refreshing the page. An indicator on each page shows the other group members that are currently looking at the same page.

1. *Dotstorm*

Dotstorm is an application inspired by sticky-note brainstorming and the Nominal Group Technique. First described by [Delbecq and Ven \[1971\]](#), the Nominal Group Technique is a group ideation process which proceeds through five stages:

1. An introductory explanation of the problem or issue.
2. Silent generation of ideas by group members.
3. Sharing and explanation of ideas by members who generated them.
4. Group discussion of the ideas.
5. Categorizing, voting, and ranking ideas.

The technique is one of the more widely used and successful formal ideation techniques discussed in the “problem solving” school of group decision making activities.⁶ Contemporary facilitators often stress the value of asking group members to express their generated ideas using visual language (e.g. drawing) as a benefit to creative thinking. The name “dotstorm” comes from the face-to-face application of this technique, where ideas are added using sticky notes or written on poster paper, and votes are added using “dot” stickers next to the ideas people like the best.

Dotstorm allows the creation of notes, which can consist of text (which is searchable), drawn images, or photos. Photos can be uploaded from one’s hard drive or shot directly with the device’s camera (laptop webcam, or camera on phones or tablets – most devices which can run a modern web browser are supported). Once added, notes can be arranged and grouped through drag and drop, and “votes” can be added to any note. Event summaries for modifications to a Dotstorm include the changed or added notes – activity summary emails thus include images showing notes that were added.

⁶ Arthur B. Vangundy. *Techniques of Structured Problem Solving*. Springer, April 1988. ISBN 9780442288471

(a) Recruiting co-ops for InterTwinkles

How can we recruit co-ops to participate in the InterTwinkles field trial?

+ Add

<p>What would it take for coops around the area to come together for a sustained joint project? We are trying to find out and build tools to help. We need your experiences to make</p>	<p>Cooperatively minded folks shouldn't have to be rich or re-invent the wheel to get new coops going. We can build on our successes and learn from our mistakes!</p>	<p>Coop house members already know the value of sharing resources. By sharing ideas and resources across coops, we can expand the movement!</p>	<p>No city has a grassroots network of co-ops yet capable of raising and distributing funds. Do you want to be part of the first?</p>
<p>We are spending over 80 hours per week for the next 3 months on this project. Would you like to benefit from our labor?</p>	<p>What worries you about participating?</p>	<p>The co-op loan fund is already well developed; but we need governance to move it forward.</p>	<p>How do you form a network?</p>

(b) InterTwinkles swag designs!

Designs ideas for potential stickers and plushes for InterTwinkles.

+ Add

Josh's favs

circle logo

Star

Black border

Georgia font

white border

Plush? <-----

plush

plush

plush

(c) Tech4Engage People

Draw, name, tag, and group each other!!!

+ Add

Tags

knight

- events MIT
- app maker
- social space
- Harvard
- policy
- solving problems
- Josie Heath
- Boulder CO
- Rockies fan
- Knight Miami
- Denver NYC
- writer

philly startu

- hacademia
- hackers
- Marci Harris Popvox
- California
- attorney
- fun despite that
- mark beadd
- gov cfa
- Bob Filbin data and new Yc
- Community informatics
- academic
- canadian
- Lucy Bernholz
- illustration

(d)

T

this is a note...

Tags

Share

Figure 4.20: (a) A largely textual brainstorming session, (b) Ideation about swag and logos for InterTwinkles, (c) Drawing faces as an ice breaker for a new group, (d) The Dotstorm note editor.

During development, I paid particular attention to ensuring that Dotstorm was usable in both synchronous and asynchronous contexts, with any web-enabled devices. In synchronous contexts, if one participant has a device with a camera, other participants can add ideas using paper notes, and use the camera-enabled device to add images of the paper to the shared board.

2. Points of Unity

Points of Unity (originally called “Ten Points”) is a tool for creating a set of shared principles or statements that everyone in a group agrees to. The design was originally inspired by a now defunct synchronous facilitated “Collaborative Democracy Workshop” developed by *may first / people link*.⁷ May first’s original design asked participants to split into groups of 4-5 people per computer, each of whom were asked to work together to author “rights” (as from a “bill of rights”) to which everyone in the small group agreed. Any rights added by the group would be projected on a shared display visible to all participants. Each subgroup was then given the opportunity to endorse a right, or if they weren’t willing to endorse it, to edit it, with the goal of finding rights that all subgroups were willing to endorse. My implementation originally followed this process – with ten hard-coded slots into which “rights” could be added.

⁷ The inspiring tool (and its facilitation guide) can be found at <http://meetings.mayfirst.org/>.

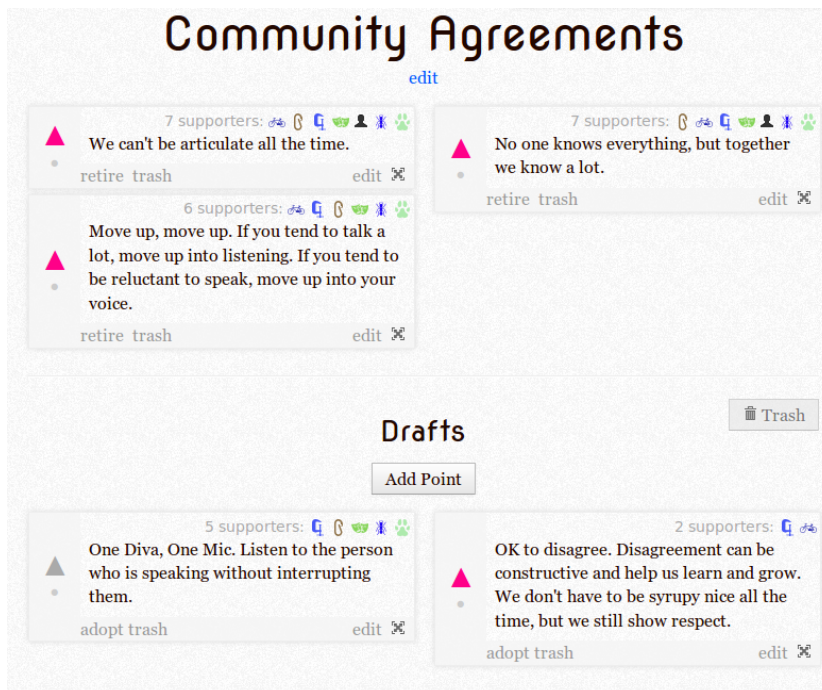


Figure 4.21: Points of Unity used to develop a set of “community agreements” for how a group speaks to each other.

Based on feedback from a number of early group tests I facilitated, as well as feedback from groups who stumbled upon the tool on the Internet and used it for their own purposes, I made several modifications to the design. First, many early groups were reluctant to edit points that others had written – it was seen as an aggressive act to change someone else’s words. To encourage people to suggest alternatives, I removed the hard-coded limit of ten points, but introduced a distinction between “approved” and “draft” points, with the idea that there would be lower social pressure to adding a “draft”.

The second major change was to change the unit of endorsement from an ephemeral subgroup of people to individuals, tied to Intertwinkles’ authentication system. This substantially improved support for a particularly compelling usage model which one early user group had innovated: the group developed a set of shared principles that everyone agreed to, and then asked new members to visit the web page and endorse all the points, or indicate the ones they disagreed with to renegotiate with the whole group. Under the original model, each subgroup was an ephemeral collection of people – so the “endorsements” of points was less useful in an asynchronous context. Tying endorsements to the individual supports asynchronous usage much more effectively. The original subgroup-based use case is still supported, since the tool supports voting on behalf of other group members – and thus a subgroup with a shared computer can record endorsements for each of its members.

3. *Resolve*

Resolve is a tool for resolving (approving or rejecting) proposals. This tool is intended for use when, after a group has discussed an issue in some other forum, there is a coherent proposal which is emerging for which the group now wants to finalize a decision.

While the basic needs for a proposal are conceptually simple (one person suggests a body of text, and other people respond with votes and statements), the details become very important for operation in asynchronous contexts. One example is the way the interface supports editing of proposals after voting has started. A key question to allowing changes to proposals after voting has started is: what should be done with the old votes? If votes are invalidated, simple grammar fixes or unimportant changes can’t be made without forcing everyone to vote again. If votes are retained, there is a danger that people will be shown to have supported something they don’t actually support. To address this, when a proposal is changed, Resolve marks any existing votes as “stale”, and tallies them separately. That way, a member of the group can make a judgment call as to whether

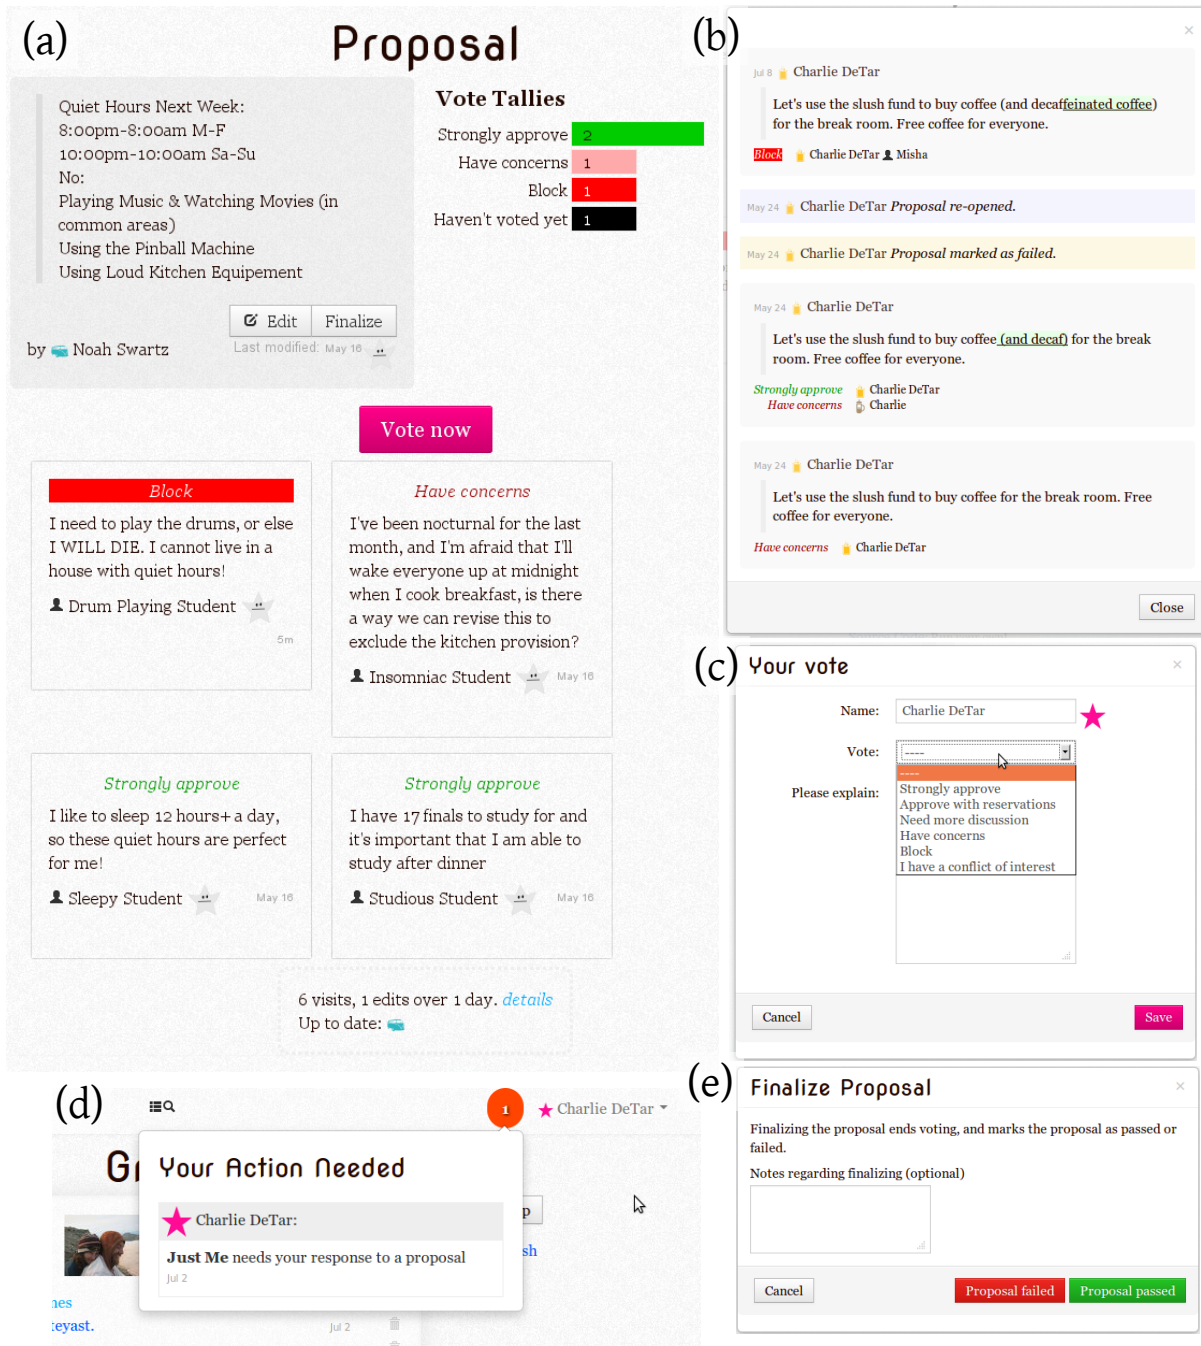


Figure 4.22: Resolve interface: (a) Board for a proposal, showing votes, opinions, and event history. (b) Dialog showing the history of a proposal, (c) Dialog for adding a vote, (d) On-site notification for a proposal, (e) Dialog for finalizing a proposal.

to accept those votes or to reject them when considering the tally. Some other interface details supporting asynchrony include:

- The interface makes it immediately apparent who has and hasn't participated, and the events dialog shows who has and hasn't visited.
- The voting interface allows group members to vote on behalf of each other, so they can fill in for people who are not near a computer.
- Votes are not binary (yes or no), but allow finer grades.
- History is tracked for all changes in opinions and proposals.
- Non-group members who have been given permission to view a group's proposal can cast advisory votes.

Resolve does not enforce any condition for whether or not a proposal passes – groups can decide their own policy for what constitutes an acceptable level of dissent.

4. Firestarter

Firestarter is a very simple tool for doing “go-arounds”, a discussion format where each member of a group speaks in turn as a way of doing introductions, or to share thoughts on a difficult issue. This tool is the least complicated of the InterTwinkles tools, and only supports adding responses and “twinkling” (similar to “liking”, but without any aggregation or side-effects outside of the context of the thing that is twinkled). (See figure 4.23)

5. Progressive Clock

Progressive Clock is a tool for tracking speaking time in groups by identity category (e.g. white, male, female, person of color). The categories are configurable – so any pertinent category that is important to a group can be included (e.g. old members, new members; or in small groups, individuals' names). This can be helpful to groups which are encountering difficulty where people with particular identities dominate meetings. Abundant analysis of speaking time in groups has found a strong correlation with behavioral dominance and speaking more in meetings (see ⁸ for a meta-study showing this result across a variety of group types). In particular, cis-maleness and whiteness correlate with dominance over people who don't share those identities – white male speakers will speak more, interrupt others, be called on more often, and speak in authoritative tones.

⁸ Marianne Schmid Mast. Dominance as expressed and inferred through speaking time. *Human Communication Research*, 28(3):420–450, 2002. ISSN 1468-2958. DOI: 10.1111/j.1468-2958.2002.tb00814.x. URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1468-2958.2002.tb00814.x/abstract>

5PM Friday, January 11, 2013

Mozilla Knight Fellows ice breaker

What is your favorite thing to pair with peanut butter?

[Add response](#)

































<p> Charlie DeTar  1 Jan 11</p> <p>PhD student at ML and C4CM. Peanut Butter and Nutella.</p>	<p> Eric  Jan 11</p> <p>PhD at CSAIL. Never acquired the taste for peanut butter, so I'll say bacon, which goes well with liver pate, which is much better.</p>	<p> Ed  1 Jan 11</p> <p>Work here at C4CM. Celery -- you can scoop a lot with a piece of celery.</p>	<p> Pablo Reyes  1 Jan 11</p> <p>Visiting scientist at C4Cm. From Spain, so I don't do peanut butter. Surprised that you can make your own peanut butter in the super market with a grinder.</p>
<p> Rogelio  Jan 11</p> <p>RA at civic, master's student at CMS. Apples (granny smith), red molle.</p>	<p> Kate  Jan 11</p> <p>IP work at Media Lab -- peanut butter and jam!</p>	<p> Denise  Jan 11</p> <p>RA in civic, first year in CMS. Noodles with peanut butter. Soy sauce, honey, etc. Big hit with peace corps.</p>	<p> Matt  Jan 11</p> <p>Allergic to peanut butter! Epi pen.</p>
<p> Roy  Jan 11</p> <p>Undergrad in civic. Peanut butter with apples.</p>	<p> Neha  Jan 11</p> <p>Peanut butter with bananas and honey.</p>	<p> Ted  Jan 11</p> <p>CSAIL. First time I'm shameful to be in love with peanut butter. Belgian waffles and yogurt.</p>	<p> Dan Schultz  1 Jan 11</p> <p>C4CM graduate, 2012 mozilla fellow. Not a great cook, but like to throw random ingredients into things including peanut butter. Peanut butter sauteed sauce.</p>
<p> Molly  Jan 11</p> <p>RA at center, grad student at CMS. Peanut butter and toast and honey.</p>	<p> Erica  Jan 11</p> <p>Community manager for open news. Peanut butter and potato chip sandwiches.</p>	<p> Charlie De Tar  Jan 11</p> <p>Nutella</p>	<p> Sonia  Jan 11</p> <p>From China. Open News fellow. Don't buy peanut butter, but anything with peanut butter like candy bars</p>

Figure 4.23: Using Firestarter as an ice breaker for the Mozilla-Knight Fellows program

Worse, these white male speakers often lack self-awareness of their dominance, and may not intend it. By measuring speaking times of people with different identity categories, one might be able to more objectively critique groups' success in confronting structural oppression.⁹ More than any of the other InterTwinkles tools, this tool is highly geared towards synchronous use – it is unlikely to be useful in an asynchronous context. (See figure 4.24)

⁹ See [Kim et al., 2008] for a paper discussing the group effects of tracking speaking times in controlled groups with no prior history.

