Group Decision-Making in Open Source Development: Putting Condorcet's Method in Practice

Christopher N. Lawrence
Department of History and Political Science
Macon State College/Middle Georgia State College
Macon, Georgia 31206
Email: c.n.lawrence@gmail.com

January 1, 2013

Abstract

Since 1999, the Debian Project, a worldwide organization of over 1000 computer programmers and other enthusiasts devoted to producing a free distribution of the Linux operating system, has been choosing its leaders and making other decisions about project governance using ranked-choice voting along with Condorcet vote counting and an effective "none of the above" option. In this paper, I describe the project's experience with using Condorcet methods, and suggest ways in which Debian's experience might be valuable to informing debates in public elections regarding the use of ranked-choice voting and "none of the above" options.

The effort to find a fair way to count votes is almost as old as the concept of group decision-making itself.[1] The expansion of the voting franchise and the rise of political parties in 19th century Europe and North America led to more impetus for electoral reform and broader interest in electoral systems beyond plurality voting, including the single transferable vote, limited voting, and cumulative voting (Droop, 2007 (1881), p. 7). Reformers, such as Droop, were particularly concerned with the need to represent minority groups and the overrepresentation of the majority within the legislature (8–10). While these issues can be addressed by using proportional representation systems, in which a

[1]For a lengthy overview of the history of voting systems, see Szpiro (2010).

number of seats are filled in proportion to their overall share of the vote, solving the problems associated with plurality elections in the event of a single-winner election is more problematic.

In fact, Arrow (1970) demonstrated that there is no single, fair solution to deciding between more than two courses of action or candidates; with three or more options, the electoral rules chosen will inevitably influence the outcome, and thus multiple possible outcomes can be considered to be "fair" using the same set of voters. For example, in the United States, where in most states a multitude of candidates qualify for placement on the presidential ballot, plurality elections are held in each state and the District of Columbia to select the members of the Electoral College;[2] however, it might be considered equally fair to allocate electors in proportion to the statewide vote, or to require a runoff election if no candidate receives an absolute majority of the vote, or perhaps many other variations and permutations of these procedures.

A three-option election can be used to demonstrate Arrow's finding. In the 1992 presidential election, Repubican incumbent George H.W. Bush, Democratic challenger Bill Clinton, and independent candidate H. Ross Perot faced each other for the presidency. Clinton was elected president using the plurality rule for winning presidential electors, even though he only received an absolute majority of votes in two jurisdictions (Arkansas and the District of Columbia) and 43% of the popular vote nationwide, raising the possibility that had the United States a rule like that of France requiring a run-off when the leading candidate does not receive an absolute majority, with Perot excluded in the second round,

[2]In the cases of Maine and Nebraska, a single election is held with votes tallied both to select per-district electors and the two statewide electors; in the other 49 jurisdictions, a single, statewide winner-takes-all contest for all of the state's electors takes place (U.S. National Archives and Records Administration, 2012).

Bush might have won reelection. It is also possible that more voters would have been happy, or at least content, had Perot been elected than either Bush or Clinton, on the presumption that most Republicans would rather have seen Perot in the White House than Clinton, and most Democrats would have rather seen Perot in the White House than Bush. Since neither of these counterfactual election scenarios came about, and the electoral system does not require voters to disclose a preference ranking of candidates, we can only conclude that the election *might* have turned out differently under an electoral system, but it is not difficult to conceive of either eventuality being true.

Despite Arrow's Theorem, scholars have continued to try to develop voting systems that will maximize the fairness of the electoral system while minimizing the prospects of strategic manipulation or perverse outcomes. In the following section, I discuss the Schulze method, an approach that has many of the desirable properties for a voting system, most notably that the winner of the election will be a *Condorcet winner*, if possible.

