Abstract

Elections are presented as analogous to digital signal processing systems, sampling continuous-time analog inputs at regular intervals and processing them into a discrete-time quantized signal. Due to how quickly the preferences of voters can change relative to how often elections are held, violations of the Nyquist criterion are inevitable. Examples of how the resulting aliasing effects manifest in elections are presented.

A system of higher-frequency elections is proposed to reduce aliasing effects. As this would result in a government strongly beholden to the instantaneous whims of the populace, further modifications are considered to add a systemic bias towards stability. These modifications include the addition of a supermajority threshold to implementing change, integrating sequential election results over time to reach that threshold, and adding an increasing bias towards change over time to emulate term limits.

Quantization error is presented to have negative effects in the electoral realm. It is proposed that these effects could be countered by injecting a random element into the determination of the outcome of the election, comparable to dithering in image processing. Simulation results indicate that this overcomes the effect of gerrymandering while maintaining single-winner districts.

Keywords

elections; engineering; digital signal processing; stochastic; nondeterministic; voting; lottery; term of office

Introduction

Any election is, fundamentally, a system for taking input, and using those to select one of several predefined output states. The inputs are the stated preferences of the individual voters, and the outputs are the possible outcomes of the election.

There are numerous important variables in this process: who gets to vote? Are votes indicated as pick-one ballots, approval ballots, ranked choices, scores? How are possible outcomes defined? How are the inputs actually translated into an output? None of which even considers manifold real-world details like ballot order, paper or electronic ballots, system transparency, or direct and opportunity cost to the voter. The space of possible election system designs is immense.

However, all election designs have one common factor: the election event measures the inputs at a particular moment in time, the output is defined based on those inputs, and the inputs have minimal impact on the outputs until the next election event. This is exactly analogous to the operation of a digital signal processing (DSP) system.

Continuous and Discrete Signals

In the realm of electrical engineering, signals are categorized as either continuous, or discrete. A continuous signal can change over arbitrarily short timescales; the signal could have one value, and then change dramatically a nanosecond later. A continuous signal can also assume any value, rather than a limited set of pre-determined values. Continuous signals are phenomena we see in the physical world, like sound, temperature, and voltage.

A discrete signal, by comparison, is quantized in both time and value. The output can only change at certain predefined intervals, and the output can only assume certain predefined values. Digital computers, for the most common example, are built of millions of individual elements, each of which can only assume the values 0 or 1. These values can only change at fixed intervals. The time between these updates defines the operational frequency of the computer, usually in gigahertz (GHz).

The preferences of individual voters can be treated as continuous signals, able to change at arbitrary times, and able to assume any value imaginable. The outcomes of the election can be treated as discrete signals, unable to change between update times, and unable to assume any value other than those pre-defined.

To interface with the continuous signals of the physical world, digital systems *sample* their inputs at each update time, and act on those sampled values until the next update time, regardless of how the inputs change in between.   This sampling process is equivalent to an election.

We will examine the implications of this analogy by applying digital signal processing principles to elections. We will observe how problematic aspects of electoral processes are further analogous to errors in the signal processing domain, and explore how means of correcting these errors in the signal processing domain may be translated to the electoral domain.

Nyquist and Aliasing

The Nyquist criterion is a fundamental mathematical law of digital sampling of continuous signals. In effect, it states that if a continuous signal is being sampled at discrete intervals, there is a minimum rate at which that sampling must take place to avoid loss of information and distortion of the sampled signal. That minimum sampling rate is related to the rate at which the input signal can change.

For example, the outdoor temperature can change by on degree over, perhaps, ten minutes; because of this, a digital system must sample the temperature no less often than every ten minutes to form an accurate digital model of how temperature was changing. If we only took readings every five hours, our resultant model would be inaccurate, because it could not include all the changes in temperature that actually took place.

The impact of a Nyquist violation of this sort is called *aliasing*; rapid changes in the inputs are misrepresented in the output as long-term effects. See Figure X for an example. In this simulation, the input signal can assume arbitrary values, and can change at arbitrary times. The input line is populated with random generated data for the purposes of illustration. For sampled outputs, two different sampling periods are used, to illustrate the differences in their effects. In one, the sample period is once every 24 seconds. In the other, the sample period is once every 72 seconds.