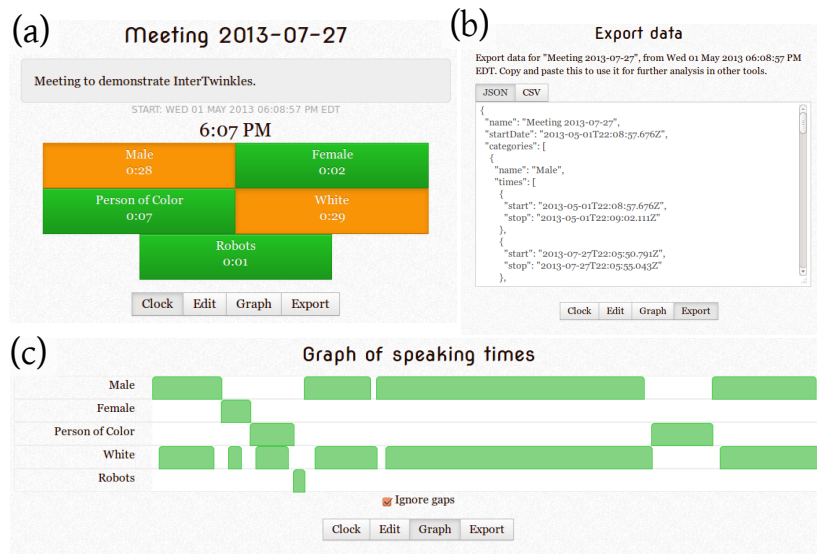


Figure 4.24: Progressive Clock interface: (a) The timer, (b) Exportable data, (c) Graph of speaking times.

6. TwinklePad

TwinklePad is simply a re-packaging of Etherpad, a free software collaborative document editor. It provides real-time collaborative document editing, but also integrates with InterTwinkles' search, event, and authentication systems. Unlike plain Etherpad, TwinklePad documents can be made private to a group, using the same sharing and privacy mechanisms as the other tools. As a demonstration of the efficiency of InterTwinkles' plugin architecture, complete Etherpad integration (including event tracking, search indexing, document archiving and deletion, etc.) was possible with only 200 lines of front-end code and 400 lines of back-end code. As with all other activities, documents edited in TwinklePad appear in the group's dashboard, and in summaries of events. (See figure 4.25)

The screenshot shows the InterTwinkles web interface. At the top, the logo 'InterTwinkles BETA' is on the left, a search icon in the center, and the user name 'Charlie DeTar' on the right. Below the header, the document title 'Code book' is displayed with a 'change' link. To the right of the title are buttons for 'Email Group' and 'Sharing'. A rich text editor toolbar is visible, containing icons for bold, italic, underline, strikethrough, bulleted list, numbered list, indent, outdent, undo, redo, background color, heading selection, link, unlink, star, settings, and user profile. The main content area shows a document with the following text:

1 **Definitions for codes to use in interview coding**

2

3 **Coarse categories:**

4 • **Facilities:** any things (meeting spaces, white boards, sticky notes, emails, computer programs, technologies, books) that people use as part of their decision making practice.

5 • **Norms:** any rules, expectations, principles, or guidelines, either explicit or implicit, which guides a group's decision making practice.

6 • **Schemas:** any qualitative interpretations of a group's decision making practice.

7 • **Meta:** comments on the interview or research process itself.

8

9 **Fine categories:**

10 • **_appropriation:** Any instances or examples of using a new facility in intended or unintended ways.

11 • **_confusion:** Uncertainty about what something is, how it works, what to expect.

12 • **_compliance:** Examples of following or failing to follow norms.

13 • **_difficulty:** Challenge and difficulty in accomplishing a desired result.

14 • **_efficiency:** Time and effort required to accomplish something, whether a lot or a little.

15 • **_influence:** Examples of people changing each other's minds or behavior through their words or actions.

16 • **_motivations:** Reasons or feelings for why something should be done a certain way.

17 • **_rule avoidings:** Examples of norms being intentionally violated.

18 • **_rule followings:** Examples of norms being followed.

19 • **_understanding:** Discussion of people either understanding or failing to understand one another's feelings, desires, thoughts, etc.

20

Figure 4.25: Example document edited in TwinklePad. Since it uses Etherpad as its base, it supports all of Etherpad's features, including coloring the background of text with a different color for each user, making it clear who has edited which parts of a document.

5 *Field Trial and Evaluation*

In March of 2013, we launched the InterTwinkles website publicly. We decided to focus efforts for a field trial on cooperative housing in Boston, for two reasons: first, I had significant existing connections with that community, making it easier to establish trust for recruitment. Second, the Boston cooperative housing community had had significant efforts from some of its members to organize a city-wide network of cooperative houses (in the form of the “Boston Collective House Assembly”). Organizers had set out to develop a cooperative startup fund, which would raise funds from within the cooperative community, and distribute them to new co-ops that were seeking to form. However, the organizing efforts lost steam; energy had substantially waned in the last year, in part due to the assembly’s lack of any governance structure and means of making decisions. We hoped that by kickstarting efforts at building community skill in online decision making, we might be able to contribute some energy back to that movement.

Our study design was to perform an analysis based in Adaptive Structuration Theory¹, with a focus on outcomes of adoption and appropriation. We would analyze the norms, facilities, and interpretive schemas that groups already operated with, and examine how they interact with the protocols afforded by InterTwinkles. In particular, we were interested in groups’ choice to adopt or not adopt the tools, and the degree of fit between the spirit for which the tools were designed, and their actual use. To inform a baseline of understanding of the groups’ current practices, we conducted a series of one-on-one interviews with members of the groups. We then conducted training workshops with each group to teach the tools and discuss ways the groups might consider relating them to their existing practice and policy. Midway through the trial period (which lasted approximately 3 months), we offered further training (though no groups took us up on the offer). Following the six week trial period, we did closing interviews with members of each of the groups to reflect on their understanding, adoption, and use of the tools.

All interviews were recorded, transcribed, and coded – in total,

¹ Gerardine DeSanctis, Marshall Poole, Ilze Zigurs, and other Associates. The minnesota GDSS research project: Group support systems, group processes, and outcomes. *Journal of the Association for Information Systems*, 9(10), October 2008. ISSN 1536-9323. URL <http://aisel.aisnet.org/jais/vol9/iss10/6>

we interviewed 19 people within the test groups, for an average of 45 minutes each, and 7 additional interviews with cooperative housing community members not in the test groups. As part of the closing interviews, we also conducted short surveys of feature comprehension.

Participating groups

Recruitment

After advertising on local email lists, cooperative housing community meetups, and in-person, we spoke directly with 29 collective and cooperative housing groups in the Boston area. Our criteria for recruitment were that the group must have a group decision making process that they regularly use, must have between 5 and 40 members, and must consider themselves to be a housing co-op or collective.

Of the groups we contacted, 6 agreed to participate, 11 declined, and the remainder did not reach a decision during the recruitment period. At least two of the groups we contacted who decided not to participate in the study decided to use the tools on their own, unsupported (an option that we suggested). Two more groups responded affirmatively, but too late for us to include them in the study. One lesson we learned from the recruitment process was that it takes a rather long time to mobilize a consensus-oriented residential cooperative to make a decision to participate – we had anticipated this taking 2-3 weeks, but it ended up being up to 3 months from our first contact before the groups finalized a decision of whether or not to participate. Factors that delayed responses included:

- Groups' meeting schedules: some groups met only once per month, and required an introduction at one meeting, and a decision at the following meeting.
- Travel schedules of group members: we conducted this trial during the summer, and residents of several of the co-ops were away for significant parts of the recruitment and study period.
- Some groups had substantial other time commitments as a group, and did not have time right away to participate in an elective study. Having a longer window during which groups could join might mitigate this.

For future studies, of this type, we would recommend a minimum 3 month recruitment period. Also, rather than planning to recruit all groups simultaneously for a parallel study, a rolling strategy where

groups on-board and finish on their own schedules may be more effective.

Participants

Six groups agreed to participate in the study. To preserve the privacy of the groups, their names are changed here. Three of these groups (Maple, Aspen, and Hemlock) are rental co-ops, one (Sapling) is a newly forming group seeking to purchase or rent property. One (Sassafras) is a group-equity non-profit that owns their property, and the last (Pine) is the Board of Directors that manages Sassafras (it partially overlaps in membership with Sassafras, but also has members from elsewhere in the Boston cooperative community). In total, these groups represent 45 unique users – 7 people were members of more than one participating group. Table 5.1 shows a summary of the membership, organization style, meeting frequency, and meeting style for each group.

Name	Members	Founded	Organization	Meeting frequency	Meeting style
Maple	10	2000	Lease co-op	Monthly	Committees, whole-house social meetings
Aspen	7	2011	Lease co-op	Occasional, as-needed	Whole-house
Sapling	9	2013	Newly forming	Scheduled as-needed	2-3 core organizers, newer recruits, ad-hoc
Sassafras	12	2005	Group equity co-op	Weekly, regular	Whole-house
Pine	9	2002	Board of Directors	Monthly, regular	Whole-board
Hemlock	5	2011	Lease co-op	Monthly, regular	Whole-house

Table 5.1: Membership, organization, and meeting characteristics of the six recruited groups.

As a more thorough description, table 5.2 summarizes an analysis of the groups' meeting practice based on interviews with members of each of the groups at the beginning of the study, using a framework from structuration theory. The table shows the facilities used for meetings, the norms (rules and protocols) for how decision making works, and interpretive schemas (assessments, subjective feelings, interpretations) around how they identify with their decision making practice. These characteristics form the basis of the analysis of appropriation on page 125.

All of the groups identify as cooperatives, but the form and formality of this was highly variable. A member of Aspen described her house as “a very loosely run co-op ... almost a punk house” – and this reflected in the very informal, “doing” oriented meeting practice (members often take action, and await statements of concern from others, rather than consulting up front). By contrast, a member of Sassafras described her house as “an intentional community ... [with] a really strong commitment to being cooperative and cooperative values.” Meetings in Sassafras are correspondingly formal and

structured, with strong facilitation. Maple was the only co-op that does not regularly have meetings to discuss house business (though they do meet socially at least once per month) – this is because they recently changed their process to an entirely committee-based system. Rather than holding a meeting to consult the group for decisions, members discuss concerns in committees, and then announce plans to the house via the mailing list, relying on members to speak up if they have concerns.

All of the co-ops identify as non-hierarchical and democratic – in every case, no member of is more empowered than other members to direct the group’s action, and this equality is held as a key value for the group. Every one of the co-ops regularly uses an internal email list as part of their group discussion practice; all of them also occasionally make decisions over email. Several of the groups use a variety of other collaborative technologies – Google Drive is particularly common for sharing minutes, agendas, and other documents. Doodle is frequently used for scheduling. Maple occasionally surveys members to inform decisions using Survey Monkey. Two of the houses regularly use custom software written by current or former members – Sassafra uses a chore scheduling tool (written by me in 2008); Maple uses custom software for finance, note-taking, and coordinating orders for a bulk food buying club.²

In addition to the groups participating in the study, as InterTwinkles is running freely and publicly on the web, a total of 278 additional people have registered user accounts on InterTwinkles, forming 60 additional groups. Their usage is not considered in this analysis.

Training

To ensure that groups had the basic understanding necessary to use InterTwinkles’ tools, we conducted a 90 minute training for each participating group. For each workshop, we followed roughly the same facilitation plan, though in some groups we had to compress the workshop into a shorter period of time. Goals for the workshop were:

- Teach the tools
- Develop a plan to integrate the tools with groups’ current policy
- Establish methods for reflection – making sure everyone has a way to give feedback.
- Ensure that the tools are accessible to all group members.
- Learn about any particular group needs in order to iterate the designs.

² Google Drive is Google’s collaborative document service, which offers shared text documents, spreadsheets, and uploaded files. Doodle is a simple scheduling service (described on page 84). Survey Monkey is a Internet survey tool that allows users to set up questionnaires and tabulate responses.

	Maple	Aspen	Sapling	Sassafras	Pine	Hemlock
	Mailing List Survey Monkey Google drive Excel documents Email newsletters Financial data analysis (custom) Online note-taking tool (custom) Bulk food ordering system (custom)	Mailing list Doodle	Mailing list Meetup.com	Mailing list Internal wiki Google calendar, Google site Doodle Online chore system (custom)	Mailing list Internal wiki (shared with Sassafras) Quickbooks Google drive Doodle	Mailing list Google calendar Google drive Email newsletter Text messages
Facilities						
	No formal house meetings; work happens in committees. Social meetings once per month, regular dinners. Committees advertise to group, but don't require participation. Consensus expected; silence implies consent. Committee meetings unstructured. Expectation of non-violent communication (though not strict).	Meetings occasional, as-needed. Free-form, loose consensus. No formal facilitation. Flexible participation. Rule: "Don't be a motherfucker".	Meetings via meetups, email. Use consensus; but norms are still being established. Group has 2-3 core organizers, other newer members.	Weekly meetings. Formal consensus. Strong facilitator. Attendance expected. Agenda prepared in advance; shared by email. "House Norms": 13 points of agreement. Occasional house retreats.	Meetings once per month. Formal consensus. Strong facilitator. Attendance expected. Preparation for discussion expected. Governed by bylaws. Legally incorporated 501(c)(3) non-profit.	Meetings once per month. Loose consensus/majority; unanimity required for membership decisions. Informal facilitation. Strive for full participation. Agenda prepared in advance.
Norms						
	Mature co-op with "lots of systems that work well". Relaxed attitude. Strong identification as cooperative; centered on food, cooking, eating. Averse to discomfort – so concerns get responses. Everyone is taken seriously, if they speak up. Free-form meetings (were structured in the past)	Informal, almost a punk house, "but more functional". Action preferred over deliberation. "Co-op lite". Avoid meetings to avoid conflict.	New group; interpretations still being established. Identifies with "paleo" – interpreted as a lifestyle, not just a diet. Aspires to be cooperative.	Identifies as intentional community focusing on respect, personal growth. Fairly structured and formal. Strong identification as cooperative. Approaches conflict by confronting and naming issues, e.g. "Blushacking".	Fairly structured and formal. Identifies with neighborhood, wider cooperative movement.	Informal, not a lot of process/rules. House identifies with its newsletter and public events. Became more cooperative after a less cooperative-minded housemate moved out. Defer decisions to avoid conflict.
Schemas						

Table 5.2: Characteristics of the groups' meeting processes.

- Find ways to share and rotate roles for facilitation and responsibility within the group.
- Inspire groups to use the same critical creativity with online tools that they already use with offline tools; and to establish the norm that failures or bumbling mistakes are normal and good parts of the learning process.

Prior to the workshops, we asked groups to individually register as users with InterTwinkles (a process that required clicking a confirmation link in an email, and thus sometimes took time). We asked everyone to come to the workshop with a web-enabled device (smart phone, tablet, or laptop). I also brought an extra tablet to each workshop to share with anyone who didn't have a device handy; though it was never needed.

The agenda for the workshops was:

- 10 min: Brief introduction to InterTwinkles, its history, and goals. Pass around paper pamphlets which introduce InterTwinkles. While this introduction is happening, ask someone in the group to begin setting up the group's membership in InterTwinkles, which involves listing every member's email addresses.
- 5 min: Discuss the importance of accessibility and full participation.
- 5 min: Ensure that everyone is now signed in, has received the invitation to join the group, and has joined. A common difficulty encountered was that the person setting up the group might use a different email address for a group member from the one with which they had registered, though this was easily fixed.
- 45 min: Practice using the tools. (We had originally intended to only highlight a couple of the tools, but we found that once groups had the tools in-hand, they were curious and demanded to know what each one was for). As a group, go through each tool, create a test activity, and practice using it.
- 5 min: Discuss the value of feedback, and point out the ways to give feedback on the site.
- 20 min: Go around to talk about ways that the group might be able to use these tools, and integrate them with policy. Answer any remaining questions.

In the process of conducting these trainings, exercising each tool on a variety of devices, we also encountered a large number of bugs to fix, as well as requests for features which we then added. Among the iterations made in response to workshop requests were:

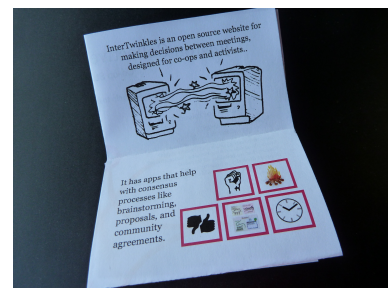


Figure 5.1: A pamphlet given to each group member as part of the training workshops. The full pamphlet can be found at <http://static.intertwinkles.org/zine.pdf>.

- Adding a “twinkles” (like) function to Firestarter.
- Changing the default sharing settings for an activity that belongs to a group from “public” to “private”.
- Reworking the graphical display of event histories.
- Changing the way lists of activities the group has worked on are displayed.
- Numerous bug fixes, from minor annoyances (such as forms missing default values) to major issues (such as data loss in Dotstorm caused by a timing bug in database updates).

Even though we had done extensive hallway testing and usability testing of all the tools prior to the workshops, the heterogeneous workshop environment with mixed computers, latencies, and user expectations proved invaluable to identifying issues with the user interface.

Ongoing support and project visibility

Mid-way through the field trial, we visited each group again to speak with them and offer additional support and training with the tools; though none of the groups took us up on the offer for additional training. We also gave each group an InterTwinkles “plush”, a star-shaped stuffed animal, as well as InterTwinkles stickers. Our hope was that by having a physical reminder of the existence of InterTwinkles which could sit in a common space in each housing co-op, we would remind members of the potential to use InterTwinkles should a group decision making need arise.

In addition to this, we maintained an active presence on Twitter, and maintained a blog with regular posts.