1 The Schulze Method

The Condorcet winner of an election is defined as the candidate (or alternative) that is preferred over all the other candidates on a pairwise basis. For example, in the 1992 presidential election example above, had more voters preferred Perot to Clinton than preferred Clinton to Perot, and more voters preferred Perot to Bush than preferred Bush to Perot, Perot would be considered the Condorcet winner. It is possible, however, that no candidate in an election will be a Condorcet winner; thus, any effective electoral system based on the Condorcet winner criterion must account for this possibility.

Determining a Condorcet winner requires voters to submit a ranked ballot, as the voter would in most other single-round, non-plurality systems (such as instant runoff voting and the Borda count[3]). Schulze (2011a) describes a specific algorithm for determining a winner, although there are a number of alternative approaches that will arrive at the same results. The Debian Constitition specifies one approach (step numbering and formatting adjusted):[4]

1. From the list of… options, we generate a list of pairwise defeats. An option A defeats an option B, if V(A,B) is strictly greater than V(B,A).
2. From the list of… pairwise defeats, we generate a set of transitive defeats. An option A transitively defeats an option C if A defeats C or if there is some other option B where A defeats B AND B transitively defeats C.
3. We construct the Schwartz set from the set of transitive defeats. An option A is in the Schwartz set if for all options B, either A transitively defeats B, or B does not transitively defeat A.
4. If there are defeats between options in the Schwartz set, we drop the weakest such defeats from the list of pairwise defeats, and return to step 5. A defeat (A,X) is weaker than a defeat (B,Y) if V(A,X) is less than V(B,Y). Also, (A,X) is weaker than (B,Y) if V(A,X) is equal to V(B,Y) and V(X,A) is greater than V(Y,B). A weakest defeat is a defeat that has no other defeat weaker than it. There may be more than one such defeat.
5. If there are no defeats within the Schwartz set, then the winner is chosen from the options in the Schwartz set. If there is only one such option, it is the winner. If there are multiple options, the elector with the casting vote[5] chooses which of those options wins. (The Debian Project, 2007)

In practice, the Schulze method can be—and has been—readily implemented in a number of computer programming languages, including Lua, Perl, and Python.

[3]The Borda count is most famously used in the United States to rank college sports teams in football and basketball.

[4]The description below has been slightly modified to omit the "local quota" provision included in Debian's adaptation of the method, discussed in greater detail below.

[5]In the case of Debian Project Leader elections, or other elections by general resolution, the incumbent DPL would have the casting vote. To date, the casting vote has never been used and its use is expected to be exceedingly rare in practice. In public elections, the drawing of lots or some other random selection would probably be desirable instead.

2 Background: The Debian Project

The development of the worldwide Internet has made more feasible new forms of social organization on a global scale. Not surprisingly, among the first people to realize the potential value of the Internet as a collaborative environment were the computer scientists, students, and hackers[6] who wrote the software and designed the hardware that made the Internet possible in the first place. However, many of these projects were either centered on development teams located geographically in one place, often within a university or corporate research lab setting (such as AT&T's Bell Labs, various research groups at the Massachusetts Institute of Technology, and the University of California at Berkeley's computer science department).[7]

One of the first large-scale projects to break this mold was the development of the Linux operating system kernel[8] to implement the core of the UNIX operating system, which at the time was mostly used on expensive workstation and minicomputer hardware, for common desktop PCs as free software. Linux's creator, Linus Torvalds, was a computer science student studying at the University of Helsinki, but his project quickly gained contributors from around the world via Usenet news[9]. Torvalds became, and remains, the leader of the development community; while anyone can contribute changes—new features, fixes for problems, and the like—to the Linux kernel, he has the ultimate power to decide

[6] I use the term "hacker" in the non-pejorative sense, referring to computing enthusiasts in general, as opposed to the narrower sense associated with computer crime, a.k.a. "cracking."

[7] For more discussion of the evolution and organization of open source software projects, see Raymond (2001) and Coleman (2012).

[8] In computing, the "kernel" of an operating system provides the basic services, like network and disk access and interaction with input and output devices like keyboards and monitors, for software running on a computer.