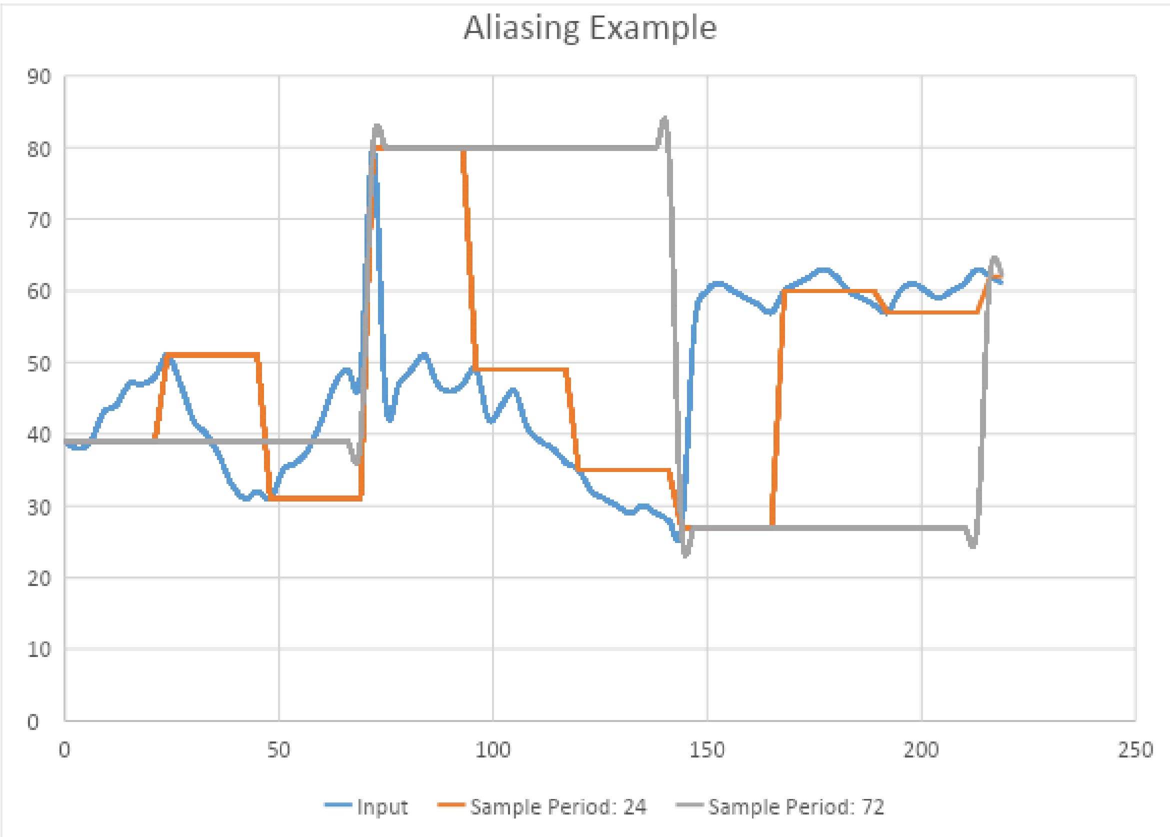

We can see that, for the first 72 seconds, the input rises, then falls. The 24-second sampling follows this input reasonably well, though with some delay. The 72-second sampling ignores all changes in the input. Depending on one's intent, both these effects may be desirable; perhaps one desires a system that follows short-term changes in the input with a delay, or one that entirely ignores short-term changes in the input, focusing only on long-term trends. A similar effect can be seen at the 144-second mark. Immediately after that point, there is a permanent long-term change in the input. The 24-second sampling system catches up with the change 24 seconds later, and the 72-second sampling system takes 72 seconds, as expected.

Unfortunately, as we see at the 72-second mark, this system design has a serious flaw. At 72 seconds, there is a momentary, extreme swing in the input signal. The input immediately returns to the previous value, but the swing happened to come at the same time as a sample was taken, and the outputs are completely unrepresentative of the inputs for 24 or 72 seconds, respectively. Rather than ignoring a short-term variation in the input, the output represents nothing *but* a short-term variation in the input. This aliasing effect exists because of the mismatch between the variation speed of the input and the sampling rate. The greater the mismatch, the worse the effects become.

We can quantify the representational error of the system by measuring the difference between the system input and each of the two system outputs, and computing a mathematical value called the root-mean-square (RMS) error. For each sample point, we subtract the input signal from the output, to create a number we call the *error* for that point in time. Once all the errors are computed, we square each value, sum those squared values, and take the square root of the sum. This process means that errors in a positive direction or a negative direction are treated only by their magnitude; an error in the positive direction does not cancel one in the negative direction.

In a perfectly responsive system, the RMS error would be zero at all times. In the system we see here, the RMS error of the 72-second sampling over the first 72-second period is 31, compared to the 24-second sampler's error of 53. The longer sampling period has resulted in a more accurate average output.

However, in the second 72-second period, the higher-frequency 24-second sampling gives an RMS error of 91, while the lower-frequency 72-second sampling gives an error of 197. During the third 72-second sampling period, the RMS errors are 84 and 158 for the higher-frequency and lower-frequency sampling, respectively. In these cases, the longer sampling period has resulted in a dramatically *less* accurate average output. This inaccuracy is the effect of aliasing; the rapid changes in the input are misrepresented in the output, and the misrepresentation gets worse the less often samples are taken.

Rather than sample less often, the solution used in DSP systems is generally to sample *more* often, but to make the system respond more slowly to changes in the input. This is generally known as a *low-pass filter*. There are numerous possible filter implementations, but the simplest version (a single-pole infinite-impulse-response filter) is implemented in Figure X for demonstration purposes.

The same input data is used from our original example, but our two outputs are now both sampling every 3 seconds. Instead of relying on sample times to act as a filter, the output at each sample time is a weighted average of the most recent sample, and the previous value of the output. In our examples, one output is 60% previous value, 40% input; the other is 90% previous value, 10% input. Making the output at each moment include a fraction of the previous value of the output gives a bias towards stability, resisting rapid changes in the inputs, while still giving some degree of immediate response.

The difference in the RMS errors is striking, with the error over any given period being reduced by between 40 and 80%.