Evaluation strategy

To evaluate InterTwinkles’ success in meeting the goals I set out with, I will consider four different dimensions:

1. Evaluation based on design principles: At the beginning of the design project, we set out 9 principles that we wished to follow. I will evaluate the extent to which the design adhered to these principles, based on subjective analysis of the outcomes, as well as quantitative metrics of use.
2. Structural analysis: I will consider the appropriation of InterTwinkles by its user groups, in light of the prior schemas, norms,



Figure 5.2: InterTwinkles plush, designed by <http://fluffyland.com>.



Figure 5.3: InterTwinkles stickers.

and facilities that the group had, and the spirit and structure of the tools.

3. User feedback: I will analyze solicited and unsolicited feedback from participants, along with data collected during the field study.
4. Analysis of control: Given the design requirement of maximizing the autonomy of groups using the tools, I will analyze the systems of control at work in InterTwinkles, using the framework developed in chapter 1.

Evaluation based on design principles

In this section, I will consider the nine design principles from which this project started, and evaluate their relation to the results. The quality of this form of evaluation depends on the validity of the original premises.

P1: Using the tools, participation should increase or remain constant. We went out of our way to ensure that participation would be as easy as possible – particularly with flexible access control, and group-aware features that allowed group members to take actions on behalf of others, as well as training to ensure that each member understood how to use the tools. This effort appears to have paid off within the study groups. Using a Pielou’s evenness index to measure the degree of difference in participation based on the logged events in InterTwinkles, we find a high degree of evenness, with an average index of 0.94 (indicating very even participation). Figure 5.4 shows the levels of participation for each user in each group. The distribution does not reflect expected power law distributions for participation in online communities; however, the sample size is far too low to draw conclusions from that.

We don’t have data for comparable offline evenness of participation; so we are unable to say whether this reflects any change from pre-study norms. From the closing interviews conducted with members of each group, we don’t believe that InterTwinkles substantially changed overall levels of group participation.

P2: Levels of participation should be visible to users of the tools. The InterTwinkles events system (described above, page 103) made semantically clear event listings available to all members, and showed who had visited since the last edit of each activity. Daily activity summaries were emailed once per day to group members – 82.9% of users elected to receive these emails (almost all of the users opting out of email notifications were from the first group to use the system, who encountered several bugs in the notification system, which might have accounted for their choice to opt out). In addition to the

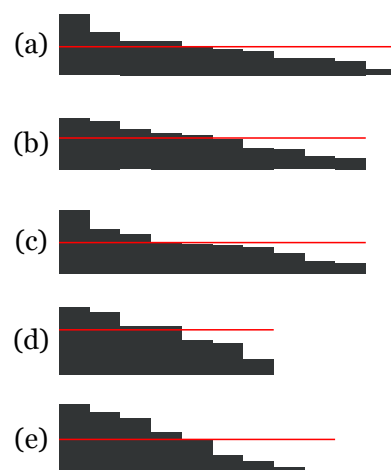


Figure 5.4: The evenness of participation in each group. The red lines show the level which would be perfectly even participation. Some of explanations of outliers: In group (a), the person with the most participation was a paid research assistant working on InterTwinkles; however, this was not the case for group (c). In group (d), the person with the least participation was traveling away from the group for most of the study period, and did not participate online during that time. In group (e), the three lowest participating members were newer members of the group who joined partway through the study period. One group did not use InterTwinkles after the initial training; they are omitted here.

activity summaries, displays of participation were present for different activities, as well as on the group's dashboard. When users were interviewed about whether they thought levels of participation of their group members were obvious, it was rated on average between "somewhat obvious" and "very obvious".

P3: Control of the tools should be accessible to all participants, regardless of their degree of technical sophistication. For use of the tools, access was demonstrated to be sufficiently easy for most participants. However, more work remains to be done to make other user interface elements more intuitive. In interviews, we asked participants about their comprehension of 11 key features. On a four point scale between "not at all obvious" and "very obvious", the average rating across all features was 3.5 (between "somewhat obvious" and "very obvious"). Participants had personally used half of these features. Some user interface elements that were listed as particularly unintuitive were the voting and "adoption" interfaces for points in Points of Unity.

The naming of individual tools was problematic – the names (e.g. "Firestarter", "Dotstorm") did not make the function of the tools sufficiently obvious. This may have led in some cases to less-optimal choice of which tools to use for particular tasks. In comprehension questions in the closing interviews, some users could not identify the function of some of the tools by name. More intuitive naming, and perhaps also a wizard-like tool to help users to choose a tool, might improve the fit between tool choice and task.

The tools were not designed such that any member of a group without knowledge of computer programming could modify their basic function – thus, for most group members, control is limited to questions of "access" rather than "evolution" of the tools themselves. Feedback systems were available to all participants; however, they still depend on the good will of developers to listen and make changes. See the analysis of control before for more on this topic.

P4: Don't be an island. The tool should integrate with existing systems. The tools successfully integrated with email and SMS. Google Hangout integration was partially completed (we hope to complete that functionality soon). Two of the tools (Dotstorm and Points of Unity) were designed carefully with the intention of supporting embedding the results in other web pages; however, in the course of the study we did not finish a requested feature to provide embed codes. As a result, users would have to know enough HTML to be able to write code for an iframe in order to embed them in blogs or other venues.

P5: Don't be everything to everyone. If the needs of different potential users are too different to reconcile, pick one, and do it

well, rather than doing many things poorly. One major choice that we made in designing these tools is to privilege the experience of small groups that trust each other over larger or more anonymous groups (e.g. choosing *affinity consensus* over *assembly consensus*). This rendered the tools less usable for one of the more visible contemporary expressions of consensus-based decision making (Occupy style General Assemblies), but conferred considerable benefits in terms of substantial simplifications in access control. We were able to demonstrate who has *not* participated in an activity precisely because the list of participants was small and known. Similarly, we were able to allow users to take actions on behalf of each other, improving participation for those with less computer access. By clearly choosing affinity consensus practitioners as our target, we fulfilled this principle.

P6: Ground design in existing practice. Preference what people actually do (and how they think about what they actually do) over what we might imagine is better. Each of the six tools in the initial InterTwinkles kit is more or less based on practices that consensus-oriented groups already do – and we took time to thoroughly consider the breadth of different techniques used for consensus, as well as existing techniques for online decision making, at the beginning of the design process. The one tool that was the least similar to existing practice is Progressive Clock. While “progressive stacks” are a common feature of affinity consensus, tracking speaking time in meetings by identity category is relatively newer. Nevertheless, I believe this is close enough to existing motivations and practice that it passes muster.

P7: Develop a theory for why something failed, and address the reasons, before building something similar. An example of this approach is found in the voting system for Resolve. There are many simple polling tools on the Internet; however, none which are suited to affinity consensus to the degree that Resolve is. We believed that the following issues with existing polling tools would impact use, including:

- No awareness of who the group is, thus no ability to show who has not yet participated.
- No privacy, or requires a new username/password for each group member, with little return for registering.
- Poll insufficiently expressive, and doesn’t allow modification of the proposal after it has started.
- Poll is difficult to set up.
- No ability or norm to vote on behalf of other group members.

- Groups lack any way to connect polls with their existing policy.

Resolve and the InterTwinkles platform address all of these issues, which to us justified building a new polling tool. We opted not to build in any mechanism for free-form discussion, however, as we believed that would be redundant to existing tools.

P8: Don't build Facebook. The tool should not attempt to replace any existing functionality, without a very good reason for doing so. The intentional choice not to include a discussion feature is an expression of this principle – groups already have mailing lists which they use; so our goal was to make InterTwinkles work along with email. The one InterTwinkles tool which is most likely to be considered redundant is TwinklePad, which simply re-packages Etherpad, allowing shared document editing (which is also duplicated by widely used tools such as Google Docs). There are two reasons why we felt that TwinklePad was appropriate: first, the existing plugin infrastructure of InterTwinkles and robust codebase of Etherpad made integrating it relatively trivial. Second, integrating Etherpad adds benefits of privacy, searchability, and indexing of shared documents, with group ownership. Neither standalone Etherpad systems nor Google Docs make this possible.

P9: Don't build anything without a clear target user, who is participating in the design. Each of the tools added to InterTwinkles were done with either direct participation of users, or in response to direct feature requests from users. The development process that was most remote from user participation was the somewhat lengthy process of determining the appropriate design of the server architecture. As this design question was based around the needs of sysadmins who would be installing instances of InterTwinkles, I consulted regularly with other web developers around these choices – I did not succeed in finding web developers within the target communities (affinity consensus practicing groups) to consult with regard to these questions.

Structurational analysis of appropriation

Analyzing group appropriation of a new technology is fraught with complexity – results are confounded by everything from the personalities of individuals in the group, to their histories and relationships, to the task type. Adaptive Structuration Theory, one attempt to grapple with this complexity, was the theoretical culmination of efforts by researchers in GDSS to approach this problem (see 72 for an overview of AST). As an analytical lens, AST considers the reflexive relationship between the facilities, norms, and interpretive schemas of the

incoming group, the “spirit” of the introduced technology, and “appropriation moves” (actions the group takes to use a technology). The result of this interaction is an emergent group with changed facilities, norms, and interpretive schemas that reflect adoption of the technology.

Historically, research into GDSS used “decision support rooms” and short synchronous decision making exercises in a laboratory setting with study groups. In such environments, researchers could videotape the sessions, and use micro-sociological techniques of conversational analysis to cite appropriation moves – identifying every moment where someone talked about or suggested that a feature of a GDSS be used. In a field study where groups are electively using tools when and how they wish (asynchronously), we don’t have that luxury – moments of appropriation happen largely in private, in spaces to which we don’t have researchers’ privilege of access. Further, intrusive efforts to observe these moments would have altered them in ways that would compromise the resulting analysis.

Instead of surveiling users in this way, we conducted interviews with each of the groups after 8-10 weeks of use, with a set of questions designed to elicit responses that characterize the interpretive schemas, appropriation moves, and norms groups established with the tools. For the most part, there was not time in the course of this study to identify organizational change or new explicit rules – however, one group did adopt explicit practices around how the tools should be used. In the subsections below, I will discuss several questions pertinent to appropriation, including:

- How did the groups use the tools?
- How did tool usage relate to the “spirit” of their design?
- What do we know about moments that impelled decisions to use the tools?
- How do the groups interpret the tools now?

How did groups use the tools?

Table 5.3 shows the number of activities completed by each group for each tool, as well as the number of events (edits and unique visits) logged for each. More events indicate higher levels of participation with that activity.

Usage of the tools covered a fairly wide range of tasks – everything from making major policy decisions to choosing songs for a playlist. A member of *Sassafras* described said they used it for things of minor importance: “It’s been a tool to use for things that aren’t

	Dotstorm	Points of Unity	Resolve	TwinklePad	Firestarter	Clock	Totals
Maple	3 (13, 95, 9)	4 (109, 18, 27, 73)	5 (61, 80, 29, 11, 20)	2 (11, 10)	5 (22, 9, 24, 33, 17)	2 (16, 8)	21 (695)
Aspen	2 (26, 19)	1 (63)	2 (28, 18)	4 (2, 28, 6, 3)	1 (10)	3 (2, 2, 3)	13 (210)
Sapling	2 (30, 18)	2 (150, 120)	1 (37)	4 (34, 29, 68, 21)	1 (52)		10 (559)
Sassafras	2 (29, 11)	4 (29, 53, 35, 37)	1 (81)	3 (27, 15, 17)	3 (33, 30, 17)	1 (12)	14 (426)
Pine	1 (35)	1 (29)	5 (33, 61, 20, 33, 26)		1 (28)		8 (265)

quite house meeting status of importance, but we haven't used it for things that we would normally bring to a house meeting." By contrast, Pine used it for some more serious decisions, where members reported that Resolve was helpful in clarifying different members' views. Members of four of the groups reported very positive outcomes and satisfaction with their experience with InterTwinkles, as well as intentions to continue using it in the future (Pine, Sassafras, Maple, Sapling), members of one reported negative outcomes (Aspen).

One group (Hemlock) did not use the tools after the initial training; they are thus omitted from this table. During the study period, they had two in-person house meetings, but did not have major decisions to make. The group began as the smallest group in the study (5 members); this was reduced further by one house member moving out, and another's travel, for a total of 3 people who felt they could have participated. In a closing questionnaire which asked if any of six potential reasons³ for not adopting InterTwinkles held, two members chose "momentum prevented us from trying something new", and one chose "I didn't have a need for InterTwinkles". One member commented, "I like the idea of it and hope to use it in the future if we had more people making decisions or a complex decision to make," implying an interpretive schema that InterTwinkles is primarily useful for difficult or complex decisions, rather than simple communication tasks. Both in this case and in the case of another group with traveling members (Aspen), groups interpreted travel as a reason for exclusion from participation in house decisions.

For the most part, tool usage was faithful to the design intent of the tools (that is, the types of activities people engaged in with the tools matched the tools' design intentions). The exception is Firestarter, which was both the first tool listed in the software interface and also the simplest in function. While the tool was originally intended to be used for introductions and go-around format discussion (e.g. "pass the conch"), groups used it for considering proposals as well as unstructured brainstorming. This may be a result of several different factors – including the unrestrictiveness of the tool, and a generally low degree of recall of the names of tools.

In addition to their usage of InterTwinkles with study groups, members of two groups introduced InterTwinkles to two other

Table 5.3: Activities completed for each tool by each group. The first number is the number of activities completed with that tool; the parenthetical lists the number of events (all edits and unique visits) by members of the group to each completed activity. (e.g. Maple completed 5 activities with Resolve, which logged 80, 61, 29, 20, and 11 unique visits and edits respectively). Progressive Clock was the least used tool; other tools show similar levels of aggregate use, though different groups did more with particular tools.

³ The closing non-use questionnaire proposed the following possible reasons for non-use:

1. I didn't have a need for InterTwinkles
2. I had a need, but InterTwinkles didn't satisfy it
3. I didn't understand InterTwinkles adequately – needed more training.
4. I didn't trust InterTwinkles.
5. Momentum prevented us from trying something new – existing strategies worked well enough.
6. I didn't have time to try something new.

groups of which they were members. One member introduced Dot-storm to a maker space with which she works, describing the value in the drawing/text combination for discussing designs. The other introduced Resolve to a co-working space, hoping to use Resolve to help with decisions. This indicates a fairly high degree of trust in the value of the tools.

How did tool usage relate to the “spirit” of the design?

The “spirit” of a technology is the set of considerations and aspirations that designers have sought to impart in its structure and function, which they hope will be reflected in organizational usage. Once used in practice, the actual outcomes and emerging structures in a group may be very different, as the technology could be used in unintended ways. Relating the spirit to the incumbent structures of the groups (such as those summarized in table 5.2) can help in interpreting outcomes.

Based the taxonomy of dimensions of spirit in DeSanctis and Poole [1994], table 5.4 shows InterTwinkles’ approach.

Decision process	Consensus; non-participation is apparent
Conflict management	Increases awareness of conflict
Leadership	Non-hierarchical; facilitation out of band
Efficiency	Increases efficiency, organizes data
Atmosphere	Easy, colorful, intuitive, emotional; structured and formal

Table 5.4: “Spirit” of InterTwinkles’ tools

Two dimensions in particular seem to relate to the reported outcomes from interviews: the degree of formality or structure in the consensus process used, and the approach the groups have to conflict management. Formality here means the extent to which the group has a culture of regularly using policy or protocol-driven practices in their meetings: as an example, members of Sassafras always begin their weekly meetings with a check-in, followed by a pre-planned agenda, followed by a check-out, with formal tests for consensus and proposals. By contrast, members of Aspen only hold meetings on an ad-hoc basis as needed, and just informally discuss issues.

The approach to conflict also varies considerably: members of both Aspen and Hemlock describe an avoidant approach to managing conflict. Aspen had previously held more formal meetings, but intentionally stopped holding them due to conflict: “For about a year, we had a housemate where there were frequently conflicts during house meetings, and that’s part of why we have fewer house meetings now, because after a certain point, it became really stressful to have this monthly meeting in which there was lots of conflict.” Similarly, Hemlock describes a strategy for dealing with conflictual

decisions of deferring a decision until the issue either becomes acute or resolves itself in other ways. Aspen describes this attitude as a “do-ocracy” – privilege is given to those who take action, though the group will address and resolve conflicting action if it arises.

Figure 5.5 maps each of the participating groups according to these axes.

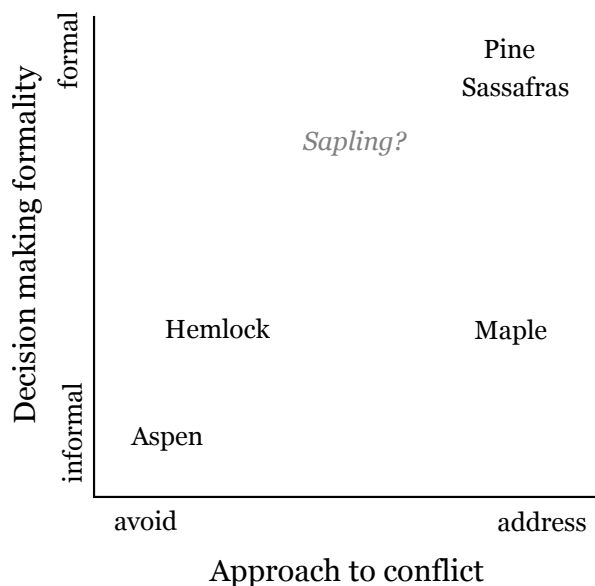


Figure 5.5: Based on pre- and post-interviews of each group, this is a subjective assessment of the degree of formality or informality of the groups' incumbent decision-making practice (the vertical axis), and the degree to which the group seeks to avoid conflict or to directly address it (horizontal axis). *Sapling?*, as a new group, has not yet strongly institutionalized approaches to conflict and decision making, so its placement is less clear.