[9] Usenet news is a worldwide computer messaging system, similar to electronic mailing lists, which is the precursor to today's web-based forums.

to accept or reject a particular change.

Many other projects today based on distributed development have followed the Linux kernel's lead, with a single leader having ultimate authority over the project; somewhat tongue-in-cheekly, this person is often called the "benevolent dictator-for-life" (BDFL), after the title adopted for Python language creator Guido van Rossum in 1995 (Rossum, 2008). Other projects have operated with a "core team" model based on a self-perpetuating steering committee, with new members occasionally being recruited by the existing members. In either scenario, the only option for someone whose ideas are not accepted by the BDFL or core team is to create their own derived project (known as a "fork") and try to recruit enough developers to make the derived project sustainable.

The Debian project started out much the same way as the Linux kernel: a leader (Ian Murdock) recruiting people to help with a project he envisioned; in this case, a project to make a better "distribution" of the Linux operating system (Fernández-Sanguino et al., 2012). A distribution combines Torvalds' Linux kernel with other free software to produce a full computing environment for end-users.

At the time Murdock proposed his project, these distributions were primitive at best: typically, the user would download a kernel binary, some basic tools[10] to get a minimal system running, and a giant archive (like a ZIP file) containing the rest of the software that would be manually installed onto the user's computer from within the minimal system. Murdock proposed instead an approach that relied on each component of the system being stored as a single "package," which could be installed or removed separately, and which

[10]These tools would include a boot loader and a "ramdisk," essentially a file that would act like a small hard drive containing the most basic commands needed to start the system. The kernel, boot loader, and ramdisk would get the user to a basic, but functional "shell," the UNIX equivalent of the DOS prompt.

would contain information about what other packages it needed to work. Any component of the system could then be updated on its own, which means users wouldn't have to download and install a whole new system archive or recompile software every time a new version was released (Fernández-Sanguino et al., 2012).

Murdock's distribution, which he named "Debian" as a portmanteau of his first name and his then-wife's name (Debra), would also prove to be innovative as a model of collaboration. From the start, Murdock favored an open development process, free of commercial control, "in the spirit of Linux and GNU" (Murdock, 1994). While initially the project followed the BDFL model, with Murdock as leader, succession came relatively early when Murdock stepped down in 1996 and appointed a replacement, Bruce Perens. Perens himself was replaced in 1998 by Ian Jackson, under whose watch the first Debian Constitution was adopted (Jackson, 1998), which called for annual elections for the position of project leader. These elections, and other project votes, were to be conducted using "Concorde [sic] Vote Counting," with the single transferable vote[11] to be used if a tiebreaker was necessary. It is unclear exactly why this procedure was adopted,[12] although presumably there was concern that there would regularly be multiple, viable candidates for the Debian Project Leader (DPL) position, or that ballots on other issues might contain multiple potential alternatives, and thus an electoral system considered more fair than simple plurality would be necessary. In 2003, this procedure was itself amended (Srivastava, 2003) to adopt Cloneproof Schwartz Sequential Dropping, an implementation of the Schulze method, after some pathologies and contraditions (which had not yet arisen

[11] In practice, STV with a single position to be filled devolves to an instant runoff vote (Farrell, 2001).

[12] An in-depth examination of the archives of the `debian-devel` mailing list from the period of the constitution's consideration and adoption finds no evidence of any discussion of the voting procedure itself, much less any rationale for it.

in practice) were identified in the vote counting system specified in the constitution (Buck, 2003). The amended procedure continues to be used to this day.

3 Condorcet in Practice within Debian

Debian's Condorcet procedure has three characteristics that are not necessarily used in other systems with ranked voting. First, voters are free to give an equal ranking to more than one option, expressing indifference between those outcomes, and unranked alternatives are implied to be a tied ranking below all voted alternatives. For example, in a contest between candidates A, B, and C, if the voter submits a ballot "A = 1; B = 2; C = 2," this ballot would be considered valid and indicating that the voter prefers A to either B or C, but that the voter has no preference between B or C.