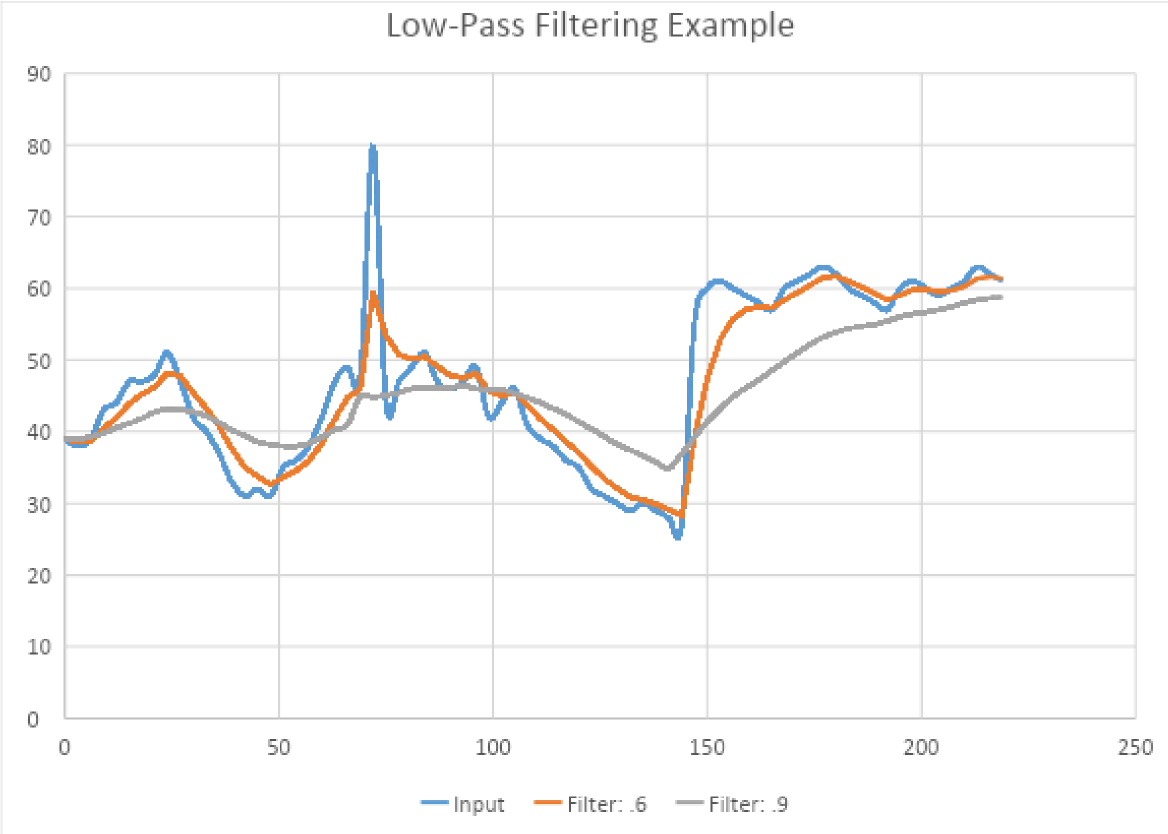

Now, by analogy, we can change the labels on our first graph and observe the effects in an election system. Instead of seconds, consider the time increments to be months. The 24-month sample period is the time between elections for the United States House of Representatives. The 72-month sample period is the time between elections for the United States Senate. Identical aliasing effects can be observed in election systems.

Suppose that at our 72-month mark, an unexpected major information release occurs immediately before an election (a so-called “October surprise”), causing a sudden shift in voter preferences that alters the outcome of the election. Immediately after the election, it is determined that the information released was false. Voter preferences return to their previous values, but it is too late; the electoral outcome based on a momentary shift in voter preference will now hold for multiple years.

In an alternate example, suppose an election takes place with no unusual events affecting the outcome. Instead, immediately after the election, the winner declares that he has received a divine revelation, abandons all previous policy positions, and declares his support of legalized slavery. Voter preferences would (one hopes) change very rapidly to oppose the elected official, but this change would not be represented in the electoral outcome for multiple years.

During the writing of the Constitution of the United States, the framers decided that the term of office of the Senate should be longer than that of the House, in order to insulate the behavior of the Senate from short-term changes in public preference. From our simulation, we can see that this works as desired in some cases, but in others, the aliasing effect gives highly un-representative outcomes. Since the timing of sudden changes in voter preferences are at best random, or at worst biased towards occurring shortly before or after an election, there is no way of guaranteeing that these aliasing events will not take place. The longer the period between elections, the greater the unrepresentative effects become. In extreme cases such as Brexit, the outcome of a single election event can have permanent and enormous effects on the entire society, with no chance of undoing the change on any timescale.

In short, the design of the United States Senate, using six-year terms to insulate the Senate from changes in public preference, comes at a cost of dramatically increased error in representation. More broadly, using length of term in office as a filtering mechanism against day-to-day changes in public mood is flawed, on a fundamental mathematical level.

This does not mean the goal is to be abandoned. Insulating elected officials from temporary changes in public mood by means of a delayed system response is a reasonable design choice. But the delay mechanism selected, that of term length, inevitably introduces substantial loss of representative *accuracy* due to aliasing effects, and should therefore be reconsidered.

In order to minimize these aliasing effects, elections must come closer to compliance with the Nyquist criterion: elections must take place *more* frequently in order to accurately represent changes in the preferences of the electorate. And some form of low-pass filtering must be introduced, to insulate the actions of government against the momentary mood of the voters.

High-Frequency Hysteretic Elections

We will now consider the side-effects of holding elections much more frequently than in our present system. For purposes of simplicity, we will discuss holding elections monthly. This is only for convenience, and not an actual proposal for the optimal frequency of elections.

Several analyses have been conducted on the effect of term lengths and the proximity of elections on the behavior of elected officials. (Ahuja, 1994) (Gaines, Nokken, & Groebe, 2012) (Bó & Rossi, 2011) (Amacher & Boyes, 1978) (Schultz, 2008) (Titiunik, 2016) However, these analyses generally involve terms of multiple years. (Fischer & Suleiman, 1997) analyzed the effects of different term lengths on the evolution of cooperation between groups. (Tridimas, 2017) addressed high-frequency elections in the direct democracy of ancient Athens, with elections taking place dozens of times each year at the extreme. But the practical implications and implementation details of high-frequency elections in a modern representative republic have not previously been addressed.

The first obvious effect is that of increased cost to the state. At present most jurisdictions hold some manner of elections once or twice a year, and the voter turnout to most of those elections is very low. If instead all offices were subject to election every month, cost to the state would presumably increase by at least a factor of ten.

A similar effect is that of increased cost of participation to the voter. In some underserved jurisdictions voting may require standing in line for hours. For low-wage workers the opportunity cost of voting instead of working can be significant. Increasing the frequency of elections also increases these opportunity costs. These topics are areas deserving further study.

To eliminate aliasing, the electorate must be able to effect change, i.e. remove an official or undo a ballot measure, rapidly. To retain stability there must be high barriers to such reversals; there must be some bias towards the present state of affairs. Systems with discrete states having such a bias are common in engineering and physics, and are said to exhibit *hysteresis*.

For example, suppose a thermostat is set to 70 degrees. The furnace can't put out 70 degree air; it puts out 100 degree air, or nothing. One way to achieve 70 degree temperature in the room would be to start heating if the temperature is even slightly below 70 degrees (69.99999…), and then stop heating if the temperature is even slightly above 70 degrees (70.0000…1). But this results in very short heating cycles. The system will be constantly switching between an on state and an off state, spending no time in either. This is undesirable for a number of reasons.

In a more realistic control system, the furnace might start to heat the room if the temperature falls below 69, and then stop heating above 71. If the temperature is below 69, the heat is definitely on. If the temperature is above 71, the heat is definitely off. But in between 69 and 71, the heat simply keeps doing the last thing it was doing; in that hysteresis range, the system is biased towards stability.

Consider a similar system for elections. At each election event, the electorate votes on whether to *retain* or *replace* each elected official (with equivalent language for ballot measures). *Replace* can only win if it reaches a predefined *supermajority threshold*. If *replace* wins, a later election is held to determine who will replace the outgoing official.