While clearly this is too small a sample to infer causation, it is notable that a design intention of some of InterTwinkles tools (particularly Resolve, Points of Unity, and Progressive Clock) is to make conflict visible in a structured, formal way – something which has the potential to increase the criticality of that conflict, ‘worrying the sore’. For groups like Sassafras, which have very explicit structures in place to deal with conflict, this can be helpful; however, for groups that have prioritized less introspective approaches, this could be harmful, making big issues out of smaller annoyances that might be better ignored. The “garbage can model” of decision making (discussed in chapter 3, page 66) raises this critique on meetings.

A notable difference in handling conflict between the members of Maple on the one hand, and Sassafras and Pine on the other, is the group's attitudes toward whether “silence implies consent”. Two members of Maple both reported a group norm that the whole house would be informed of decisions by Maple's committees, but if the house members failed to speak up with concerns, the decision would be considered valid. By contrast, members of Pine and Sassafras both reported that their groups saw the equation of silence and consent as problematic, a potential source of coercion. Nevertheless, all three groups reported positive experiences using Resolve to decide issues

– in the case of *Sassafras* and *Pine*, the group considered non-voting members to be a problem to be addressed, whereas members of *Maple* were comfortable ignoring them.

What do we know about moments that impelled decisions to use the tools?

A critical moment in the usage of any structured communication tool, whether an in-person protocol or online, is the moment in which a member of the group decides to use the particular tool and asks others to participate with them. When one is proposing the use of some process, one must have a level of expectation that other members of the group will be sufficiently prepared (or sufficiently willing to learn) that it can functionally use the process. If the process is complex and requires a fair degree of knowledge on the part of group members to use the process (e.g. Robert's Rules of Order, or software that requires logins and training to use), one must expect that other group members are trained, or that it would be possible to train them sufficiently in time.

In asynchronous contexts where one doesn't have body language from other group members to gauge comfort with a process, asking explicitly whether people are willing to use a particular new process via a more familiar medium (like email) can be more effort than just skipping the new process and using the familiar medium to take on the task, even if the old process seems less optimal. *InterTwinkles'* training workshops were designed to provide the necessary training and setup so that groups could easily use the tools, and so they would *know* all group members were similarly trained. *InterTwinkles*, however, cannot impel the group to action: it can provide a structure for communication, but human users still need to provide the inspiration and motivation to use it.

Groups found a variety of strategies to initiate activities with *InterTwinkles* – but generally, the workflow consisted of one member creating an activity, and then emailing a link to the group. A member of *Maple* reported that they would discuss an issue in person with people who happened to be in the room to “get something rolling”, and then send an email to the house with a link to the activity. This process closely mirrored *Maple's* committee-based decision making structure – but using a more structured channel for feedback from the wider group, rather than just email. A member of *Sassafras* reported that particular members who were taking on facilitative roles would be more interested in using the tools, and would initiate activities, then email links to the activity to the house. Similarly, one of the early organizers of *Sapling* mentioned that he would start an activity, email the group about it, and then nag members who

weren't responding: "I would send out another email saying, 'I really need everyone to respond to this thing.'" Users reported a great deal of satisfaction when they had started something, and saw others' contributions; but also frustration when others did not respond quickly.

Personal relationships were also a factor contributing to people's decision to use tools. A member of Aspen (which reported generally negative outcomes with InterTwinkles) said that she used the tools largely because of her friendship with another member of the house who was a research assistant working on InterTwinkles. In an early training session with Pine, one member remarked to me that she wanted to be sure the group used the tools well enough for my needs in a research study – I tried to assure her that they should use the tools to the extent that they found them useful, and not on account of my need to complete a study, as my study would be complete whether or not they decided to use the tools. Nevertheless, this aspect of the native participation of me and three research assistants in these groups likely contributed to the groups' decisions to adopt them. Groups may wish to leverage this effect by encouraging their members to develop bespoke tools, which will see higher potential as effective tools for coordination due to personal relationships with the designers.

Hemlock, which did not end up using the tools beyond the study period, cited issues with the authentication system as a factor contributing to their non-use. Two members of the house were away, and two subletters did not have logins or membership in the group's InterTwinkles account. Though InterTwinkles was designed with the intent that authentication remain optional, this was perceived as a barrier. Members of several other groups commented on problems with authentication: a member of Sassafras mentioned that she did not use the website directly (but instead asked others to take actions on her behalf) due in part to the hassle with logging in – though this also indicates success in our explicit design choices which made proxy usage possible. A member of Aspen expressed concerns about privacy and her increased "digital footprint" based on authenticated use of the site (she also expressed particular concern over the use of Mozilla Persona rather than a site-specific authentication system).

The one example we have of a group adopting an explicit norm around how to use InterTwinkles was Pine. They developed a prescriptive workflow that integrated InterTwinkles with their current practice – and due to the partial overlap in membership between Pine and Sassafras, Sassafras members picked up similar practices:

When a member starts a new activity on InterTwinkles, they should email the group using the mailing list to announce the activity, pur-

pose, and when people need to respond by. That member is then responsible for finalizing and follow-through.

It would likely be beneficial for future training sessions and documentation to include explicit processes like this as recommendations for better integration with existing systems.

How do the groups interpret the tools now?

As a way of encapsulating groups' overall interpretation of InterTwinkles, we asked interviewees from each group to characterize how they would describe what InterTwinkles is to new housemates or members of other co-ops. They responded:

- **Maple** "InterTwinkles are good tools for groups to use, but you have to put time into learning how to use them. Every tool has a learning curve, and it's a short one for InterTwinkles, but in a group, everyone has to learn the tools at their own pace in their own way, so it can take a while for the group to be together on the same page, the same level."
- **Aspen** "We have a former housemate who's really interested in things like policy and verbal co-op-ness and non-violent communication and all of that, and if I was talking to her about it, I would definitely have this proposition: 'Oh this is this online tool to help facilitate all of the formal decision making processes ... that come with formal consensus based co-ops.' ... [For others] it would probably be something like, 'Here's this tool, it has some options for drawing ... it has this open Etherpad option ... these two different decision making options and if you want to try to use it, go for it. I'll be supportive of you, go for it.'"
- **Sapling** "I would probably walk them through, I would probably want to be there with them saying, 'Okay, let's go through, we're going to make your account, we're going to set that up. Okay, here's where we are. Here's things that we're working on and if you wanted to start something you would click this,' and I would show them."
- **Sassafras** "InterTwinkles is a tool we <laughs> sometimes use when we remember it exists to deal with house stuff, and it's online and you login. Everyone is part of the [Sassafras] group, and we can do things like brainstorm ideas to bringing proposals that we make decisions about."
- **Pine** "InterTwinkles is an online set of tools for making decisions by consensus, so it's something we can use when we're not able to

meet or sort of in the interim between meetings; and also a cleaner way for making decisions and actually being able to visualize them better or in a different way; and also ... consolidate the information in a place about different decisions.”

User feedback

In this section, I will summarize some of the key areas of design and usability feedback from the study.

One issue that became apparent from several of the groups is that name comprehension for the tools was very low. The names were chosen both because of their relationship to the activity and for their uniqueness and evocativeness (e.g. “firestarter” was chosen over synonyms like “ice breaker” and “go arounds”, even though it is a less common term for the activity). However, I suspect that naming it with a more mundane name (“go arounds”) would have resulted in more comprehension. Users expressed confusion over the other names as well. One user commented, “I don’t like the names – they don’t say what they do. I know it’s hard to pick unique names, but just use something simple like ‘draw’”. This confusion extended beyond simple recall: when faced with the menu of tools and a task to complete, users were unsure which tool to pick, even if they had familiarity with the basic function of each of the tools. With that in mind, the tools could be renamed:

- Dotstorm -> Brainstorms, or Sticky Notes
- Resolve -> Proposals
- TwinklePad -> Documents
- Firestarter -> Go arounds

This difficulty speaks more generally to the strategy of small, composable, structured communication tasks. In order for a group to benefit from using a light weight tool for a simple task, they must have shared awareness of the affordances and limitations of the tool. The palette of techniques available for face-to-face groups are not all learned at once, but practiced over many years of experience in facilitated group discussions – similarly, it may take time and training for groups to be able to productively recall the use cases of different tools. From an experimental perspective, we would benefit from a more ubiquitous platform which takes care of authentication needs and group modeling, but which allows groups to establish a base of familiarity in some well-known tools, and experiment with others occasionally.

Fragmentation of communication channels was an issue for some users. A member of Aspen commented: “A lot of those conversations ... were happening simultaneously on InterTwinkles and on email, and people were in different places at different times so it was hard to tell.” Worse still, where users failed to perceive any value in the structure provided by the tools (as in Aspen’s case), this made the tools more harmful than beneficial: “InterTwinkles isn’t doing anything that email isn’t already doing.” Beyond fragmentation of channels, in Sassafras, InterTwinkles exposed a lack of certainty among some members of the house about what subjects were appropriate for house meetings, versus uncoordinated action. “I think in [Sassafras] there’s confusion about things that don’t really require being brought to a house meeting and consented over, but still need some input from everybody.” This ambiguity of which venue was appropriate for what type of communication, and whether the communication was required in the first place, was made more acute by the presence of an always-available “meeting” channel.

The perception of InterTwinkles as a tool capable of addressing urgent needs was mixed. A member from Maple believed that for a very urgent matter, “We would likely get faster responses from our email list.” On the other hand, for issues that required explicit procedural action (such as a decision that was required by Pine’s bylaws), members felt that the ability to efficiently collate responses in one place was more effective than email – and ‘it allowed us to talk about other things at our [Pine] meeting’ instead. Similarly, a member of Sapling found InterTwinkles very useful for helping to handle the large amount of work in forming a new group:

Yes, we have regular in person meetings. So we were using this sort of in between our meetings, because as we were forming our meetings, so much time was taken up with interviewing new people. We literally didn’t have time to sit around and hash some of this stuff out. And we did, we were having three and a half hour meetings. We were just in sort of a rush to define ourselves in order to be able to do our fundraising and communicate clearly with the outsiders that we needed to get through this stuff pretty quickly. And it really was helpful for that.

A member of Maple commented that they really appreciated the design and intention behind InterTwinkles, and saw it as a better fit for the way their group operates. While Maple had used many online tools for group use in the past (including Google Docs, Survey Monkey, etc.), he felt that they were forced to bend what they did to fit those tools, whereas InterTwinkles was a closer match: “We’ve been forcing the way that we do things as a house to fit into what technology can offer us. And I wanted you to know that I see how you have created something with this technology that is built around

how people actually make decisions.”

Ethical fit, and analysis of control

Members of each of the study groups – and members of the target communities for which InterTwinkles was designed in general – hold strong values for autonomy and democratic control. Indeed, as discussed in chapters 1 and 2, values of consensus decision making and non-hierarchicalism can run contrary to the interests of centralized online service providers, who hold the power to engage in surveillance or withhold service. For group-oriented software, this problem is more acute, as changing service providers requires more than just individual choice – it requires changes to group momentum.

In chapter 1 (page 27) I provided an analytical framework to consider the dimensions of control in decentralized systems. I will now apply that framework to InterTwinkles, to examine the ways in which the project has succeeded in ceding control to the user groups, as well as the areas where more work is needed to develop community control. This framework considers three dimensions (Access, Evolution, and Vision) through each of four analytical lenses (technical capabilities, rules and protocols, exclusive meanings, and structural inequalities).

*How do **technical capabilities** impact who participates in...*

Access: Currently, InterTwinkles is free to use, and open to anyone on the web. It is designed to function on modern web browsers running on any device (phone, tablet, desktop, etc.). InterTwinkles is open source, and designed to be easily replicated. The codebase can be installed by any sysadmin with basic unix skills. Currently, InterTwinkles is only running at <https://intertwinkles.org>, however it is designed with the intention of supporting many other installations.

Evolution: So far, I am the sole developer of the entire InterTwinkles code base (not counting 3rd-party libraries on which it depends) – this is a problem, and one which would be fixed by attracting more developers to the project. The danger of a single developer is that the codebase will reflect the idiosyncrasies and mental models of that one developer, making contributions from others more difficult.

During development, I made several design choices to make participation by developers easier: each of the tools is designed as a plugin to the core platform. A plugin system enables developers to take on a particular component without having to learn the entire codebase. However, this plugin API is still not well encapsulated and entirely undocumented. Documentation and isolation of this API

will make development of modules substantially more likely. The system also has an installation script that installs and configures all necessary components in one operation.

InterTwinkles is developed in the open using git as a revision control system, and hosted publicly.⁴ Anyone can download and begin editing the code immediately. There are many members of the target user groups who have the technical means to help develop InterTwinkles. However, the lack of their participation indicates a deficiency in efforts to successfully decentralize evolution.

⁴ Code is available at <http://github.com/yourcelf/intertwinkles>

Vision: An InterTwinkles project site (<http://project.intertwinkles.org>) and mailing list were set up early in the project's life to facilitate collaboration and input around project directions, including questions of access. However, this site has languished, and the mailing list is unused. Both should be developed further to allow input from user groups on questions of access.

The InterTwinkles site integrates with UserVoice for feedback (both for authenticated and anonymous users). However, this integration has been demonstrated to have usability issues – people have tried to use it ask tech support questions, but have done so in ways that prevented us from responding to them.

*How do **rules and protocols** impact who participates in...*

Access: Users can authenticate if they wish, but they don't need to do so if they only wish to access activities that are not private. If they authenticate, InterTwinkles uses Mozilla's Persona service for authentication – a protocol that is still new, but which is designed to ultimately allow federated authentication without loss of privacy. Nevertheless, the use of Mozilla's Persona currently requires that users accessing InterTwinkles place some trust in Mozilla.

Using software as the source for protocological structure in a group presents a challenge not present in low-technology alternatives: while anyone with little training can replicate, modify, adapt, and alter common low-tech facilitation techniques, software requires a degree of expertise to modify.

Evolution: InterTwinkles' software is a freely licensed, which is a minimum first step to the ability of others to participate in production. The code also includes a full test suite with continuous integration testing, making it easier for others to contribute – one rung higher than the minimum. There is currently no guide for contribution, no documentation of a development workflow, nor recommendations for the use of issues for bug reports. All of this documentary material might assist other developers in contributing.

Vision: InterTwinkles began as a participatory design project, including participants from several groups who were target users for the project. However, as described in chapter 4, development moved to an iterative model due to limitations in the time and capacity of target groups to participate. It is not too late for InterTwinkles to re-connect with communities of users, developers, and thinkers about online decision making practices to engage a more diverse leadership.

InterTwinkles design decisions have largely happened at MIT in the context of this dissertation. Exceptions to this include design meetings hosted at community meeting spaces in 2012, and a more recent design meeting in July, 2013. It would be beneficial to add components regarding the vision for InterTwinkles to public project websites.

To date, visioning has been guided by the constraints of academic work. Post this dissertation, it could take alternate directions. There are currently no protocols beyond the dissertation process for evaluation, planning, and direction setting process for the project.

*How do **exclusive meanings and practices** impact who participates in...*

Access: The documentation about how to use each of the tools is inadequate, and the tool's structure relies on groups to have effective out-of-band facilitation techniques that are not explained within the tools. This means that effective use of the tools is likely to be difficult for groups that have not participated in the training workshops – their effective use is an illegible, exclusive practice.

The design of InterTwinkles site incorporates language about consensus, activists, and affinity groups, even though there is nothing that structurally requires groups to use the tools as part of a consensus process. This is likely to make the tools less attractive to groups that use majority voting or other systems for decision making. The design also strives for a visual aesthetic of warmth and cuteness, with lots of pink. This may be unwelcoming to people who would prefer a more austere aesthetic experience (such as those in traditional corporate contexts).

Evolution: InterTwinkles currently follows standard development practices for small open source projects, including publishing a public repository, maintaining a test suite, and keeping a public bug tracker. However, but this still leaves much to be desired to make the code accessible to non-developers. The codebase is primarily written in a somewhat esoteric language (one that might be considered somewhat “hipster”), which will limit accessibility to developers who

either know the language or have the time and inclination to learn.

Vision: As InterTwinkles lacks governance, there is no mechanism in place beyond my desires as lead project developer for mechanisms of planning and evaluation. Consequently, mechanisms of participation in development of vision for InterTwinkles are likely to be illegible and inaccessible to others.

*How do **structural inequalities** impact who participates in...*

Access: The requirement for devices supporting modern web browsers limits InterTwinkles' usability to those who can afford modern computers or phones with data connections.

Evolution and Vision: While InterTwinkles began as a participatory design project, with concerted efforts to involve people with all levels of computer skill and from a diversity of backgrounds in the processes of ideation, planning, and early development, as the project has progressed, it has become less open. As a result, the visioning process is currently largely driven by me (a rather privileged individual), and does not currently have mechanisms for involving more people. While I don't intentionally structurally exclude others, if the project grows, it will be necessary to establish mechanisms of governance early on to ensure that it maintains the goal of working against systems of structural inequality.

In summary, this evaluation of control demonstrates that InterTwinkles has ceded control to groups at the most basic level (sysadmins could download and install the software; people who trust our installation can use it); however, much work remains to be done in developing the project infrastructure to make the code base accessible to other developers, and to develop a shared process for project visioning and planning. As it stands now, I have an uncomfortably high degree of control over users of this tool set as a sole developer and project lead – control which could be abused by forces beyond my control.⁵

⁵ Colin Moynihan. Tad hirsch, media lab grad. student, subpoenaed for users of TXTmob, system employed by protesters. *The New York Times*, April 2008. URL <http://tech.mit.edu/V128/N15/txtmob.html>

6 Conclusion

In this project, I set out to address the question: *How can we design effective online tools for democratic consultation in non-hierarchical, consensus-oriented groups?* As background to this design work, I explored the meanings of non-hierarchicalism and consensus, and the history of computer systems that aid decision making. I then designed, implemented, and tested a platform of online tools.

I developed an analytical framework to help explain how control operates in decentralized systems. Many groups aspire to be inclusive, open, and non-hierarchical – just as many software projects strive to promote freedom. But the means by which they are “free” or “open” are often limited to naïve notions of access or intention, rather than deeper questions of whether they actually succeed in maintaining inclusive processes for use, development, and leadership. I believe that the framework I developed can help groups and system designers to dig deeper into the difficult territory of challenging systems of structural oppression, grappling with tension between identity and exclusivity, and recognizing the role of technologies, procedural rules, social meanings and practices in determining outcomes.