Second, each ballot (whether for a substantive vote or electing the DPL) includes as its last option "further discussion." In the event "further discussion" wins the election, in the case of a DPL election the election will be re-run, with a new nomination period beforehand; in a proposal to amend the Debian Constitution or take some other action, "further discussion" amounts to a rejection of all of the proposals on the ballot in favor of reopening the (virtual) floor for discussion of additional proposals. In the context of a DPL election, this permits voters to indicate that a particular candidate is unacceptable by ranking him or her below the "further discussion" option.

Finally, there is a "quorum" requirement that excludes options that do not receive a minimum number of votes, based on the number of eligible voters, and certain votes (such as amendments to the Debian Constitution and other foundational documents) require

supermajority support to carry; the latter issue does not occur in DPL elections but has arisen in other votes. However, Voss (2012) points out that the way the quorum requirement is implemented in Debian votes may occasionally lead to the premature elimination of options that could influence the vote's outcome.

As the prospect of fradulent ballots transmitted over the Internet is extremely high, votes in Debian's elections must be cast by electronic mail using a ballot that is crypotographically signed using the individual voter's secret PGP key; each voter also uses this key for other interactions with the project (such as uploading new software to the distribution and managing their user account), and the keys are validated to belong to the person they purportedly belong to by one or more existing developers before an individual is allowed to join the project as a voting developer. At the conclusion of each vote, voters can verify that their individual votes were tallied correctly and anyone can validate that the public tally sheet matches the outcome of the election as announced by the project secretary.

Since 1999, Debian has held 13 elections for the project leader position. Of these, 12 elections were contested. Unlike other votes within Debian, elections of the project leader are conducted by secret ballot; only the project secretary knows which votes come from which voters. A summary of these 13 elections appears in Table 1.

The election results from the table above reflect a number of important features of Debian elections over the years. One notable feature is the number of elections (8 of 13) in which an incumbent DPL did not seek reelection. However, as is common in public elections, incumbents who *did* seek reelection were generally successful (only in one year, 2003, did an incumbent seeking reelection—Bdale Garbee—fail to get reelected, among the

Year	Candidates	# < Further Discussion	Incumbent	Winner	Margin
1999	4	0	No	Wichert Akkerman	94
2000	4	0	Yes	Wichert Akkerman	94
2001	4	0	No	Ben Collins	37
2002	3	0	No	"Bdale" Garbee	113
2003	4	0	Yes	Martin Michlmayr	4
2004	3	1	Yes	Martin Michlmayr	105
2005	6	1	No	Branden Robinson	23
2006	7	2	No	Anthony Towns	6
2007	8	2	No	Sam Hocevar	8
2008	3	0	No	Steve McIntyre	53
2009	2	0	Yes	Steve McIntyre	43
2010	4	0	No	Stefano Zacchiroli	147
2011	1	0	Yes	Stefano Zacchiroli	380
2012	3	0	Yes	Stefano Zacchiroli	310

Table 1: Summary of Debian Project Leader Elections, 1999–2012. Source: The Debian Project (2012)

six elections where an incumbent was a candidate). A proliferation of candidates is also fairly common; only two elections had 2 or fewer candidates, with a mode of 4 candidates across all elections. Given the absence of political parties, an electoral system that does not turn candidates with low support into potential "spoilers" who might be discouraged from running, and the fairly open nomination process, the high number of candidates on average is not particularly surprising.

Another notable feature of these elections has been the widely varying margins of victory over the second-place candidate; three elections have seen the winner's margin over the second-ranked candidate be eight votes or fewer. On the other hand, the three most recent elections have seen the current DPL, Stefano Zacchiroli, win by overwhelming margins of 147 votes or higher. We can also note that substantial numbers of voters have been able to indicate in several elections that certain candidates were unacceptable (indicated in the table as "# < Further Discussion"); some of these candidacies were

obvious novelty or parody efforts, in line with those of Britain's Monster Raving Looney Party, while others had either not established much of a reputation within the project or had a track record of problematic behavior.