In this system the supermajority threshold effectively dictates how difficult it is to remove an elected official, and therefore how closely that official must follow momentary public opinion to retain his office. Different offices would presumably have different thresholds, depending on how responsive that office is intended to be to the immediate will of the electorate. For the purposes of discussion, we will suppose that a senator has a threshold of 60%, while a congressman has a threshold of 53%.

Example 1: A new Senator is elected. Each following month the voters are asked to vote to "replace" or "retain" that Senator.

Month 1: 54/46 retain. The senator is still relatively popular.

Month 2: 55/45 retain.

Month 3: 53/47 retain.

Month 4: 50/50 tie.

Month 5: 49/51 replace. After this election, the senator is caught on tape disparaging a popular public figure.

Month 6: 45/55 replace.

Month 7: 42/58 replace. The senators' supporters start to worry and show up to the next election in greater numbers.

Month 8: 44/56 replace. The senator is caught on tape again saying undesirable things. His supporters don't bother showing up next month.

Month 9: 39/61 replace, triggering a replacement election.

The bias towards stability is present; the Senator was able to retain his office despite momentary public opposition. But there is an obvious flaw in this system: the bias towards stability is overwhelming. The senator could survive having 59% of the electorate in favor of replacing him, *forever*. The majority vote *must* eventually win out in simple electoral matters; we sacrifice this principle at the risk of destabilizing the entire electoral system. We have overcome aliasing to a substantial degree, by allowing the system to have much faster response, but the purely hysteretic system requires further modification.

Integrating Electoral Outcomes

We require a system that is biased towards stability in the short term, but in longer terms still implements majority rule. To that end, we propose integrating (that is, summing) electoral outcomes over time. In this system, the difference between the "replace" and "retain" percentages at the end of each election is added to a running total, limited to some pre-defined range. If this running total ever reaches the limit of its range towards "replace", a new election is held to replace the official.

Having been expressly voted out of office, the replaced official could reasonably be barred from participating in the replacement election. This would universally eliminate any incumbent advantage.

In this system the time it takes to effect a political change is entirely dependent on the size of the majority favoring that change. Suppose senators have a point range of +-20, and congressmen have a point range of +-6. A senator starting with zero points and opposed by a consistent 51% of the electorate would be removed after 20 months, while a congressman opposed by that same margin would be removed after only six months.

Conversely, a large enough supermajority can always remove an official after a single election. A popular senator with the maximal score of 20 needs to lose 40 points to be removed. If his popularity shifts to a 51% opposition, it would take 40 months to remove the senator. But a shift to 70% opposition could remove him after a single month. The speed of the system response is partially determined by the magnitude in the shift in popular attitude.

In theory, it would be possible for the elected officials to be removed and replaced every month, resulting in substantial difficulty in the always-freshman representatives being unable to accomplish anything before being replaced. In practice, however, this would require that an individual win election as a replacement official, and then immediately have 70% of voters vote against them, in a significant portion of the total number of elections. This would imply that the elections selecting replacement candidates were designed such that they regularly selected highly unpopular candidates. Avoiding the use of plurality voting in favor of majoritarian methods would seem to greatly reduce the probability of such outcomes. We also discuss a stochastic method later in this paper, with more complex implications.

Another effect is that the elected official can literally build up quantifiable political capital, measured by their electoral scores. Unpopular actions would have immediate quantifiable results, and the official would have clear warning that their choices may soon cost them their job. But they can also take unpopular actions they deem necessary if they have enough capital built up to take the hit, especially if they believe they will be proven right in the longer term.

Finally, the electorate has clear and legally meaningful indication of the likely outcome of the election before it happens, without relying entirely on unofficial polls. In some cases under present systems, many potential voters do not participate (or they vote for a third party) because they believe the result of the election to be a foregone conclusion, and that their vote will have zero effect on that outcome.

In our proposed system, these drivers of voter apathy are reduced. Voters are more likely to vote if they believe their vote has an impact. Each monthly election event has a quantifiable, long-term impact on the electoral prospects of that official. If an official's cumulative score is changing month-by-month, both opposing and supporting voters are incentivized to vote, to precipitate or prevent a recall election.

Example 2: A new senator is elected. His score starts at zero, and the scores for senators are limited to between 20 and -20.

Month 1: 54/46 retain. Score: 8.

Month 2: 55/45 retain. Score: 18.

Month 3: 53/47 retain. Score: 20 (maximum).

Month 4: 50/50 tie. Score: 20.

Month 5: 49/51 replace. Score: 18.

Month 6: 45/55 replace. Score: 8.

Month 7: 45/55 replace. Score: -2.

Month 8: 45/55 replace. Score: -12.

Month 9: 45/55 replace. Score: -20, triggering a replacement election.

This system gives a degree of election-to-election stability, while still retaining majority rule on longer timescales. This overcomes the primary theoretical difficulty with holding elections more often, while still allowing us to mitigate aliasing effects by ignoring brief changes in public opinion.

Term limit are often proposed as a solution to various electoral and political problems. Their utility as an appropriate solution to any of those problems is debatable, and it is not the purpose of this paper to address any of those issues. However, should a similar bias towards long-term change be deemed desirable by those designing a hysteretic election system, this can be accomplished by adding a bias factor to each election result, with the bias factor increasing over time. As the bias increases, the candidate would eventually require continual supermajorities to retain their office. This has the effect of allowing experienced legislators to remain in office, but only with exceptional popularity and service records.

Example 3: A senator starts his fourth year in office with a score of zero. The scores for senators are limited to between 20 and -20. Each month some bias towards replacement is added to the election results, the bias increases by one point every year. So in this year the bias is -4, and after twelve months the bias will increase to -5..

Month 1: 53/47 retain. Bias: -4. Score: 2.

Month 2: 52/48 retain. Bias: -4. Score: 2.

Month 3: 51/49 retain. Bias: -4. Score: 0.

Month 4: 51/49 retain. Bias: -4. Score: -2.

Month 5: 51/49 retain. Bias: -4. Score: -4.

.

.

.

Month 12: 51/49 retain. Bias: -4. Score: -18.

Month 13: 52/48 retain. Bias: -5. Score: -19.