To improve clarity around what consensus-orientation entails, I developed a taxonomy of different forms of consensus. The different expressions of consensus vary enough in scale, membership, and outcomes that they warrant different designs for tools to support their operation. I zeroed in on *affinity consensus*, and explored the motivations that drive its use and the protocols by which it functions. As a set of protocols, the practices of affinity consensus are one possible solution to the problem of choosing an equitable way to structure decision making processes in non-hierarchical groups.

After analyzing the history of efforts to support group decision making with computers, I developed the working hypothesis that the majority of past academic work in the field had failed due to a lack of attention to usability and human-centered design. While hundreds of studies were completed in the 1980’s and 1990’s with early “Group Decision Support Systems”, these systems were not

designed with participation of users who were asked to use them, and were not designed with the needs of extant groups in mind. Furthermore, the studies were primarily conducted from a frame of positivistic lab research rather than field study. Looking at the trend of contemporary successful projects toward simple, tightly focused systems that strive to solve narrowly defined problems, I adopted this design ethos for tools in InterTwinkles.

Through a participatory and iterative design process, I developed a suite of tools and a platform that connected them. An inherent tension between user needs for an integrated tool set and a desire for atomic and tightly focused tools led me to develop an integrated platform with a plugin architecture. The resulting tools support a range of techniques found both in previous academic work in decision support systems, as well as meeting techniques found in groups that practice consensus.

Six existing consensus-oriented groups participated in a field study over the course of three months, using the tools electively for endogenous tasks. The results of the study were largely positive for four of the participating groups, and less so for two of them. The success of the tools seems to track closely with the degree of fit between the design intention for the tools and groups practice: the groups which used more formal decision making processes and confrontive approaches to conflict (an intention shared by the design of the tools) had better success. This result demonstrates that there may be more work to be done in tuning tools to the needs of groups with different approaches and styles.

Future directions

Beyond “Decision Making”: One thing that has been increasingly apparent to me in the course of this research is that the frame of “decision making” as a description of group activity may have outlived its usefulness. The practices that take place in meetings are diverse and complicated enough that the only real commonality between them is the need to have bodies proximal in space and time. Once we remove that requirement, and develop systems that work asynchronously, we might find that it is less useful to group all of those activities under a single banner. Just as “computer aided communication” has now been splintered into email, instant messaging, chat, video conferencing, Internet telephony, message boards, and more; we might find that “decision making” will retain the basic notion of *consultation*, but will be described and conceived of as a wider variety of more task-specific activities. The language of decision making may be actively harmful to this effort, both by constraining developers’

conceptions of what they are designing for, and by constraining users' interpretations of how to use the systems.

It would also be valuable for future tool design to begin to specifically target ways to build groups' awareness of practices that make them more exclusive, and ways that they might be replicating structural inequalities. While InterTwinkles scratched the surface of this, there is much more work to be done. A "progressive stack" tool could assist with real-time meetings. A facilitator rotation log could help groups manage a rotation of leadership roles, nudging groups toward more effective shared leadership. A "conflict communication" tool could allow group members to consider and share how they prefer to communicate in conflict situations (and to prime them to consider that communication about conflict is necessary). Policy-oriented tools could help groups to craft useful policies, and audit whether written policies are actually functioning. In addition, further development of a notion of the *quantified group* could help groups to reflect on structural inequalities or exclusive meanings that they embody – for example, analysis of group mailing list traffic could reveal communication dynamics that reflect structural inequalities. All of this work would likely benefit from escaping the frame of "decision making" to allow a much richer variety of tools that help groups to reflect on their own practices.

Group training: Currently, InterTwinkles relies very heavily on in-person workshops for training of groups. It would be valuable to develop tools for in-band training, so that groups could increase their aptitude with tools outside of workshop contexts. This work could be expanded to share prescriptive models of organizational forms and decision making strategies that other groups have found to be successful. For newer groups, a platform could be developed as a "governance-in-a-box" which provides assistance with setting up shared values, developing norms for communication, and mechanisms for shared leadership. The insights developed in this dissertation around the structure of different types of groups could be reflected back in a palette of organizational forms that groups might choose.

Ecosystem of development: To be successful, InterTwinkles will need to diversify development to a wider array of contributors. One particularly promising direction is to work on integrating with other online platforms rather than continuing to build walled gardens. An initial step would be to develop tools to auto-populate groups from their existing mailing lists, or from other online community tools. However, a more robust approach would be to work with other providers to develop a federated mechanism for sharing group membership among different platforms – essentially, an "OAuth

for groups”, where the agent sharing data is a group rather than an individual.

The systems for ongoing feedback, design consultation, and support built into InterTwinkles need considerable attention. For the initial development, we used UserVoice (a commercial provider that offers a simple feedback and support mechanism) as an available integrated feedback tool – but this has proven to have substantial usability issues, including people attempting to use it to ask support questions, but in such a way that we are unable to respond to the people who tried to reach out. A more complete solution would include mechanisms to include users in planning processes for design and development.

Autonomy and security: It is clear that there is currently no good solution for data security in an era of PRISM and XKeyscore. The challenges of maintaining the rigorous practices necessary for even two-way encryption (as with PGP) compound exponentially when communicating with a group. There are neither adequate technical systems nor user skill to encrypt group data such that it is invulnerable to subpoenas or National Security Letters served to server operators. As a result, the only option available to developers to protect users is to encourage them to either install and maintain their own servers (and hope they can maintain the security of those systems and resist subpoenas), or to recommend that they stay off the Internet when engaging in politically contentious activity. Developing viable crypto-systems and user education for group-based communication is an important and needed area for future work.

To reduce the barrier to running your own instance, InterTwinkles could be developed into a stand-alone installer usable by unsophisticated computer users, which runs on a local wifi network. This could allow groups to use the tools without a sysadmin or access to their own server, and without needing to trust a central service provider to have successfully avoided government coercion to share data.

*

**

This work has been an effort to merge multiple areas of deep interest to me personally into an integrated project – including the cooperative housing community, activism, community organizing, free software, social justice, anti-oppression, and philosophy. After years of work, I still feel that it is only at the beginning. My hope is that the analysis, source code, and experiences resulting from this work will help to inspire renewed efforts to work toward our collective liberation. The activist website riseup.net used to have a tag line “Get off the Internet, I’ll see you in the streets.” In my lifetime, I hope to see this revised to: “Whose Internet? Our Internet!”

7 *Acknowledgments*

I would like to express my deepest gratitude to the following people and groups, without whom this project would have been impossible.

Thanks to Mary Regan, Josh Beckmann, Noah Swartz, and Kendra Moyer for putting enough faith in me to devote some of your vocational livelihood to this. I hope I was a fair and friendly boss, even though we didn't operate non-hierarchically.

Thanks to the six brave co-ops who agreed to participate in this project. You put your trust, time, and energy into supporting my work, for which I'm grateful.

Thanks to Eric French Monge for joining me in the earliest stages of conception and development of what became InterTwinkles, and to Becky Hurwitz for early moral support. And the rest of the codesign 2012 seminar who helped trial early prototypes.

Thanks to all of the groups and people who participated in design workshops in Boston, Providence, and elsewhere. Particular thanks to North American Students of Cooperation for inspiring me to dedicate myself to non-hierarchicalism and consensus, and the Boston Collective House Assembly, whose members entertained my crazy thoughts about cluster organizing and online tools.

Thanks to all the members of AORTA – in particular kiran nigam and Esteban Kelly, whose deep experience and valuable insights into group work have enormously influenced this work and my outlook on life.

Thanks to my adviser, Chris Schmandt, for putting up with me for these 6 years. Without your trust and support I would never have been here.

Thanks to Sasha Costanza-Chock for advising this project from its humble beginnings in your codesign class. And especially thanks for helping to keep the social justice frame front and center throughout the work – a much needed booster to ward off techno-centric positivism.

Thanks to Sep Kamvar for helping me to break out of the narrative ruts that I forged myself into while writing the dissertation.

Thanks to Ethan Zuckerman and Lorrie LeJeune for supporting

me as a full member of the Center for Civic Media, administration be damned. I felt proud to be a member of the group; you made me feel welcome there. Thanks also to Chris Csikszentmihalyi, who got me started on this path.

Thanks to Ricardo Dominguez for helping me to think past the boxes of traditional research warrants, and inspired an attitude toward art that I still strive to emulate.

Thanks to Denise Cheng and Jeffrey Warren who selflessly donated their design skills to help InterTwinkles get its distinctive look and feel.

Thanks to Molly Sauter, Erhardt Graeff, Chris Peterson, Willow Brugh, and Ed Platt for the best kind of critical feedback throughout the project. Late night consults around the civic table were a joy with your presence. Thanks to Nathan Matias for that plus invaluable help in designing the Flame War card game and thinking about server architectures. Thanks to Drew Harry, Finn Brunton, Benjamin Mako Hill, Jérôme Hergueux, and Andrei Boutyline for valuable help with the discussion of ideas and critiquing drafts.

Thanks to the Speech + Mobility group members Andrea Colaço, Misha Sra, Hye-Soo Yang, Wu-Hsi Li, and Jaewoo Chung for helping with practice talks and critiques.

Thanks to Carl McLaren and Benjamin Sugar for holding down the fort at Between the Bars while I was too preoccupied to help out.

Thanks to my love Laura Burkle for putting up with me during the trials and tribulations of dissertating, as well as lending your adviserly insight and critical reading.

Thanks to Voqal and the Media Ideation Fellowship for believing enough in this work to give it financial support and mentorship.

A Consensus flowcharts

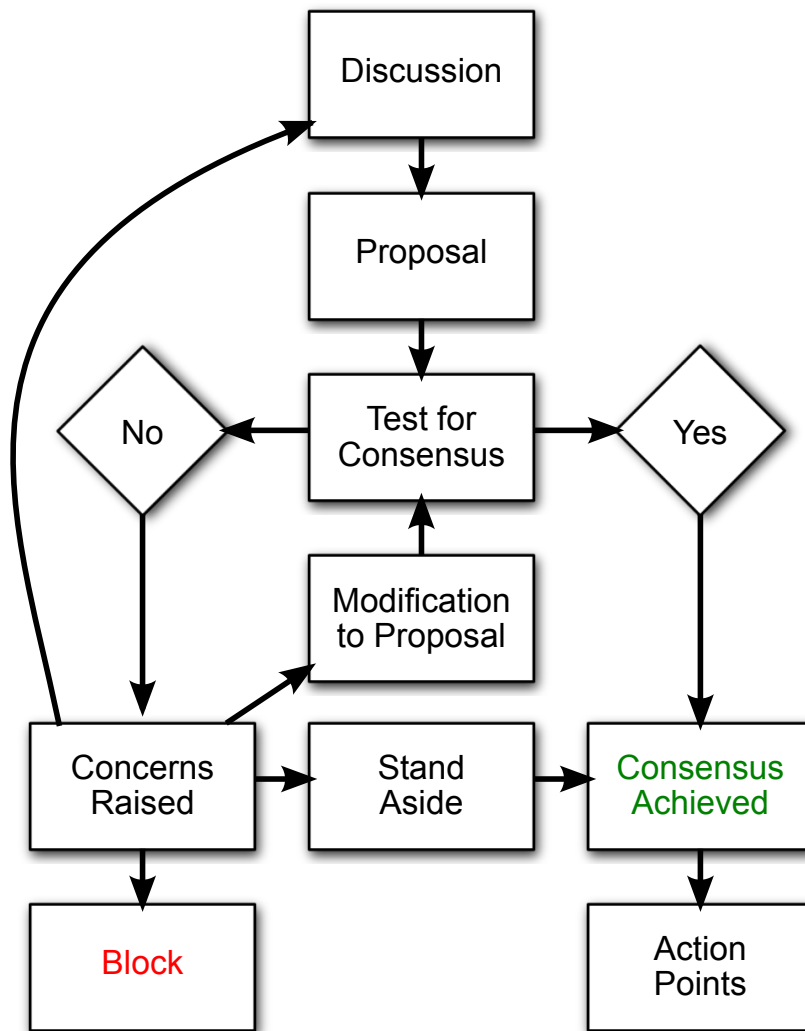


Figure A.1: A fairly standard flowchart for the process of consensus decision making, from https://en.wikipedia.org/wiki/Consensus_decision_making.

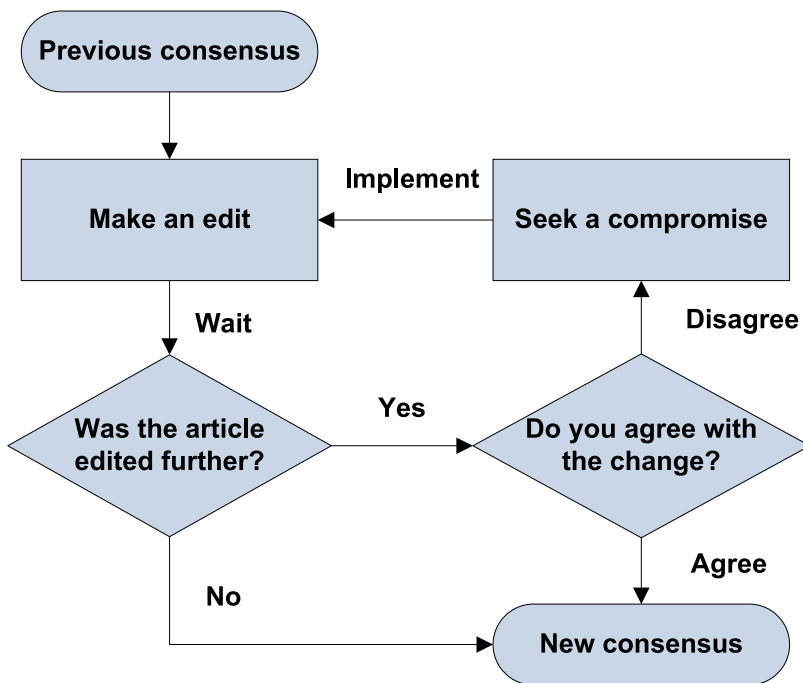


Figure A.2: A flowchart of the process Wikipedia uses for consensus around articles. This differs from many processes in that it explicitly acknowledges asynchronous time, as well as a changing cast of editors – after waiting a “reasonable” amount of time, if an article wasn’t reverted, the assumption is that new consensus has been reached. Under normal process, there is never any explicit ritual of approval that indicates consensus has been reached. From <https://en.wikipedia.org/wiki/Wikipedia:Consensus>.

NASCO BOARD MEETING PROCESS

A work in progress, Feb. 2007

Overview: The following approach to facilitation and decision-making outlined in this document is a modified consensus process which came out of the June 2006 NASCO Board meeting after a good deal of discussion and experimentation with different meeting processes. The process outlined below is something of a hybrid, which attempts to combine some of the clarity and formality of Roberts Rules with the participatory and flexible nature of consensus processes.

Guidelines for this NASCO Board meeting process are broken down into:

- Decision-making process flow chart
- Overview of facilitation roles
- Additional guidelines

DECISION-MAKING PROCESS FLOW CHART:

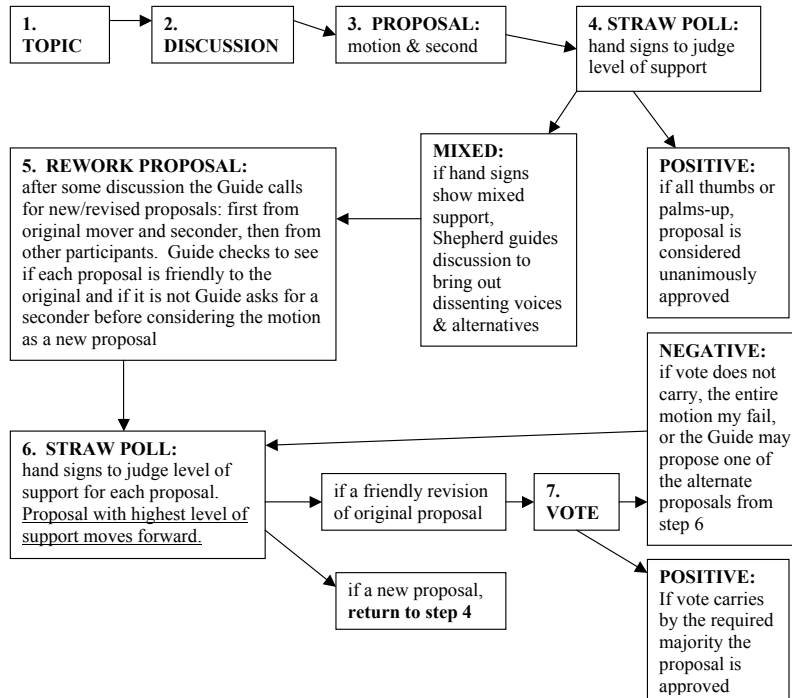


Figure A.3: The decision making flowchart used by the North American Students of Cooperation (NASCO, <http://nasco.coop>), a non-profit membership organization which supports housing cooperatives across the US and Canada. This process starts with an effort at building consensus, but with majority voting as a fallback when consensus can't be reached.