4 Lessons for Public Elections

There are some important caveats to any effort to apply the lessons of the Debian Project in a public election setting. Debian is a voluntary organization; governments, by and large, are not, particularly at higher levels of organization—in most societies, people do not get to choose the government jurisdiction they will fall under,[13] but instead a single government holds a geographic monopoly on the provision of governmental services. Thus the "exit option" (Hirschman, 1970) is much more viable for a person dissatisfied with Debian than a public government, and perhaps it is more important for the DPL to be broadly popular than a leader in a public setting (particularly at the national or state level) as a result.

This distinction suggests Condorcet methods may be particularly well-suited to elections to offices that are designed to either mediate between conflicting parties or promote a consensus-based approach to decision-making. For example, in countries or political subdivisions using a parliamentary system of government, it might be desirable to elect the head of state and legislative speakers—offices that typically rely on a reputation of being neutral or fair for their effectiveness—using a Condorcet method like Schulze's. Similarly, to the extent the election of judges, court officers, and election administators is

[13]Even in societies where family law is based on religious authority rather than civil authority, and parallel systems of general and family law exist, like Egypt, citizens rarely have the option to opt into a different family law system, but instead are born into a particular one based on the religious affiliation of their parents.

desirable (as is common practice at the local and state level in U.S. states), Condorcet methods may be preferable in promoting the election of candidates who can command broad community support.

On the other hand, single-winner Condorcet methods might be less appropriate in cases where the purpose of the election is to choose representatives who will resolve conflicts in an elected body, although a Condorcet winner might be considered more "representative" of their geographic constituency than a plurality winner. In these cases, an electoral system using elements of proportional representation (perhaps with some members selected from single-member districts to ensure geographic representation) might be appropriate; alternatively, Schulze (2011b) suggests a modified version of the single transferable vote that would combine proportionality with the benefits of a Condorcet method.

One concern about Condorcet methods that has been expressed is the fear that it may lead to the election of obscure or "spoiler" candidates who lack core support (see, for example, Dennis 2008). In addition to the behavioral argument that relatively few voters would actually express a higher preference for an obscure candidate they knew little or nothing about over one they knew more about and would find acceptable, Debian's process includes the "further discussion" option that functions in leadership elections effectively as a "none of the above" option. As noted above, weak candidates have ranked below "further discussion" in past Debian elections, indicating that joke and obscure candidates can be effectively weeded out with this provision. In public elections, a requirement that the election be re-run with a new slate of candidates if "none of the above" wins would effectively preclude candidates without much support from gaining public office.

Alternatively, a system modeled after the "deposit" used in the United Kingdom and other countries could be adopted to discourage the participation of obscure candidates, in which only candidates who receive more votes than "none of the above" would receive a refund of their qualifying fees.

Ultimately the biggest obstacle to the adoption of Condorcet methods, or even any other electoral reform, may be public acceptance. Smith (2006) describes two surveys of U.S. voters in which the respondents expressed a preference for plurality voting over both approval and range voting, suggesting that (at least in the United States) the perceived complexity of fairer voting systems outweighs their potential support. Given the existing length and frequency of U.S. ballots, which often leads to long lines at general elections, this finding may not be particularly surprising, and indicates that any successful move to a fairer electoral system should be paired with a reduction in the number of elected offices, in order to ensure voting queues remain manageable.

5 Conclusions

In this paper, I have examined the experience of the Debian Project with the use of a Condorcet method for its elections since 1999. The findings suggest that Debian's use of the Schulze method has generally been effective in choosing project leaders with fairly broad support, even from fields with a large number of candidates and containing potential spoilers. I have also discussed some of the ways in which Debian's use of elections differ from their use in public settings, which may limit the generalizability of these findings to potential public elections.