Month 14: 52/48 retain. Bias: -5. Score: -20, triggering a replacement election.

Even with a consistent majority voting to retain the senator, “replace” eventually won the election due to the additional bias towards change over longer periods. If the senator’s support base had been larger, his term in office would have been longer.

By application of DSP techniques, we have designed a novel system of elections with improved representative accuracy, finer control of the delay between public opinion and legislative outcome, reduction in voter apathy, and elimination of incumbency advantage. We will next explore other aspects of digital signal processing systems, which when applied to election design can address issues such as winner-take-all elections and gerrymandering.

Quantization Error

Digital systems, by their nature, can only store a finite number of possible values, which places limits on the precision they are able to represent. Due to this, when a sample is taken of the input signal, a DSP system can only approximate the value of that signal by selecting a storable value near the real value of the input. For example, if a system can only store integers, an input of 1.5 must be stored as either 1 or 2. A degree of precision in the input signal, and thus data, is inevitably lost, a phenomenon known as *quantization error*.

Quantization error is well-known in both audio and video signals. Consider Figure 1, which, like all digital photographs, is represented as an array of pixels. Each pixel is itself a single grayscale value; it may be black (0), white (255), or any value in between.

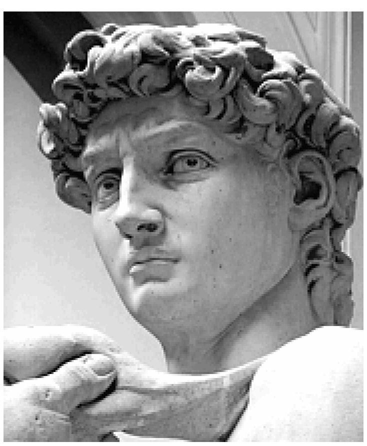

Figure 1 - Image represented in grayscale (Toussaint, 2007)

Now assume we want to represent this same image, but now with only black (0) and white (1) pixels, with no grays in between. To perform this simplification, each individual pixel must be changed to either black or white, and we must define some algorithm to execute this transformation. A simple approach would be to say that darker grays become black, and lighter grays become white. In other words, we *quantize* the analog grays to the nearest allowed value. The exact grayscale value of each pixel is lost, introducing quantization error. The results of this particular approach are not appealing, as shown in Figure 2.

Figure 2 – Image quantized to black and white (Toussaint, 2007)

In this image, we see large areas of black and white, with all gradations lost. Note in particular the ears, which are indistinguishable, having been completely subsumed in a solid black mass.

In elections a similar phenomenon is at work. In any non-unanimous case, the vote tally will not match perfectly with any actual candidate, but the outcome of the election can only be a selection among the available candidates. Due to this, some metric is used to determine which candidate is closest to the vote tally.

In the simplest case, consider a two-candidate election with plurality voting. The percentage of the vote directed at either candidate can be any value between 0 and 100%. The outcome of the election, by comparison, can only assume one of two values: a win for candidate “A”, or a win for candidate “B”. There is a dramatic difference between candidate “A” receiving 51% of the vote or 99% of the vote, but the outcome is that candidate “A” wins, regardless of that difference.

Consider the typical map used in American media to portray outcomes in presidential elections, with red and blue representing the two parties. As pointed out in (Barron, 2016), no state or district votes unanimously for either party, and thus would more accurately be represented as a shade of purple. The results of the 2012 Presidential election are thus represented in Figure 3. However, the nature of the winner-take-all electoral process used in most states converts those shades of purple to either red or blue, as if the minority party did not exist in that district. This is shown for the 2012 Presidential election in Figure 4. Information is lost between the vote tabulation and the outcome of the election, identical to the quantization error we observed between Figure 1 and Figure 2.

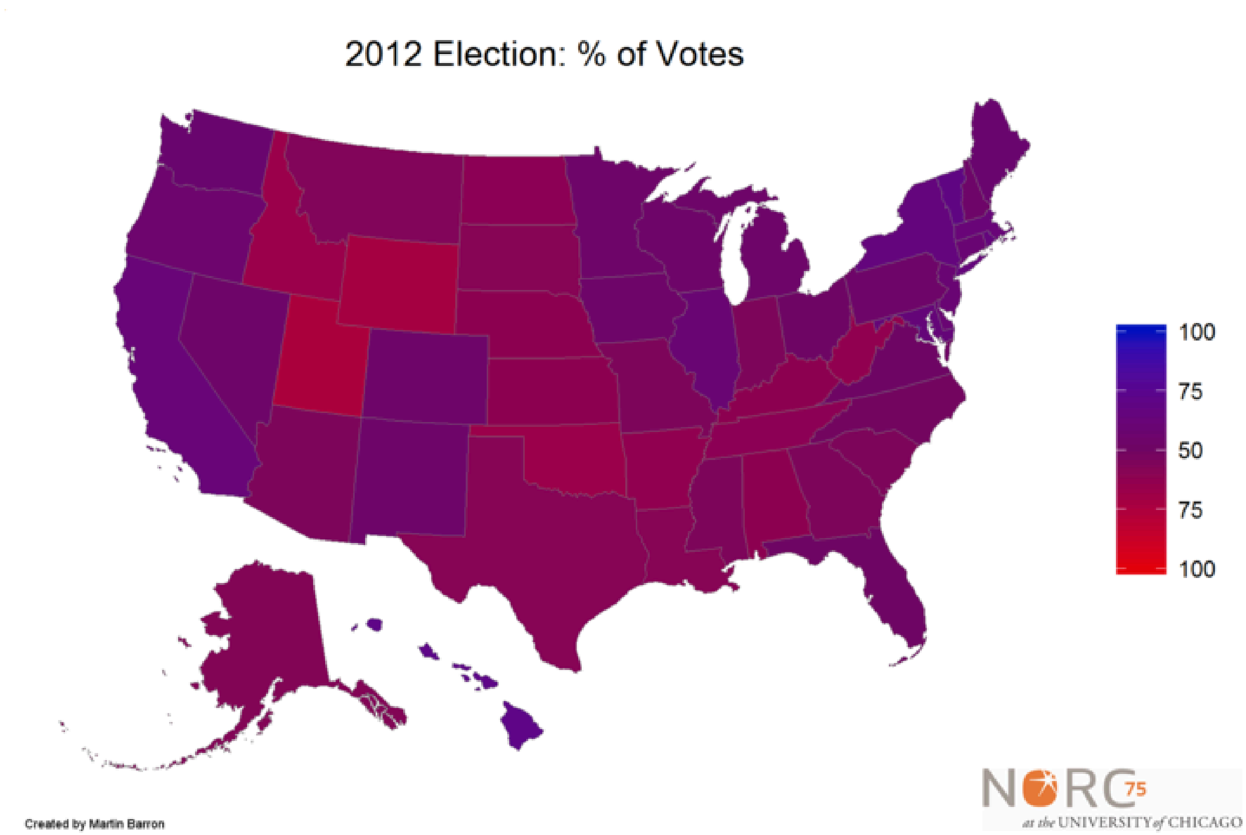

Figure 3 - 2012 US Presidential election results in purple-scale (Barron, 2016)