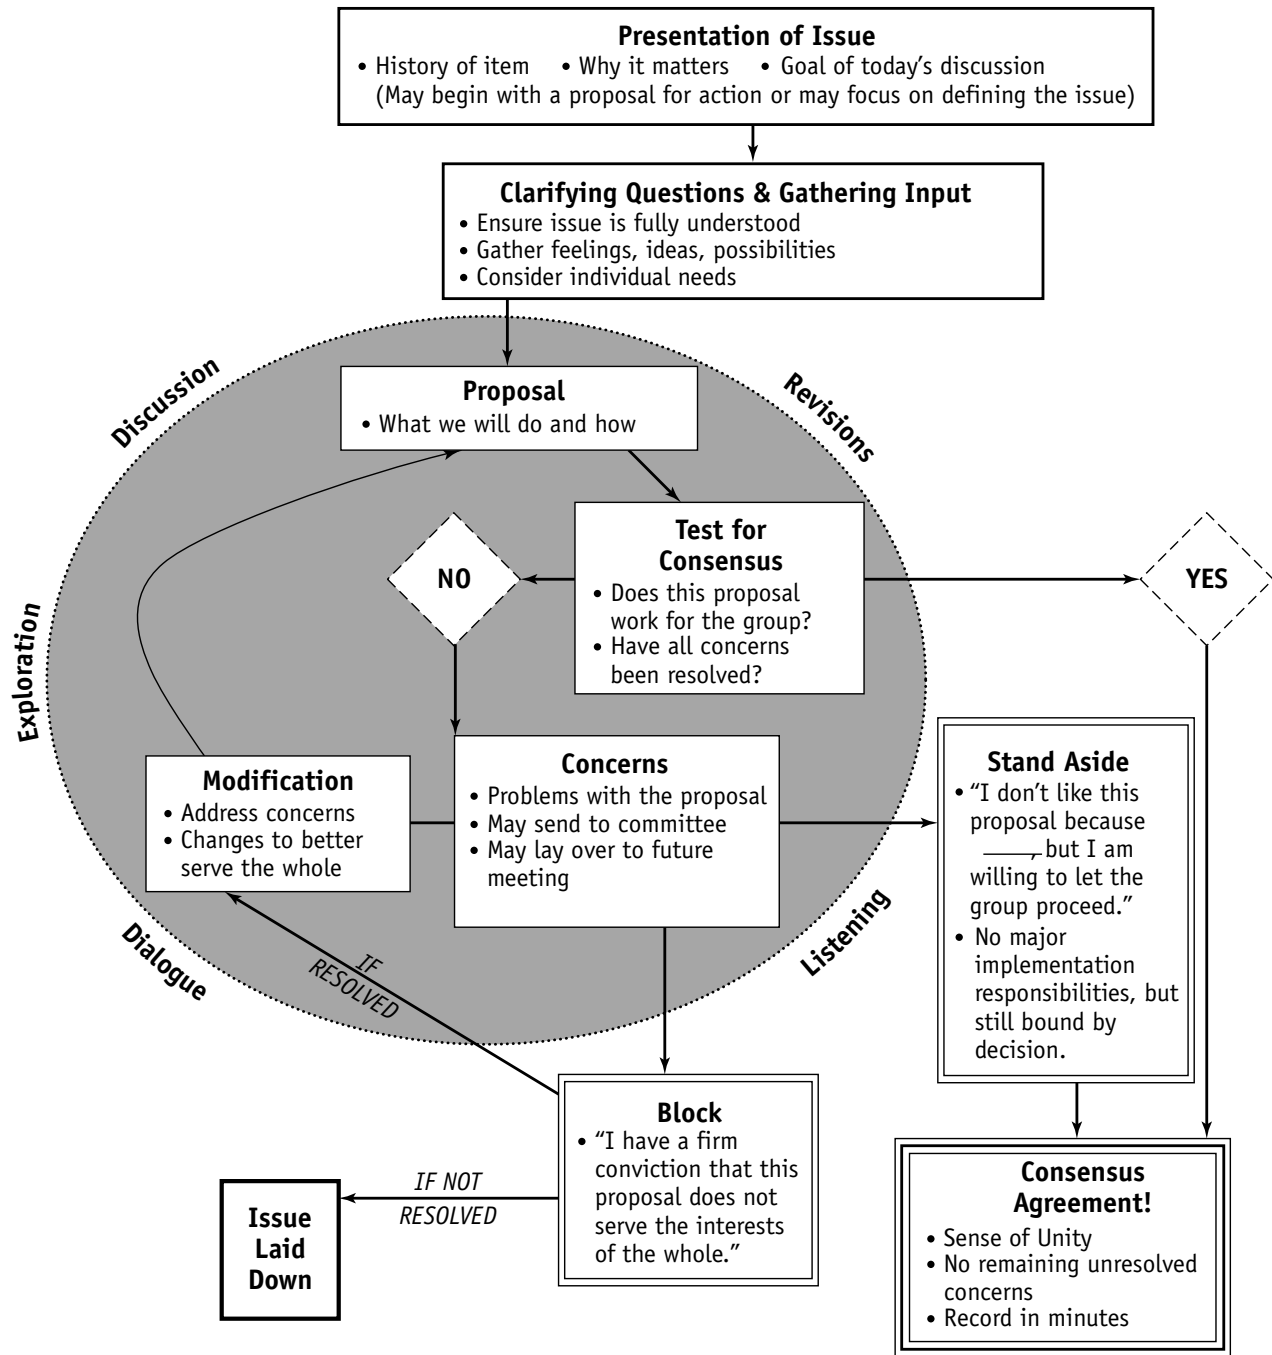


Figure A.4: Flowchart by Tree Bresson, from “The Change Handbook” [Peggy, 2009], emphasizes the details of background that informs proposals, as well as the role of standing aside – in particular, the value that those who stand aside from an issue are not tasked with implementation responsibilities.

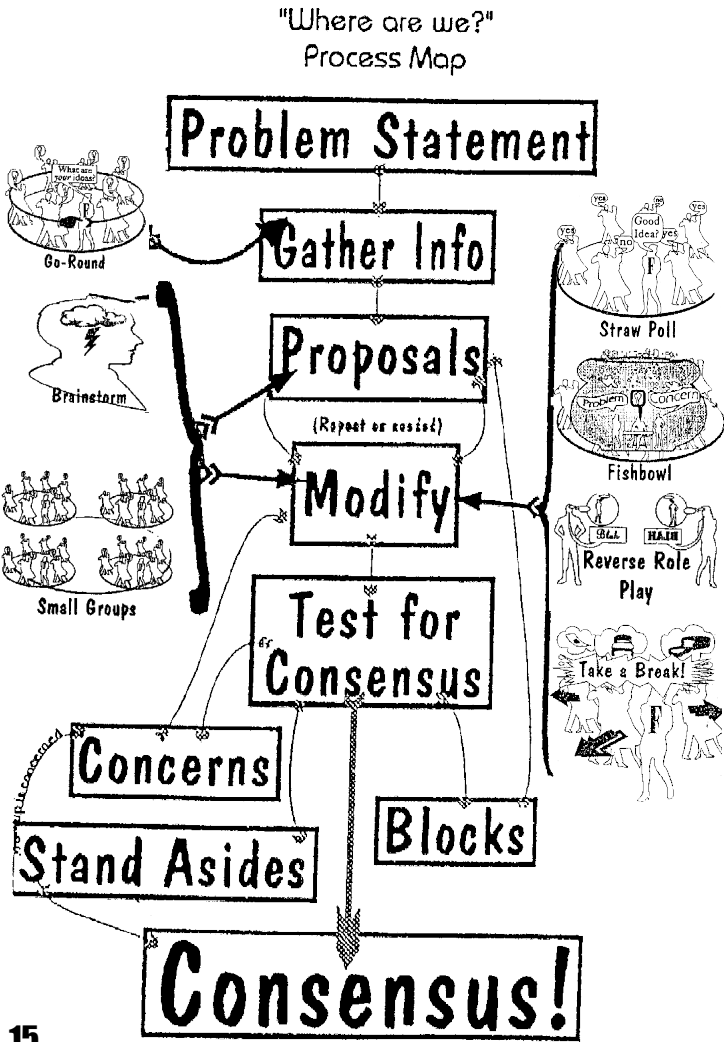


Figure A.5: Flowchart from the Direct Action Conference's "Shared Path" booklet, which emphasizes the different techniques used in the "Modify" stage of a proposal, including brainstorms, small groups, straw polls, fishbowls, reverse role-plays, and taking breaks. International [1995]

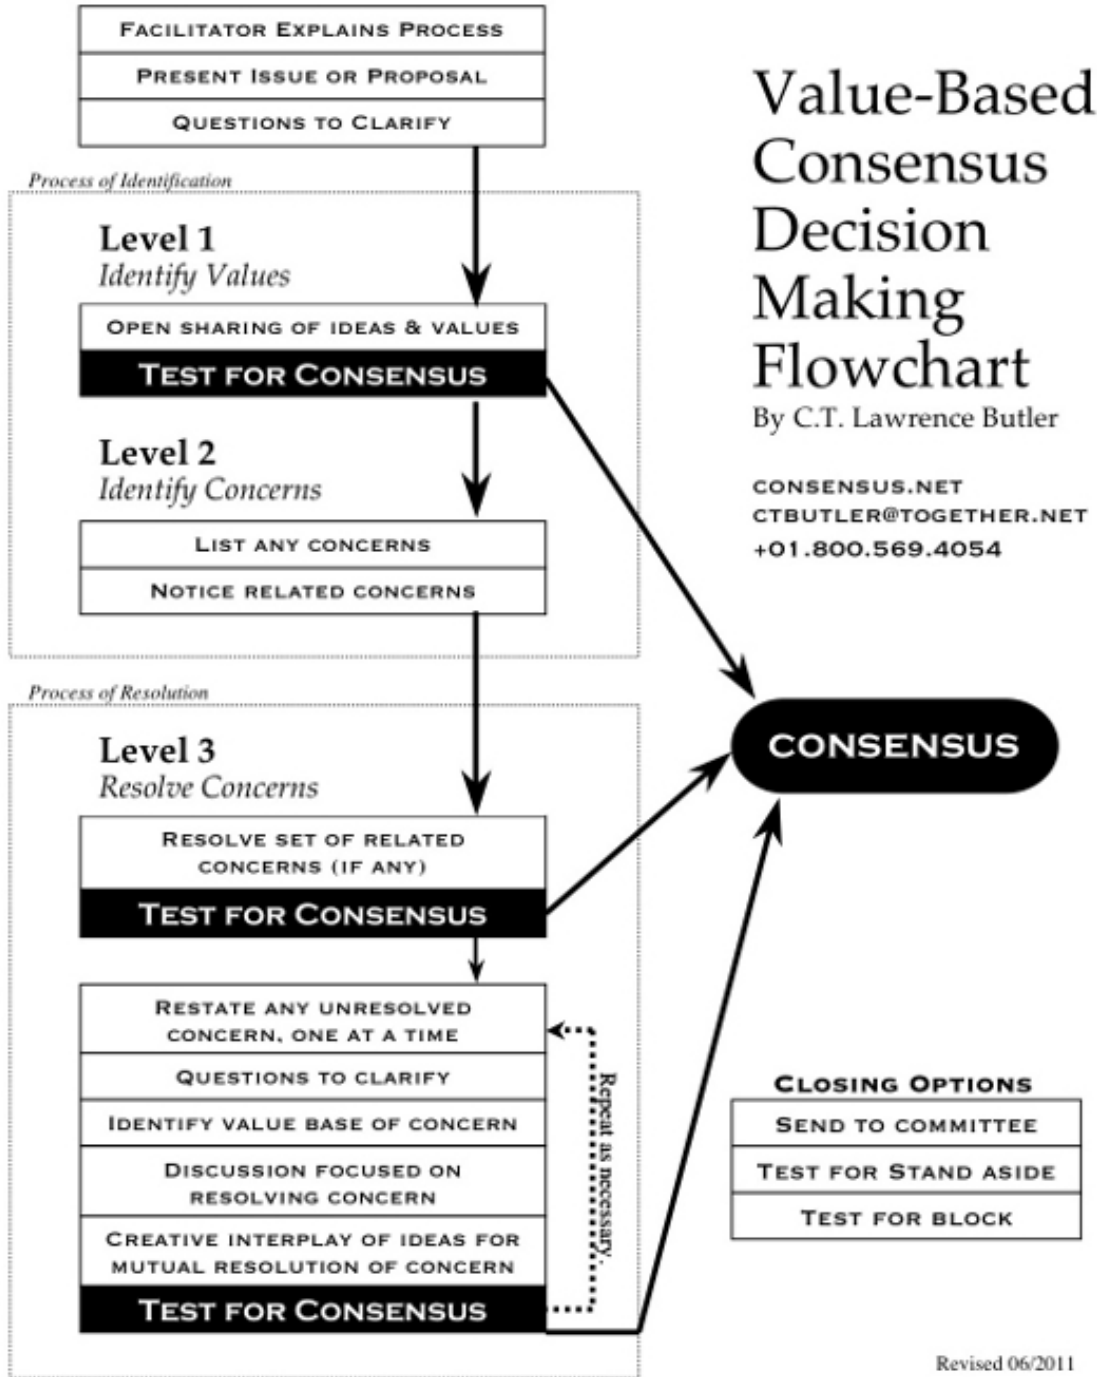


Figure A.6: C.T. Lawrence Butler's "Value-Based Consensus Decision Making Flowchart". <http://www.consensus.net>

Robert's Rules of Order, Abbreviated

Motion	Language	Interrupt speaker?	Second needed?	Debatable?	Ammend-able?	Vote needed?
PRIVILEGED MOTIONS deal with the welfare of the group, rather than with any specific proposal. They must be disposed of before the group can consider any other motion.						
Adjourn the meeting	I motion that we adjourn	no	yes	no	no	majority
Recess the meeting	I move we recess until	no	yes	no	yes	majority
Complain about noise, room temperature, etc.	Point of privilege	yes	no	no	no	none, chair rules
SUBSIDIARY MOTIONS provide various ways of modifying or disposing of main motions. They must be acted upon before all other motions except privileged motions.						
Suspend debate on a matter without calling for a vote	I move we table the matter	no	yes	no	no	majority
End debate	I move the previous question	no	yes	no	no	2/3 majority
Limit length of debate	I move debate on this matter be limited to...	no	yes	no	yes	2/3 majority
Ask for a vote by actual count, to verify a voice vote	I call for a division of the house	no	no	no	no	none, chair rules (majority if someone objects)
Postpone consideration of a matter to a specific time	I move we postpone the matter until...	no	yes	yes	yes	majority
Have a matter studied further	I move we refer this matter to a committee	no	yes	yes	yes	majority
Consider a matter informally	I move the question be considered informally	no	yes	yes	no	majority
Amend a motion	I move that this motion be amended by...	no	yes	yes	yes	majority
Reject a main motion without voting on the motion itself	I move the question be postponed indefinitely	no	yes	yes	no	majority
INCIDENTAL MOTIONS grow out of other business that the group is considering. They must be decided before the group can return to the question that brought them up.						
Correct an error in parliamentary procedure	Point of order	yes	no	no	no	none, chair rules
Object to a ruling by the chair	I appeal the chair's decision	yes	yes	yes	no	majority
Consider a matter that violates normal procedure, but does not violate the constitution or bylaws	I move we suspend the rules	no	yes	no	no	2/3 majority
Object to considering some matter	I object to the consideration of this matter	yes	no	no	no	2/3 majority
Obtain advice on proper procedure	I raise a parliamentary inquiry	yes	no	no	no	none, chair rules
Request information	Point of information	yes	no	no	no	none
Withdraw a motion	I request leave to withdraw the motion	no	no	no	no	majority
MAIN MOTIONS are the tools used to introduce new business.						
Introduce business	I move that...	no	yes	yes	yes	majority
Take on a matter previously tabled	I move we take from the table...	no	yes	no	no	majority
Reconsider a matter	I move we reconsider our already disposed of action relative to...	yes	yes	yes	no	majority
Strike out a motion previously passed	I move we reconsider our action relative to...	no	yes	yes	yes	majority
Consider a matter out of its scheduled order	I move we suspend the rules and consider...	no	yes	no	no	2/3 majority

Table A.1: Robert's Rules, abbreviated

B Meeting practices used by consensus-oriented groups

The following is a brief catalog of consensus practices used by groups practicing consensus which I have observed over the last 15 years of membership in such groups. Many of these techniques can be found described in greater detail in facilitation manuals and group work guides.¹

Roles

Facilitation roles are a critical component of a successful meeting. Facilitators help to direct the process of a meeting, but should ideally not contribute to the content of the meeting. Some groups formalize this by asking facilitators to formally “step down” from their role as facilitator, handing it to another group member, if they find they need to participate in the content of something that has come up in the meeting, and only resume their role as facilitators when they are again ready to give up participation in the discussion.

It is important for groups to frequently rotate facilitation roles. Despite peoples’ best intentions to the contrary, the role of a facilitator has a potential to accumulate power. But more importantly, meeting attendees who are familiar with techniques of facilitation are often better participants – more competently able to reason about group dynamics.

While smaller groups may consolidate all of the following roles into one, larger groups might separate them into multiple people.

Facilitators

The basic duties of a facilitator include the following:

- Explain the procedures, hand signals, and roles used by the group to any newcomers.
- Familiarize themselves with the agenda before the meeting, and direct progress through it during.

¹ <http://cultivate.coop> is a wiki devoted to cooperatives, including pages on consensus, and links to many further resources. Tree Bressen is a professional facilitator who has written several guides for consensus process and common pitfalls, including [Bressen, 2011a] and [Bressen, 2011b]. Educator Randy Schutt has also written several papers about consensus process, notably [Schutt, 2001] and [Schutt, 2007]. Longer-form instructional narratives about consensus technique include [Butler and Rothstein, 1991], [Gelderloos, 2006], and [Susskind et al., 1999]. “The Group Works Deck” is a pattern language style card-deck with a variety of techniques for group meetings, including consensus; found at <http://groupworksdeck.org>.

- Choose and implement the meeting and discussion formats to ask the group to use for different items.
- Keep discussion focused on the topic at hand.
- Summarize discussions, and keep track of minority viewpoints.
- Count votes or assess the group opinion at the resolution of proposals.

In addition to those roles, the facilitator may also take on some of the following.

Stack keepers

Stack keepers keep a stack (see below). They are responsible for ensuring that the full breadth of the group's diversity is represented in the discussion, and that no speaker is dominating the meeting. A stack keeper might also suggest different meeting formats (such as go-arounds, popcorn, etc.).

Time keepers

Time keepers watch the clock and remind the group of the remaining time in each agenda item. Groups face a fundamental tension between allowing a discussion to take the full space they need to elicit all of its issues, and ending the meeting in a reasonable time. Meetings that run too long can decrease accessibility, impinging on peoples' needs to work, take care of families, etc. But meetings that are too time-focused can lead to increased stress and poorer solutions. Groups have to strike a balance between these needs. A time keeper can help reflect that balance to a group, helping members to focus their comments. If time for an item runs out, the group can decide to extend time, to table the issue, or defer it to a committee.