Given the increasing number of organizations using Condorcet methods, including the Schulze method, there may be greater opportunities to research how ranked choice voting systems work in practice. These opportunities may come in tandem with more use of these methods in public elections in the United States and elsewhere, as increasing public and elite concern over the distortions in electoral outcomes resulting from gerrymandering and single-member district systems may lead to greater pressure for electoral reform from the public and the judiciary.

References

Arrow, Kenneth J. (1970): *Social Choice and Individual Values*. 2nd. New Haven, Ct.: Yale University Press.

Buck, Buddha (2003): *Re: Call for votes for the Condorcet/Clone proof SSD voting methods GR*. URL: http://lists.debian.org/debian-vote/2003/06/msg00036.html.

Coleman, E. Gabriella (2012): *Coding Freedom: The Ethics and Aesthetics of Hacking*. Princeton University Press.

Dennis, Greg (2008): *Why I Prefer IRV to Condorcet*. URL: http://www.fairvote.org/why-i-prefer-irv-to-condorcet/.

Droop, Henry Richard (Oct. 2007 (1881)): "On Methods of Electing Representatives". In: *Voting Matters*, no. 24, pp. 7–46. URL: http://www.votingmatters.org.uk/ISSUE24/I24P3.pdf.

Farrell, David M. (2001): *Electoral Systems: A Comparative Introduction*. Palgrave Macmillan.

Fernández-Sanguino, Javier, Bdale Garbee, Hartmut Koptein, Nils Lohner, Will Lowe, Bill Mitchell, Ian A. Murdock, Martin Schulze, and Craig Small (2012): *A Brief History of Debian, version 2.17*. URL: http://www.debian.org/doc/manuals/project-history/.

Hirschman, Albert O. (1970): *Exit, voice, and loyalty: responses to decline in firms, organizations, and states*. Cambridge, Mass.: Harvard University Press.

Jackson, Ian (1998): *Historical version of the Constitution for the Debian Project (v1.0)*. URL: http://www.debian.org/devel/constitution.1.0.

Murdock, Ian A. (1994): *The Debian Manifesto*. URL: http://www.debian.org/doc/manuals/project-history/ap-manifesto.en.html.

Raymond, Eric S. (2001): *The Cathedral & the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary*. O'Reilly Media.

Rossum, Guido van (2008): *All Things Pythonic: Origin of BDFL*. URL: http://www.artima.com/weblogs/viewpost.jsp?thread=235725.

Schulze, Markus (2011a): "A new monotonic, clone-independent, reversal symmetric, and Condorcet-consistent single-winner election method". English. In: *Social Choice and Welfare*, vol. 36, iss. 2, pp. 267–303.

— (2011b): "Free Riding and Vote Management under Proportional Representation by the Single Transferable Vote". URL: http://m-schulze.webhop.net/schulze2.pdf.

Smith, Warren D. (2006): *Descriptions of single-winner voting systems*. URL: http://home.versanet.de/~chris1-schulze/votedesc.pdf.

Srivastava, Manoj (2003): *Constitutional Amendment: Condorcet/Clone Proof SSD Voting Method*. URL: http://www.debian.org/vote/2003/vote%5f0002.

Szpiro, George G. (2010): *Numbers Rule: The Vexing Mathematics of Democracy, from Plato to the Present*. Princeton University Press. Princeton University Press.

The Debian Project (2007): *Constitution for the Debian Project (v1.4)*. URL: http://www.debian.org/devel/constitution.

— (2012): *Debian Voting Information*. URL: http://www.debian.org/vote.

U.S. National Archives and Records Administration (2012): *U.S. Electoral College: Frequently Asked Questions*. URL: http://www.archives.gov/federal-register/electoral-college/faq.html#wtapv (visited on 12/31/2012).

Voss, Jochen (Aug. 2012): *Quorum in the Debian Voting System*. URL: http://www.seehuhn.de/pages/quorum (visited on 12/31/2012).