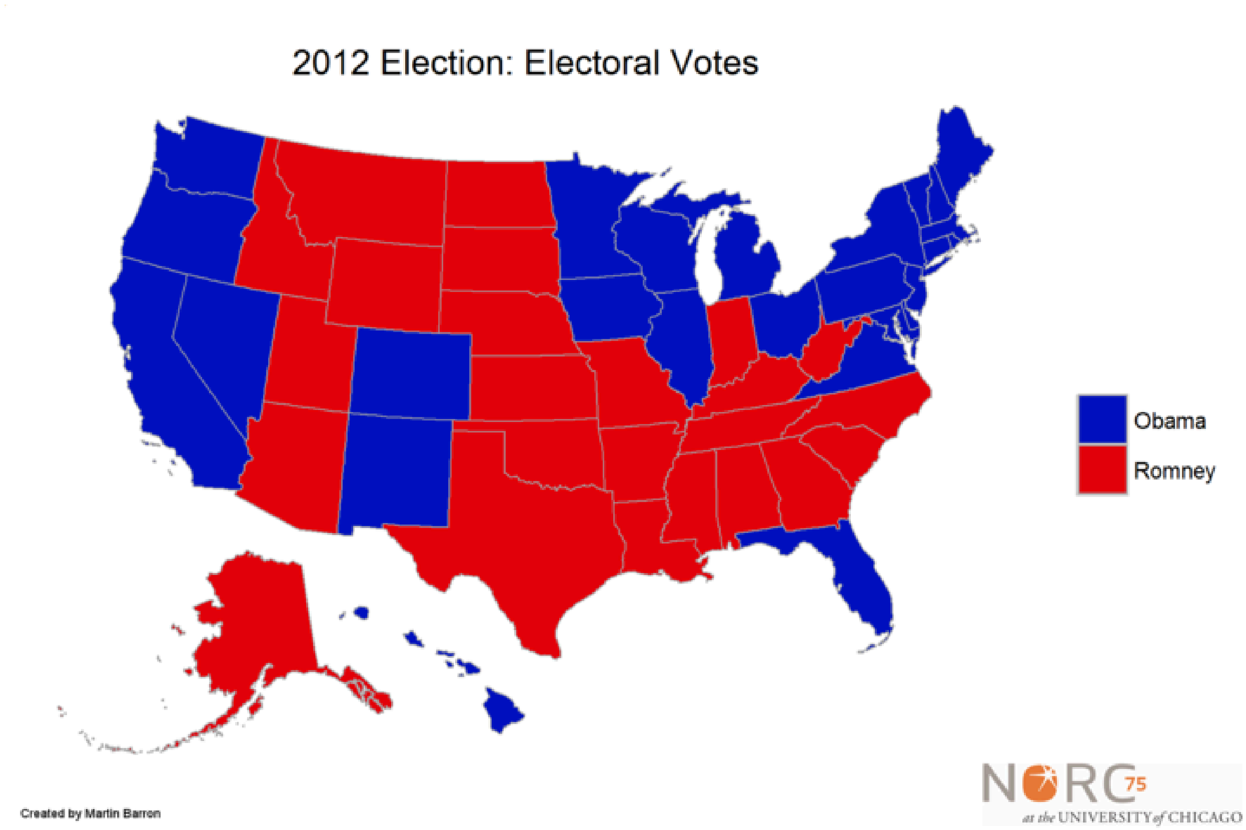

Figure 4 – US 2012 Presidential election quantized to red and blue (Barron, 2016)

The implications of this quantization are notable, particularly with regard to geographically-defined voting districts and gerrymandering. If there is no difference between winning with 51% of the vote and winning with 99% of the vote, there is no incentive for a candidate to build the largest possible coalition, or to do their best to represent all voters equally. Instead, a candidate will tend to ignore the needs of minority groups in their constituency, focusing on remaining the preferred candidate of the majority group. For example, a congressman from a district which consistently votes 60/40 Republican need never pay any mind to the wishes of the Democrat 40% of their electorate, and instead will focus all their attention on being the most appealing Republican possible. In short, quantization error leaves large numbers of voters without representation, and strips their votes of any weight in the practice of governing. The dominant party's primary election becomes the true determinant of who represents the district, with the general election reduced to a formality.

Dithering

Digital signal processing systems deal with quantization error by a process called *dithering*, which injects a degree of random noise into the signal. Without dithering, a pixel that is 80% black and 20% white always becomes 100% black. With dithering, that same pixel is subject to a random selection, with an 80% chance of becoming black, and a 20% chance of becoming white. Over large numbers of samples, this random element causes the blacks and whites to average back to the original gray. Dithering results in a much more representative image, as shown in Figure 5.

Figure 5 - Grayscale image quantized to black and white, with dithering (Toussaint, 2007)

The dithered image is still represented using only black and white pixels, but it much more accurately represents the appearance of the original grayscale image. The addition of a random element to the process distributes the quantization error in a less predictable fashion, resulting in a more desirable translation between the input and the output.

In electoral terms, the implementation of dither would add a degree of randomness to the electoral process, between the determination of the vote tally and the selection of a winner to represent the electorate. (Note that this is distinct from sortition, which does away with voting entirely in favor of every candidate having equal odds of winning.) Various weightings are imaginable, but in a simple model, a candidate who obtained 51% of the vote would have a 51% chance of winning the election. Elections would become stochastic processes, instead of the deterministic processes they have heretofore been. The implications of such a change are worthy of consideration.