Vibes watchers

A vibes watcher keeps track of the emotional state and energy of the group. If people's attention is beginning to wander (blank stares, fiddling with phones or bags, fidgeting), they can recommend that the group take a break, stand to stretch, or do some other focusing activity. If discussion becomes overly heated, they can recommend that the group take a moment to pause, breathe, and collect their thoughts before proceeding.

Note takers

Many groups suffer from short institutional memory, and issues can be rehashed again and again even after they have been thoroughly addressed. Keeping well-organized minutes of meetings can help avoid this, as well as offering transparency of the meeting's activity, and a record for people who missed the meeting of what went on. Different groups have different styles that they prefer for minutes – ranging from verbatim transcriptions of what people say to terse summaries of key points and decisions. Most important is that minutes be organized in an easy, available way for group members to find.

Since note taking is a fairly attention-intensive activity, note takers may need to hand off the duty if they wish to participate in a particular discussion. Shared documents can be particularly helpful for this.

Buddies

Buddies can be a useful way for groups to bring new members on board, and also to keep group members accountable to each other in between meetings. By pairing a members together, buddies can check in with each other and make sure that they understand everything they need, feel welcome and productive in the group, and follow through with any responsibilities. To overcome social resistance to buddies, it's helpful for the group to institutionalize the practice among all members.

Facilitation tools

Community agreements

Community agreements are a set of principles or guidelines that everyone in the group agrees to. They form a baseline that defines a group (similar to a bill of rights or a constitution). They are often rendered as short, pithy statements that can be easily remembered or referenced; short-hand for larger concepts or discussions. Some example items from the community agreements of groups I've worked with:

- *OK to Disagree*: It isn't necessary for everyone to agree all the time; disagreement is natural and healthy.
- *Try it on*: People should welcome chances to try new things, though that doesn't mean you need to stick with them.

- *1 diva, 1 mic*: Don't interrupt people who are speaking.
- *No one knows everything, together we know a lot*: A reminder to listen to what other people have to say.
- *We can't be articulate all the time*: Give others the benefit of the doubt, even if words don't flow smoothly.
- *Move up, Move back*: Consider whether you tend to dominate conversations, or speak less, and work on stretching out of your comfort zone to focusing more on listening or speaking. One group I participated in used the more cryptic *Move up, Move up* to emphasize that you would "move up into speaking", or "move up into listening", to convey the importance of careful listening.
- *Avoid blushacking*: "Blushacking" is a portmanteau of "blame", "shame", and "attack". Members are encouraged to take responsibility for their actions and feelings, and avoid transferring those to others.

In addition to agreements that provide norms for how the group should behave, groups that frequently work with new or changing membership might also develop *assumptions* which encapsulate or head-off unproductive discussions that distract from the group's main work. The following are assumptions from an anti-oppression training by AORTA²):

² <http://aortacollective.org>

- Systems of oppression exist
- It's not useful to argue about hierarchies of oppression
- All systems of oppression are interconnected
- Most of us have experienced being both targets and agents of oppression
- Dismantling systems of oppression will benefit everyone
- Placing blame helps no one, taking responsibility helps everyone.
- Confronting social injustice is painful and joyful

While it can sometimes be useful to have "101" style discussions that establish baselines and bring people up to speed on theory of intersectional oppression, this can also distract from other work. Explicit assumptions can help to avoid getting bogged down.

Stacks

A stack is a list of speakers who wish to speak. Most groups ask participants to raise their hand; the stack keeper (a role sometimes taken on by the facilitator, or sometimes by another meeting participant) will then write their name down. People might be called on in the order they raised their hands, or in a different order. Some common alternate orderings are:

- Progressive stack: Call on people who are less well represented in the group – whether by identity category (e.g. gender, race, old members, new members), or based on relative speaking time (calling on people who have spoken least).
- Preferencing speaking types: If the group uses hand signals to indicate the type of comment one wishes to make, the stack keeper might call on people with specific types first. A typical order of preference is: points of process, points of information, direct responses, clarifying questions, general comment.

The main advantage of a stack is that it prevents people from interrupting each other. Interruptions derail trains of thought and allow particular people to dominate conversations. They also can replicate male privilege, as men tend to interrupt women more than the reverse.

Proposals

Proposals are the main way that groups approve or reject courses of action and changes to policy. A proposal is a succinct statement of what the proposer would like the group to do. Ideally, proposals are shared with the group before the meeting, so that members can have a chance to think about or research the issue before making a decision about it. A well-formed proposal might include each of the following:

- A succinct description or name for the issue
- Pertinent background of the issue, or why the proposal is needed.
- A comparison of the status quo with the proposed changes.
- A list of alternate options.
- An estimate of the cost in money (if any) or time to implement the proposal.
- Any additional considerations that might influence peoples' opinions.
- Suggestions for how the proposal will be implemented.

Friendly amendments

Friendly amendments are a request by a group member that is mostly satisfied with a proposal to make a change that they believe would be welcomed by the person bringing the proposal. They are a way to more efficiently handle concerns without a contestational amendment process. After the proposal is made to the group, the amender asks the proposer if a particular amendment would be “friendly”, which they can then immediately incorporate into the proposal.

Straw polls

Straw polls (also called “temperature checks”) are non-binding votes which gauge the current sense of the group. They can be used to determine whether an issue is contentious or close to consensus. Groups might take a straw poll through an up-down vote (e.g. “raise your hand if you agree”), or using hand signals of approval or disapproval (see below). Straw polls can be attractive because they help to avoid lengthy discussions when the group doesn’t need to have them. However, taken prematurely, they can also lead groups to avoid discussing minority opinions that might change minds.

Tabling

When a group runs out of time to consider a particular issue, it can be “tabled”, or deferred for discussion at a later date. Every group I’ve ever been a part of has had more things to do than time to do them in. Tabling can help groups to prioritize the most important things to discuss at the current meeting.

Parking Lot / Garden / Bike Rack

A “parking lot”, “garden”, or “bike rack” are all names for a strategy for capturing side issues that need more discussion in the future (but which aren’t yet well developed enough to count as proposals or agenda items). In the course of discussion, groups might identify important issues for the group to consider, but which are not critical to the core of the issue being discussed. Rather than derailing discussion to address these issues or forgetting about them, a group can write them out on a piece of poster paper in the meeting space. Listing the issues visibly helps to capture their importance, but also to keep the group’s discussion focused on the topic at hand. Over a long meeting, a group might build time into the agenda to address these items; or they could be deferred to future meetings.

Deferring to committee

If a group has topical committees or working groups (e.g. finance, maintenance, outreach, etc.), rather than using the full group's limited time to dig into details and minutia of issues pertinent to those topics, the group might delegate the work to the appropriate committee. This is key to the basic exercise of *subsidiarity* – there is no need for a large group to micromanage what a smaller group is more nimbly and competently able to handle.

Comment types

Groups often use a taxonomy of different types of comments (speaking turns) in order to help manage speaking order. Typically, general comments of opinion or statements of concerns or support of a proposal are given the lowest precedence; while procedural or informational comments and questions are given higher precedence. The following are some common types.

Clarifying Question

A *clarifying question* (commonly signaled using a “C” shape with the hand, or a question-mark hook with the index finger) is a request that the person bringing a proposal, or someone who recently spoke, clarify some point of confusion in their comment or proposal. The question should not be editorial.

Points of process

A *point of process* (commonly signaled using a “triangle” shape with two hands) is a type of comment which is used to suggest a change in facilitation or to address concerns about the meeting venue and circumstance. One might raise a point of process to suggest a facilitation strategy for a particular issue, to call for a time check, to point out problems with the meeting space (noise, temperature, disruptions, etc), or to indicate when the group is violating its own rules or policies. In terms of Robert's Rules of Order, points of process cover *privileged motions* and *incidental motions*.

Direct responses

A *direct response* (commonly signaled by swinging pointed fingers back and forth, sawing the air) is a type of comment in which one participant requests to “jump the stack” (speak immediately rather than waiting) in order to respond directly to something another

participant has said. While this tool can sometimes allow more efficient conversation when someone has pertinent information that responds to something someone has just said, it can also lead to unproductive back-and-forth arguments. Many groups choose not to use this for that reason.

Point of Information

A *point of information* (commonly signaled by pointing an index finger vertically in the air) is a type of comment in which one has a pertinent, short, factual piece of information to share. Groups might allow commenters making points of information to use points of information to “jump the stack” when the facts presented might change the course of discussion.

Wrap-it-up

When discussion is running too long (or someone is rambling on), swinging a finger in a circular “lasso” shape can indicate “wrap-it-up”.

Approval / disapproval

“Twinkles” (after which this dissertation is named) are a widely used signal for approval, loosely based on American Sign Language for “applause”. Group members hold their hands up in the air and waggle their fingers back and forth. Some groups use “down twinkles”, or fingers wagging while pointed downward, as a signal for disapproval (though some groups find that distasteful). Another variant for approval is a knocking motion with a closed fist, which is American Sign Language for “yes”, or snapping fingers (audible, but less obtrusive than applause).

Hand signals are also used for voting on proposals. The simplest form is raising hands for approval or keeping them down for disapproval, but other forms include the “fist of five” (holding up a number of fingers expressing a degree of support), thumbs (up for yes, down for no, sideways for “unsure” or “abstain”), or more complicated variants such as the five-set of hand positions: thumbs up for approve, palm up for approve with concerns, flat hand with sideways palm for unsure, palm down for disapprove, and thumbs down for block.

Other symbols used for “blocking” by some groups include a fist raised in the air, or arms crossed in front of the chest.

Meeting phases

One of the important jobs of a facilitator is to establish a rhythm in the meeting which corresponds to a group's energy and available time. Breaking the meeting up into distinct phases can help with this. Some phases might include:

- **Orientation:** If there are new members in a group, it's important to orient them to how the group works, including any procedural issues, hand signals, or group norms around communication.
- **Agenda:** Read over and describe the agenda, including estimates of time for each item. This can be a time when group members can request that new things be added, though it is always preferable to have the agenda constructed in advance of the meeting so that group members know what to expect. Some groups will formally "approve" an agenda before beginning the meeting.
- **Check-in:** This is an opportunity for an ice breaker or go-around to get people started in the meeting, as well as a chance for members to reflect on any stressors or emotional concerns that might impact their participation in the meeting.
- **Announcements:** It is helpful to use a separate space for announcements that do not require discussion, so that groups can satisfy the need for keeping each other informed without resorting to lengthy agenda items. Each announcement might be time-limited (e.g. no more than 2 minutes per person).
- **Breaks:** Regular breaks are critical to maintaining group energy in long meetings.
- **Check-out:** A final go-around in which meeting participants have an opportunity to reflect on the meeting and their experience can be valuable for giving group members a sense of closure, as well as helping the group to learn from its process.

Formats

While the bread-and-butter of discussion formats for most consensus-oriented groups is the "open stack" in which a stack keeper calls on people in the order in which they raise their hand, a variety of other meeting formats can help break the monotony, ensure that every voice is heard, and more nimbly or appropriately handle different types of tasks. Here is a short list of some widely used formats.

Icebreakers / fire starters

An icebreaker or fire starter is a go-around in which members of the group introduce themselves. Many groups also ask a silly, casual, or meaningful question for members to answer. Groups that prioritize trans liberation also ask members to announce their *preferred gender pronouns*, the pronouns they wish others to use to refer to them by (e.g. “he/him/his”, “she/her/hers”, “they/them/theirs”, “ze/zir”, etc.), to normalize the practice of using pronouns that people prefer, rather than the ones that others might assume. Some reasons to begin meetings with icebreakers include:

- To introduce each member to the group, so that everyone at least has some familiarity with who the other group members are.
- To give people a chance to let the group know about any needs they have (e.g. “please speak loudly, I’m hard of hearing”, or “I can’t sit for long periods, so if I stand up, that’s why”, or “I will need to step out to breast feed my child”).
- To give each member moment to clear their throat and speak, increasing the likelihood that they will do so later in the meeting.
- To introduce levity and greater group familiarity.

Go-around

A go-around (e.g. “passing the conch”) is like an ice breaker, but used for moments when it is important for every member of a group to speak. This can be effective when the group is addressing contentious issues and wants to be sure that each viewpoint is heard. It can also be a useful tool to draw out members of the group who are reluctant to speak.

Popcorn

Popcorn is a discussion format where members of the group shout out short ideas when they think of them, without waiting to be called on. It is useful for brainstorming or collecting ideas (where the activity is *generative*); it is less useful when one participant might want to argue against what another participant has said.

Break-out groups

Large groups can be cumbersome – the larger the group gets, the less time each member has to discuss an issue, and the more time proportionally goes to facilitation. Large groups can break into

smaller groups to discuss an issue. After a predefined period of time, a spokesperson selected within each small group then reports the topics of the group's discussion back to the larger group.

World cafe

A world cafe is a refinement of break-out groups for generating ideas about complex, multi-faceted topics. The facilitator sets up several stations around the room, each of which has a topic. Each station is outfitted with poster paper, markers, and (if the group is large enough) a station facilitator. The whole group then splits into smaller groups, each of which visits one of the stations for a set period of time, and then rotates to the others. At each rotation, members of the small groups can see the ideas written by previous groups, and add their own.

Dotstorm

A dotstorm (also "dotmocracy", or less evocatively, "Nominal Group Technique") is a meeting format for generating ideas and then narrowing the choice of ideas to a few. (It is also implemented as an InterTwinkles tool, discussed on page 107). It proceeds in five stages:

1. The facilitator announces a problem, idea, or challenge. It is helpful to phrase it as "In what ways might we...?" The facilitator distributes sticky notes or paper.
2. Members of the group spend time drawing or writing their ideas on the note paper, silently; one idea per page.
3. Members of the group (perhaps using a popcorn style) explain the ideas they have generated, and place them on a common board or table.
4. The group discusses the ideas, and adds any new ideas they might think of. The facilitator guides the categorization of ideas into any similar groupings. To help elicit more ideas, they might suggest organizing the ideas spatially across two conceptual axes (e.g. "long term vs short term plans", or "cheap vs expensive plans"), to see if there's a corner that has fewer ideas.
5. Members vote on the ideas they like the most by distributing 4 or 5 "dots" (stickers or small pieces of tape) on the ideas or groups of ideas they like the best.

Spectrogram

A spectrogram is a type of icebreaker which can help to normalize disagreement and difference of opinion. The facilitator asks people in the group to come up with the most controversial and incendiary ideas they can, which they expect would divide opinion in the group. If the group is large (more than 20 or 30 people), they might develop these ideas in small groups. The facilitator then chooses people one at a time to announce their controversial idea, and asks people to arrange themselves on a line in the room ranging from “strongly agree” to “strongly disagree”. The facilitator then asks people at the extremes along the spectrogram to explain their reasoning. Spectrograms can be useful in preempting passive aggression and false consensus by normalizing public airing of differences of opinion.

Neighbor interviews

Neighbor interviews are a technique for people to gain deeper understanding of what particular people think about an issue. Members of the group are asked to pair off, and interview each other about the topic under consideration. Each person is then asked to explain their partner’s position to the whole group.

Solipsist’s meeting

A solipsist’s meeting is a technique in which each member of the group is asked to write their opinion about a particularly controversial issue on paper, and to hand this paper to the facilitator. The facilitator then anonymously reads the ideas back to the group. This technique can help when dealing with issues that are very difficult to talk about, especially where there are strong social divisions in the group. However, it can also lead to passive aggression if members take it as an opportunity to snipe others from the safety of anonymity. Topics should be chosen judiciously, with extra group work before and after the meeting to address these risks.

Bibliography

Censorship of wikipedia - wikipedia, the free encyclopedia. URL https://en.wikipedia.org/wiki/Censorship_of_Wikipedia.

The diaspora project. URL <http://diasporaproject.org/>.

Free software is a matter of liberty, not price - free software foundation - working together for free software. URL <https://www.fsf.org/about/>.

Wikipedia:Wikipedians, August 2013. URL <https://en.wikipedia.org/w/index.php?title=Wikipedia:Wikipedians&oldid=566411546>. Page Version ID: 566411546.

Gabriel Abend. The meaning of 'Theory'. *Sociological Theory*, 26(2):173–199, 2008. ISSN 1467-9558. DOI: 10.1111/j.1467-9558.2008.00324.x. URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1467-9558.2008.00324.x/abstract>.

Mumia Abu-Jamal. *We Want Freedom: A Life in the Black Panther Party*. South End Press, January 2004. ISBN 9780896087187.

John Adams. *A Defence of the Constitutions of Government of the United States of America: Against the Attack of M. Turgot in His Letter to Dr. Price, Dated the Twenty-second Day of March, 1778*. J. Stockdale, 1794.

Yochai Benkler. *The wealth of networks: how social production transforms markets and freedom*. Yale University Press, New Haven, 2006. ISBN 9780300125771 0300125771 9780300110562 0300110561.

Kenneth Bettenhausen and J. Keith Murnighan. The emergence of norms in competitive decision-making groups. *Administrative Science Quarterly*, 30(3):350, September 1985. ISSN 00018392. DOI: 10.2307/2392667. URL <http://www.jstor.org/discover/10.2307/2392667?uid=3739696&uid=2&uid=4&uid=3739256&sid=21102549878047>.

Scott Bradner. IETF working group guidelines and procedures, September 1998. URL <https://tools.ietf.org/html/rfc2418>.