First, the concept of a “safe district” becomes meaningless. Under a deterministic system, a district where one party has a 60% majority will always be represented by the majority party, leaving the 40% minority without representation in any future election, barring shifts in population preference over time. Using stochastic elections, the 40% minority would instead select the district representative 40% of the time. Third parties and other political minorities would similarly have representation not afforded them in deterministic elections; a party with the support of 5% of the populace would be expected to win one election out of every twenty.

Secondly, the importance of party primaries is reduced. Under deterministic systems in safe districts, the district representative is decided by the primary election of the majority party, requiring the winner to appeal only to a majority of voters in the majority party, rather than to the district as a whole. This incentivizes candidates to tend toward extremism. Under stochastic systems, candidates would instead be incentivized to build the largest coalition possible within the electorate, in order to maximize their odds of winning. Unlike in single-winner deterministic elections, every vote genuinely has equal probability of affecting the outcome of the election.

Thirdly, even candidates with extreme minority support would have some probability of winning the election. For example, a candidate with .1% of the vote would win one out of every 1000 elections. In a country with 435 congressional districts, there would be an 89% chance of such a candidate winning once in five election cycles.

This consequence potentially gives unpopular candidates disproportionate time in office. Suppose the probabilities of each candidate winning are intentionally devised such that a candidate holds office for a duty cycle approximately proportional to their support among voters; a district split 60/40 between Republicans and Democrats would be represented by a Republican 60% of the time, and by a Democrat 40% of the time. Such a system is comparable to the DSP concepts of pulse-width modulation (PWM) and space vector modulation (SVM), which use a series of pulses between available values to represent an analog signal between those values.

Any such system demands that we also identify the period over which such averaging is taken. If a Republican is to represent a district for 60% of the time, this criterion is equally met by representing six out of every ten seconds, or by six out of every ten centuries. Neither results in a functional or responsive system of government; the length of the term of office becomes a critical factor. For similar reasons, if a candidate with .1% support wins an election with a two-year term of office, their time in office is only proportional to their support if the candidates' time in office are averaged out over two thousand years!

One way of minimizing this effect is to hold elections far more often than we presently do, with the addition of some filtering mechanism to average the last several electoral results, as proposed earlier. A second way is to require any winner to have some minimum percentage of the vote. If we assume that the representation of a district should average out over, perhaps, twenty years, then a two-year term is 10% of that, meaning a candidate with less than 10% support can never win. With the one-month terms contemplated earlier, a twenty-year averaging time would dictate a minimum support threshold of .28%.

A fourth consideration is that it becomes possible for any official to be replaced at any electoral cycle. This effect becomes especially pronounced with higher-frequency elections. With one-month terms and stochastic elections a district could have twelve different representatives in a single year! Parliamentary procedure would require a complete redesign to handle such rapid changes, including some filtering mechanism to average the votes from a district over time. This would potentially be equivalent to a district casting fractional votes on any given issue.

Executive offices would be impractical to subject to such rapid changes; month-to-month shifts in executive domestic policies, staffing of appointed offices, or foreign policy would lead to chaos in short order. Therefore, a stochastic system cannot be directly applied to executive offices. However, indirect application is possible in systems where the executive is selected by the legislature, rather than directly by the people. The incentives for candidates to appeal to a maximal segment of the populace (as represented by the legislature) remain, while minimizing the probability of immediate replacement of the executive following every monthly election due to the averaging effect of the size of the legislature. There would remain a non-zero probability that a legislature dominated by a political minority could come into office by a statistical fluke, and thereby temporarily remove an otherwise broadly popular executive.

Finally, stochastic elections using plurality voting are naturally clone-independent, so long as each clone's popularity exceeds the minimum support threshold. Vote-splitting between two similar candidates only results in each of them having a lesser probability of winning, while in no way increasing the probability of a mutual opponent winning.

Many of these effects were previously noted in (Amar, 1984). Indeed without the minimum support threshold, stochastic elections devolve into lottery ballots, with every candidate having odds of winning exactly proportional to the number of votes received.

Simulations

Simulations were run to compare various electoral systems. Popular vote data was acquired for each US House district election result for 2014 and 2016. The margin of victory for the winning candidate was used to compute the candidate's vote percentage, on the simplifying assumption that only the Republican and Democratic parties ran candidates. Incumbents running unopposed were presumed to have no change in their vote percentage since their last opposed election. The vote percentages were interpolated and extrapolated on a month-by-month basis over the period from November 2014 to November 2018.

These data were used as input to a series of simulations of various election systems, including:

- Monthly deterministic elections with no filtering
- Monthly deterministic elections with hysteretic filtering
    - Integrated unpopularity must exceed 300% to trigger replacement
- Monthly deterministic elections with single-pole infinite impulse response (IIR) filtering
    - Filter pole at ten months
- Proportional representation based on total national popular vote, updated monthly
    - With the simplifying assumption that all districts are of equal population
- Monthly stochastic elections with no filtering
- Monthly stochastic elections with single-pole IIR filtering
    - Filter pole at ten months

The number of seats in the US House held by Republicans was calculated in each simulation. For stochastic simulations, 1,000 simulation runs were executed, and the average of those runs taken for each month of the simulation. The standard deviation among the 1,000 runs for each month was between nine and ten seats, regardless of simulation type. The results of these simulations are shown in Figure 6.

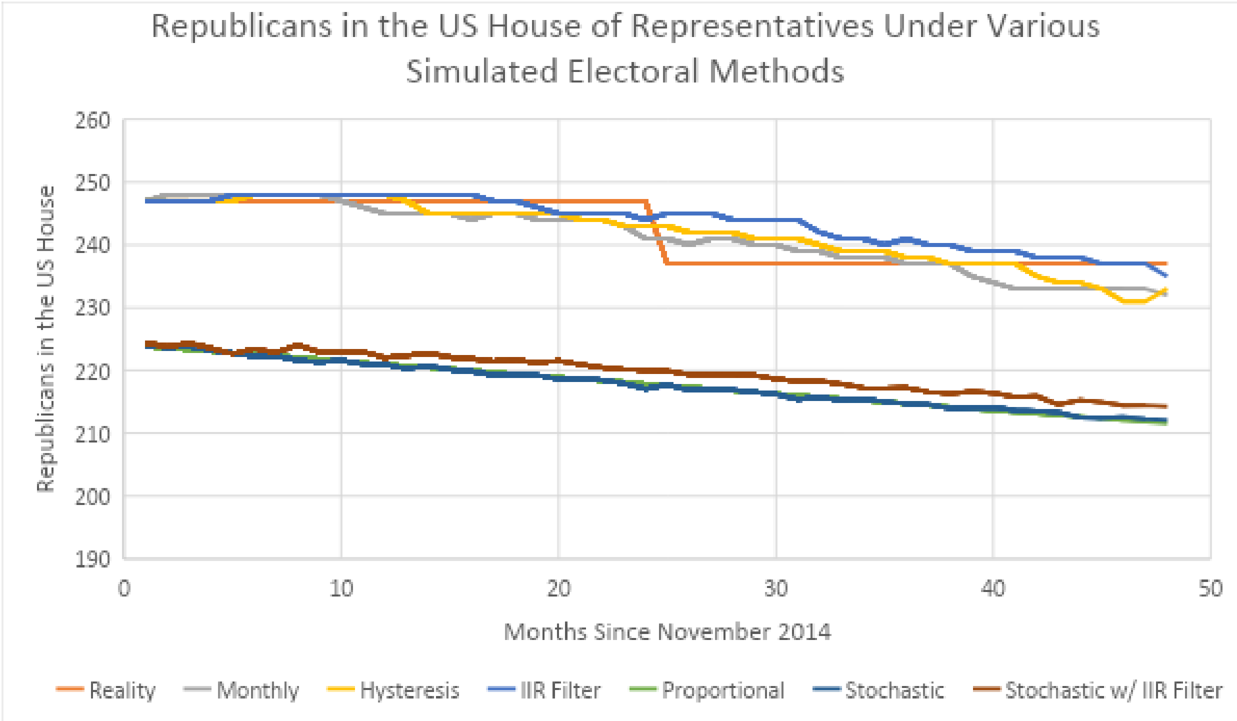

Figure 6 - Simulation Results

Simulation results indicated that the Republican party derives a substantial advantage over proportional allocation by the use of geographic districts, due to Democratic voters tending to be more densely concentrated, combined with gerrymandering of the district lines. Monthly elections with or without filtering did not mitigate this effect. Averaged stochastic elections, however, closely followed the theoretical proportional allocation curve, completely overcoming the effect of geographic districts. Adding a filter to either stochastic or deterministic elections did not change the overall trends; it simply delayed their translation from voter preference into outcomes, as intended.

<u>Conclusion</u>

Elections have much in common with digital signal processing systems. Both sample continuous-time analog inputs at a fixed rate and process them into a discrete-time quantized output. The problem of aliasing, caused by sampling the input at a rate too slow to capture rapid changes, exists in both domains; outputs can fail to represent changes in input, depending entirely on when the change came relative to the sample time.

In both domains aliasing is solved by an increased sampling rate. However, it is often not desirable to have the output of the electoral system, government actions, change as quickly as the input to the system, public opinion. This is why elected representatives have multi-year terms; it is a mathematically invalid solution to a very real problem.

A more correct solution would be to hold elections with higher frequency, and to insert a filter after the sampling takes place, causing the system to be biased towards stability and away from change. By requiring a supermajority to effect change, and by making the time necessary to execute a change inversely proportional to the level of the supermajority present, majority rule and the ability to make rapid changes are retained, while still damping the response of the system to smaller and more temporary shifts in public opinion.

Quantization error has a substantial effect on electoral outcomes, and a variant of dithering has the potential to correct those effects. Based on simulation results, stochastic elections will, if averaged over time, provide equivalent results to a deterministic proportional allocation system. This allows for proportionality in representation and for a degree of clone-independence, while still maintaining the existence of single-winner geographic districts. Gerrymandering ceases to have meaningful effect.

Many questions remain for future study. To use the Nyquist criterion to select an appropriate sampling rate, one must first determine whether there is a limit on how quickly public opinion can, in practice, change. Should it exist, this limit should be quantified. The detailed effects of having different sample times or threshold levels have not been fully treated here, nor have the increased costs to the voter and the state been effectively addressed. Other filters besides hysteresis and integration are worthy of consideration. Further research is needed on the implications of dithering in various voting systems and in elections involving more than two candidates. Work is also required to adapt parliamentary procedures to the possibility of rapid changes in membership.

References

Ahuja, S. (1994). Electoral Status and Representation in the United States Senate: Does Temporal Proximity to Election Matter? *American Politics Research*, 104-118.

Amacher, R. C., & Boyes, W. J. (1978). Cycles in senatorial voting behavior: implications for the optimal frequency of elections. *Public Choice*, 5-13.

Amar, A. R. (1984). Choosing Representatives by Lottery Voting. *The Yale Law Journal*, 1283-1308.

Barron, M. (2016, November). *Sparks.* Retrieved from A Better Way to Visualize Presidential Election Results: http://www.norc.org/Blog/pages/sparks-post.aspx?BlogId=89

Bó, E. D., & Rossi, M. A. (2011). Term Length and the Effort of Politicians. *The Review of Economic Studies*, 1237-1263.

Fischer, I., & Suleiman, R. (1997). Election Frequency and the Emergence of Cooperation in a Simulated Intergroup Conflict. *Journal of Conflict Resolution*, 483-508.

Gaines, B. J., Nokken, T. P., & Groebe, C. (2012). Is Four Twice as Nice as Two? A Natural Experiment on the Electoral Effects of Legislative Term Length. *State Politics & Policy Quarterly*, 43-57.

Schultz, C. (2008). Information, polarization and term length in democracy. *Journal of Public Economics*, 1078-1091.

Titiunik, R. (2016). Drawing Your Senator from a Jar:Term Length and Legislative Behavior. *Political Science Research and Methods*, 293-316.

Toussaint, D. (2007, September 7). *Dither.* Retrieved from Wikipedia: https://en.wikipedia.org/wiki/Dither

Tridimas, G. (2017). Constitutional choice in ancient Athens: the evolution of the frequency of decision making. *Constitutional Political Economy*, 209-230.