- Tree Bressen. Consensus decision-making, November 2011a. URL <http://treegroup.info/topics/handout-consensus.pdf>.
- Tree Bressen. The top 10 most common mistakes in consensus process and how to avoid them, October 2011b. URL <http://treegroup.info/topics/Top-10-Consensus-Mistakes.pdf>.
- Joe Brockmeier. Apple and KHTML [LWN.net], May 2005. URL <https://lwn.net/Articles/136538/>.
- Jon Brodtkin. Linus torvalds defends his right to shame linux kernel developers, July 2013. URL <http://arstechnica.com/information-technology/2013/07/linus-torvalds-defends-his-right-to-shame-linux-kernel-developers/>.
- C. T. L. Butler and A. Rothstein. *On conflict and consensus: A handbook on formal consensus decisionmaking*. Food Not Bombs, 1991. URL http://theanarchistlibrary.org/library/C.T._Butler_and_Amy_Rothstein__On_Conflict_and_Consensus__a_handbook_on_Forma_Conensus_decisionmaking.pdf.
- C.T. Butler. A guide to formal consensus. URL <http://consensus.net/revolutionary.html>.
- R. P. Carasik and C. E. Grantham. A case study of CSCW in a dispersed organization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '88*, pages 61–66, New York, NY, USA, 1988. ACM. ISBN 0-201-14237-6. DOI: 10.1145/57167.57177. URL <http://doi.acm.org/10.1145/57167.57177>.
- Carlos Flores, Juan Ludlow, Chauncey Bell, Raul Mora, Terry Winograd, and Michael Graves. Method and apparatus for structuring and managing human communications by explicitly defining the types of communications permitted between participants, June 1993.
- clenchner. The 7 sins of occupy, September 2012. URL <http://www.dailykos.com/story/2012/09/20/1134321/-The-7-Sins-of-Occupy>.
- Michael D. Cohen, James G. March, and Johan P. Olsen. A garbage can model of organizational choice. *Administrative Science Quarterly*, 17(1):1, March 1972. ISSN 00018392. DOI: 10.2307/2392088. URL <http://www.jstor.org/discover/10.2307/2392088?uid=3739696&uid=2&uid=4&uid=3739256&sid=21102485258521>.
- Andrew Cornell. *Oppose and Propose: Lessons from Movement for a New Society*. AK Press, April 2011. ISBN 9781849350662.

- Rubén Cuevas, Michal Kryczka, Ángel Cuevas, Sebastian Kaune, Carmen Guerrero, and Reza Rejaie. Is content publishing in BitTorrent altruistic or profit-driven?, December 2010. URL <http://e-archivo.uc3m.es/handle/10016/10116>. 13 pages, 5 figures.– ArXiv pre-print version available at: <http://arxiv.org/abs/1007.2327>.
- Todd R. Davies and Seeta Peña Gangadharan. *Online Deliberation: Design, Research, and Practice*. Center for the Study of Language and Information, Leland Stanford Junior University, 2009. ISBN 9781575865546.
- André L. Delbecq and Andrew H. Van de Ven. A group process model for problem identification and program planning. *The Journal of Applied Behavioral Science*, 7(4):466–492, July 1971. ISSN 0021-8863, 1552-6879. DOI: 10.1177/002188637100700404. URL <http://jab.sagepub.com/content/7/4/466>.
- G. Desanctis and R.B. Gallupe. A foundation for the study of group decision support systems. *Management science*, pages 589–609, 1987.
- Gerardine DeSanctis and Marshall Scott Poole. Capturing the complexity in advanced technology use: Adaptive structuration theory. *Organization Science*, 5(2):121–147, May 1994. ISSN 1047-7039. DOI: 10.2307/2635011. URL <http://www.jstor.org/stable/2635011>. ArticleType: research-article / Full publication date: May, 1994 / Copyright © 1994 INFORMS.
- Gerardine DeSanctis, James R. Snyder, and Marshall Scott Poole. The meaning of the interface: A functional and holistic evaluation of a meeting software system. *Decision Support Systems*, 11(4):319–335, May 1994. ISSN 0167-9236. DOI: 10.1016/0167-9236(94)90079-5. URL <http://www.sciencedirect.com/science/article/pii/0167923694900795>.
- Gerardine DeSanctis, Marshall Poole, Ilze Zigurs, and other Associates. The minnesota GDSS research project: Group support systems, group processes, and outcomes. *Journal of the Association for Information Systems*, 9(10), October 2008. ISSN 1536-9323. URL <http://aisel.aisnet.org/jais/vol9/iss10/6>.
- The Economist. Leaderless, consensus-based participatory democracy and its discontents, October 2011. URL <http://www.economist.com/blogs/democracyinamerica/2011/10/occupy-wall-street-3>.
- Douglas Engelbart. Augmenting human intellect: A conceptual framework. Technical report, October 1962. URL http://sloan.stanford.edu/mousesite/EngelbartPapers/B5_F18_ConceptFrameworkInd.html.

Thomas Ferguson. Game theory. URL http://www.math.ucla.edu/~tom/Game_Theory/Contents.html.

Jerry Fjermestad and Starr Roxanne Hiltz. An assessment of group support systems experiment research: Methodology and results. *Journal of Management Information Systems*, 15(3):7–149, 1998. ISSN 07421222. URL <http://search.proquest.com/docview/218925443/abstract?accountid=12492>.

Fernando Flores, Michael Graves, Brad Hartfield, and Terry Winograd. Computer systems and the design of organizational interaction. *ACM Trans. Inf. Syst.*, 6(2):153–172, April 1988. ISSN 1046-8188. DOI: 10.1145/45941.45943. URL <http://doi.acm.org/10.1145/45941.45943>.

Seeds for Change. Consensus decision making. URL <http://seedsforchange.org.uk/consensus.pdf>.

Marcus Foth, Laura Forlano, and Christine Satchell. *From Social Butterfly to Engaged Citizen: Urban Informatics, Social Media, Ubiquitous Computing, and Mobile Technology to Support Citizen Engagement*. MIT Press, 2011. ISBN 9780262016513.

Jo Freeman. The tyranny of stuctureless, 1972. URL <http://www.jofreeman.com/joreen/tyranny.htm>.

Ina Fried. Behind bing’s blue links, March 2010. URL http://news.cnet.com/8301-13860_3-20000623-56.html.

Jason Fried. Why i run a flat company, April 2011. URL <http://www.inc.com/magazine/20110401/jason-fried-why-i-run-a-flat-company.html>.

Christopher Fry. Justify: A web-based tool for consensus decision making. February 2012. URL <http://web.media.mit.edu/~cfry/justify/Justify-IUI-Submitted.doc>.

Alexander R. Galloway. *Protocol: How Control Exists After Decentralization*. MIT Press, 2004. ISBN 9780262072472.

Sarah N. Gatson and Amanda Zweerink. Ethnography online: ‘Natives’ practising and inscribing community. *Qualitative Research*, 4(2):179–200, August 2004. ISSN 1468-7941, 1741-3109. DOI: 10.1177/1468794104044431. URL <http://qrj.sagepub.com/content/4/2/179>.

Peter Gelderloos. *Consensus: A New Handbook for Grassroots Social, Political, and Environmental Groups*. See Sharp Press, September 2006. ISBN 9781884365393.

- Anthony Giddens. *The Constitution of Society: Outline of the Theory of Structuration*. University of California Press, January 1986. ISBN 9780520057289.
- Anthony Giddens. *New Rules of Sociological Method: A Positive Critique of Interpretative Sociologies*. Stanford University Press, 1993. ISBN 9780804722261.
- Doris A Graber. *Verbal behavior and politics*. University of Illinois Press, Urbana, 1976. ISBN 0252002628 9780252002625.
- David Graeber. *The democracy project: a history, a crisis, a movement*. Allen Lane ; Penguin, London, England; New York, 2012. ISBN 9781846146633 1846146631.
- David Graeber. Some remarks on consensus, February 2013. URL <http://occupywallst.org/article/david-graeber-some-remarks-consensus/>.
- Kenneth L. Hacker and Jan van Dijk. *Digital Democracy: Issues of Theory and Practice*. SAGE, December 2000. ISBN 1446264823.
- Eran Hammer. OAuth 2.0 and the road to hell, July 2012. URL <http://hueniverse.com/2012/07/oauth-2-0-and-the-road-to-hell/>.
- Benjamin Mako Hill. Taking a principled position on software freedom | copyrighteous, July 2009. URL <http://mako.cc/copyrighteous/taking-a-principled-position-on-software-freedom>.
- Benjamin Mako Hill. Almost wikipedia: What eight collaborative encyclopedia projects reveal about mechanisms of collective action | berkman center, October 2011. URL <https://cyber.law.harvard.edu/events/luncheon/2011/10/makohill>.
- Benjamin Mako Hill. The institute for cultural diplomacy and wikipedia, March 2013. URL <http://mako.cc/copyrighteous/the-institute-for-cultural-diplomacy-and-wikipedia>.
- Starr Roxanne Hiltz. *The Network Nation: Human Communication Via Computer*. MIT Press, 1993. ISBN 9780262581202.
- DAC International. Shared path, 1995. URL http://www.zhaba.cz/uploads/media/Shared_Path.pdf.
- Sam Kaner. *Facilitator's Guide to Participatory Decision-Making*. John Wiley and Sons, March 2011. ISBN 9781118047019.
- Kate Khatib, Margaret Killjoy, and Mike McGuire. *We Are Many: Reflections on Movement Strategy From Occupation to Liberation*. AK Press, 2012. ISBN 9781849351164.

- Taemie Kim, Agnes Chang, Lindsey Holland, and Alex Sandy Pentland. Meeting mediator: enhancing group collaboration using sociometric feedback. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work, CSCW '08*, pages 457–466, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-007-4. DOI: 10.1145/1460563.1460636. URL <http://doi.acm.org/10.1145/1460563.1460636>.
- K.L. Kraemer and J.L. King. Computer-based systems for cooperative work and group decision making. *ACM Computing Surveys (CSUR)*, 20(2):115–146, 1988.
- Travis Kriplean, Jonathan Morgan, Deen Freelon, Alan Borning, and Lance Bennett. Supporting reflective public thought with considerit. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work, CSCW '12*, pages 265–274, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1086-4. DOI: 10.1145/2145204.2145249. URL <http://doi.acm.org/10.1145/2145204.2145249>.
- J. Clayton Lafferty, Patrick M. Eady, Alonzo William Pond, and Human Synergetics. *The Desert Survival Problem: A Group Decision Making Experience for Examining and Increasing Individual and Team Effectiveness: Manual*. Experimental Learning Methods, 1974.
- J. C R Licklider. Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics*, HFE-1(1):4–11, 1960. ISSN 0099-4561. DOI: 10.1109/THFE2.1960.4503259.
- Joseph CR Licklider and Robert W Taylor. The computer as a communication device. *Science and technology*, 76(2):1–3, 1968.
- John Van Maanen. *Tales of the Field: On Writing Ethnography, Second Edition*. University of Chicago Press, July 2011. ISBN 9780226849645.
- David Marcus. The horizontalists. *Dissent*, 59(4):54–59, 2012. ISSN 1946-0910. DOI: 10.1353/dss.2012.0092. URL <http://muse.jhu.edu/journals/dissent/v059/59.4.marcus.html>. <p>Volume 59, Number 4, Fall 2012</p>.
- Marianne Schmid Mast. Dominance as expressed and inferred through speaking time. *Human Communication Research*, 28(3):420–450, 2002. ISSN 1468-2958. DOI: 10.1111/j.1468-2958.2002.tb00814.x. URL <http://onlinelibrary.wiley.com/doi/10.1111/j.1468-2958.2002.tb00814.x/abstract>.
- Joseph Edward McGrath. *Groups: Interaction and Performance*. Prentice-Hall, 1984. ISBN 9780133657005.
- Evgeny Morozov. The meme hustler. *The Baffler*, 22, March 2013a.

- Evgeny Morozov. Why social movements should ignore social media. *The New Republic*, February 2013b. ISSN 0028-6583. URL <http://www.newrepublic.com/article/112189/>.
- Chantal Mouffe. Deliberative democracy or agonistic pluralism, 2000.
- Colin Moynihan. Tad hirsch, media lab grad. student, subpoenaed for users of TXTmob, system employed by protesters. *The New York Times*, April 2008. URL <http://tech.mit.edu/V128/N15/txtmob.html>.
- Milton L. Mueller and Hadi Asghari. Deep packet inspection and bandwidth management: Battles over BitTorrent in canada and the united states. *Telecommunications Policy*, 36(6):462–475, July 2012. ISSN 0308-5961. DOI: 10.1016/j.telpol.2012.04.003. URL <http://www.sciencedirect.com/science/article/pii/S0308596112000523>.
- Sidney John Roderick Noel. *From power sharing to democracy: post-conflict institutions in ethnically divided societies*. McGill-Queen's Press - MQUP, 2005. ISBN 9780773529489.
- Holman Peggy. *The Change Handbook: The Definitive Resource on Today's Best Methods for Engaging Whole Systems: Easyread Edition*. ReadHowYouWant.com, June 2009. ISBN 9781442978386.
- Francesca Polletta. *Freedom Is an Endless Meeting: Democracy in American Social Movements*. University Of Chicago Press, 1 edition, May 2004. ISBN 0226674495.
- Barbara Ransby. *Ella Baker And The Black Freedom Movement: A Radical Democratic Vision*. Univ of North Carolina Press, 2003. ISBN 9780807862704.
- Joseph Reagle. "Free as in sexist?" free culture and the gender gap. *First Monday*, 18(1), December 2012. ISSN 13960466. URL <http://firstmonday.org/ojs/index.php/fm/article/view/4291>.
- GUNNAR Ristroph. Debian's democracy. *Online Deliberation: Design, Research and Practice*. Chicago, Illinois, USA: Center for the Study of Language, pages 207–212, 2009. URL http://odbook.stanford.edu/static/filedocument/2009/11/15/Chapter_17._Ristroph.pdf.
- Henry M Robert and Sarah Corbin Robert. *Robert's rules of order newly revised*. Da Capo Press, Philadelphia, Pa., 2011. ISBN 9780306820205 030682020X.
- Howard Ryan. Blocking progress: Consensus decision making in the anti-nuclear movement, 1983. URL <http://www.connexions.org/CxLibrary/CX6187.htm>.

Elizabeth B.-N. Sanders and Pieter Jan Stappers. Co-creation and the new landscapes of design. *CoDesign*, 4(1):5–18, 2008. ISSN 1571-0882. DOI: 10.1080/15710880701875068. URL <http://www.tandfonline.com/doi/abs/10.1080/15710880701875068>.

Randy Schutt. Notes on consensus decision-making, August 2001. URL <http://www.vernalproject.org/papers/process/ConsensNotes.pdf>.

Randy Schutt. Consensus is not unanimity, December 2007. URL <http://www.vernalproject.org/papers/process/ConsensusNotUnanimity.pdf>.

Andrew Sellars. Seized sites: The in rem forfeiture of copyright-infringing domain names. SSRN Scholarly Paper ID 1835604, Social Science Research Network, Rochester, NY, May 2011. URL <http://papers.ssrn.com/abstract=1835604>.

Seuss. *Yertle the turtle, and other stories*,. Random House, New York, 1958. ISBN 0394800877 9780394800875.

William H. Sewell. A theory of structure: Duality, agency, and transformation. *The American Journal of Sociology*, 98(1):1–29, July 1992. ISSN 0002-9602. URL <http://www.jstor.org/stable/2781191>. ArticleType: research-article / Full publication date: Jul., 1992 / Copyright © 1992 The University of Chicago Press.

Uri Shwed and Peter S. Bearman. The temporal structure of scientific consensus formation. *American Sociological Review*, 75(6):817–840, December 2010. ISSN 0003-1224, 1939-8271. DOI: 10.1177/0003122410388488. URL <http://asr.sagepub.com/content/75/6/817>.

Marina Sitrin. *Horizontalism: voices of popular power in Argentina*. AK Press, Edinburgh, Scotland; Oakland, CA, 2006. ISBN 9781904859581 1904859585.

Jim Slaughter. Parliamentary procedure in the 21st century.

Jessamyn Smith. What she really said: Fighting sexist jokes the geeky way!, March 2012. URL <http://geekfeminism.org/2012/03/19/what-she-really-said-fighting-sexist-jokes-the-geeky-way/>.

Werner Sperschneider and Kirsten Bagger. Ethnographic fieldwork under industrial constraints: Toward design-in-context. *International Journal of Human-Computer Interaction*, 15(1):41–50, 2003. ISSN 1044-7318. URL http://www.tandfonline.com/doi/abs/10.1207/S15327590IJHC1501_04.

- Clay Spinuzzi. The methodology of participatory design. *Technical Communication*, 52(2):163–174, May 2005.
- Garold Stasser and Dennis Stewart. Discovery of hidden profiles by decision-making groups: Solving a problem versus making a judgment. *Journal of Personality and Social Psychology*, 63(3):426–434, 1992. ISSN 1939-1315(Electronic);0022-3514(Print). DOI: 10.1037/0022-3514.63.3.426.
- Mark Stiegler. Introduction to petname systems, June 2010. URL <http://www.skyhunter.com/marcs/petnames/IntroPetNames.html>.
- Lawrence Susskind, Sarah McKernan, and Jennifer Thomas-Larmer. *The consensus building handbook: a comprehensive guide to reaching agreement*. SAGE, August 1999. ISBN 9780761908449.
- Aaron Swartz. Squaring the triangle: Secure, decentralized, human-readable names (aaron swartz's raw thought), January 2011. URL <http://www.aaronsw.com/weblog/squarezooko>.
- Justine Alexandra Roberts Tunney. Occupiers! stop using consensus!, February 2013. URL <http://occupywallst.org/article/occupiers-stop-using-consensus/>.
- Jeff Tyson. How the old napster worked, October 2000. URL <http://computer.howstuffworks.com/napster.htm>.
- Philly Stands Up. A stand up start-up. URL <http://www.phillystandsup.com/PDFS/AStandUpStartUp.PDF>.
- Arthur B. Vangundy. *Techniques of Structured Problem Solving*. Springer, April 1988. ISBN 9780442288471.
- Jim Whitescarver. Electronic information exchange system II, February 2002. URL <http://web.archive.org/web/20020220171051/http://www.njit.edu/njIT/Department/CCCC/eies.html>.
- Ilze Zigurs, Bonnie K. Buckland, James R. Connolly, and E. Vance Wilson. A test of task-technology fit theory for group support systems. *SIGMIS Database*, 30(3-4):34–50, September 1999. ISSN 0095-0033. DOI: 10.1145/344241.344244. URL <http://doi.acm.org/10.1145/344241.344244>.